

AVERTISSEMENT

Ce document est le fruit d'un long travail approuvé par le jury de soutenance et mis à disposition de l'ensemble de la communauté universitaire élargie.

Il est soumis à la propriété intellectuelle de l'auteur : ceci implique une obligation de citation et de référencement lors de l'utilisation de ce document.

D'autre part, toute contrefaçon, plagiat, reproduction illicite de ce travail expose à des poursuites pénales.

Contact : portail-publi@ut-capitole.fr

LIENS

Code la Propriété Intellectuelle – Articles L. 122-4 et L. 335-1 à L. 335-10

Loi n° 92-597 du 1^{er} juillet 1992, publiée au *Journal Officiel* du 2 juillet 1992

<http://www.cfcopies.com/V2/leg/leg-droi.php>

<http://www.culture.gouv.fr/culture/infos-pratiques/droits/protection.htm>



THÈSE

En vue de l'obtention du
DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE
Délivré par l'Université Toulouse 1 Capitole

Présentée et soutenue par
Christopher SANDMANN

Le 29 juin 2021

essais en théorie économique

Ecole doctorale : **Toulouse Sciences Economiques**

Spécialité : **Sciences Economiques - Toulouse**

Unité de recherche :
UMR 5314 -TSE-R

Thèse dirigée par
Thomas MARIOTTI

Jury

M. Lucas MAESTRI, Rapporteur
Mme Leeat YARIV, Rapporteur
Mme Marina HALAC, Rapporteur
M. Bruno JULLIEN, Examineur
M. Thomas MARIOTTI, Directeur de thèse

Essays in Economic Theory

Christopher Sandmann

Directeur de thèse : Thomas Mariotti,
Toulouse School of Economics, Directeur de Recherche, CNRS

Membres du jury : Marina Halac, Bruno Jullien, Lucas Maestri, Leeat Yariv

Thèse
En vue de l'obtention du
Doctorat de l'Université de Toulouse
École doctorale : Sciences Economiques-Toulouse School of Economics

2021

Contents

1	When are posted-price mechanisms profit-maximizing?	3
1	Introduction	3
2	Second-degree price discrimination	6
2.1	Set-up	6
2.2	The optimality of posted-price mechanisms	6
2.3	Production Cost	11
3	Durable good monopoly	12
4	Conclusion	16
A	Random Mechanisms	17
2	Non-stationary Search and Assortative Matching	19
1	Introduction	19
2	The model	24
3	Equilibrium existence	30
4	Positive Assortative Matching	34
4.1	Definition of positive assortative matching	34
4.2	The mimicking argument	35
4.3	Stationary environment	36
4.4	Non-stationary environments	38
4.5	Necessity	41
5	Explicit search cost	42
6	Conclusion	45
A	Set-up: Proofs	46
A.1	Population Dynamics	46
A.2	The value of search	48
B	Existence of an equilibrium	50
B.1	Compactness of strategy subspaces	50
B.2	Continuity a.e. of the value function	52
B.3	The approximate fixed-point mapping	57
B.4	Convergence of the approximate fixed points	59
B.5	Limiting arguments	61
C	PAM: miscellaneous proofs	62
D	Necessity	64

3	Probabilistic Assortative Matching under Nash Bargaining	69
1	Introduction	69
2	The model	73
2.1	Set-up	73
3	Equilibrium existence and uniqueness	77
3.1	Definition of equilibrium	77
3.2	System of FBSDEs	77
3.3	Regularity conditions	78
3.4	Existence and uniqueness of equilibrium	80
4	Single-peaked preferences and single-crossing	80
4.1	The preference relation	80
4.2	Single-crossing	81
4.3	Single-peaked preferences	82
5	Probabilistic Assortative Matching	87
5.1	Definition of probabilistic assortative matching	87
5.2	Preferences for assortative matching	89
5.3	Main result	91
5.4	Revisiting Shimer and Smith (2000)	93
6	Conclusion	96
A	Introductory proofs	97
A.1	Remark 1	97
A.2	Proof of Lemma 2	97
A.3	Proof of Lemma 1	100
B	Uniqueness and existence: proof of theorem 1	100
B.1	Delarue's theorem	100
B.2	Transfer of notation	102
B.3	Regularity conditions and proof	103
C	Single-peaked preferences	106
C.1	Proof of Lemma 3	106
C.2	Proof of theorem 2	108
C.3	Observation: non single-peaked wages	111
D	Assortative matching: proofs	112
D.1	Illustrative results	112
D.2	Proof of the preference representation	113
D.3	Characterization of reciprocity via \mathcal{P}_t	114
D.4	Proof of theorem 4, non-stationary Shimer and Smith	115
D.5	Self-acceptance at the boundary: steady state	116
D.6	Counter-example	119
E	Discussion of Related Literature	120
E.1	Existence and uniqueness	120
E.2	The case for pair-specific production shocks	120

4	Recursive Information Design	123
1	Introduction	123
2	Information Structures	129
3	Belief Structures	129
4	Information Design in Supermodular Games	134
5	An Application to the Infection Argument	141
6	Conclusion	154
A	Appendix	155
	A.1 Higher-order Beliefs	155
	A.2 Characterization of believable belief structures	156

Acknowledgment

The journey which brought this thesis into being was shared. I am grateful to my advisor, Thomas Mariotti. His work became my first compass, the reassurance that the nitty-gritty holds the most rewarding and deep insights. Thomas' acute perception of economic problems and his aesthetic determination to resolve them has been the example to which I sought to live up to. I wish to express my gratitude for our conversations, his encouragement, trust, patience, and dedication to me; finally, his invaluable support during the job market, through which he steered me without me noticing.

I am also grateful to my good friend and co-author Nicolas Bonneton. Half of this thesis was written and envisioned jointly across three continents, simultaneously, to the point of exasperation. He is like the twin who looks left where I look right. Neither of us would have gotten nearly as far without the other. To be able to work so closely together is very special to me.

In a fortunate turn of events, Lucas Maestri saw to it that I could come visit him at Fundação Getulio Vargas in 2018 and early 2019. FGV was the perfect environment to think carefully about what turned out to be a months-long endeavor to prove existence of a non-stationary search and matching equilibrium. Lucas provided thoughtful and very valuable comments, Humberto Moreira let me see him weekly to discuss the latest progress. Paulo Monteiro Klinger was willing to assess the entire proof upon its first completion. The small but active internal seminars, also joined by Carlos da Costa, provided just the right amount of distraction and motivation. I am very grateful for their support.

I also wish to express my gratitude to Johannes Hörner who was instrumental to shape what was to become my job market paper. He pointed out the lack of proof of existence and motivated its resolution. Johannes was of tremendous help once more a year later when he agreed to meet me at Yale in a short window of time despite being jetlagged and having suffered a terrible bike accident earlier. I greatly benefited from his technical remarks that day and cannot thank him enough for his invaluable support during the 2019/2020 job market.

I have benefited enormously from the intellectual environment in Toulouse. I thank Andrea Attar, Jacques Crémer, Daniel Garrett, Bertrand Gobillard, Renato Gomes, Christian Hellwig, Bruno Jullien, Francois Salanié, Patrick Rey, Andrew Rhodes, Anna Sanktjohanser, Jean Tirole and Takuro Yamashita. Their questions, ideas and thoughtfulness during our meetings but also experienced during seminars held the most valuable lessons for me. I thank Patrick Rey for entrusting me the role as teaching assistant for his micro PhD course on incentives. I am also thankful to the theory students in Toulouse; especially to Alae Baha (who provided a key idea and impetus about random contracts) and Dakang Huang (who has taught me so much about nonexclusive competition).

I thank TSE chairman Jean Tirole and TSE director Christian Gollier. I thank Patrick Fève and David Alary for opening me the door, PhD directors Wilfried Sand-Zantman for his benevolence and Nour Meddahi for his extensive support and advice. I thank Alexander Gümbel, aided by Daniel Ershov, for the countless hours spent on easing our way on the job market; I thank the administrative staff, Laurence Delorme, Aude Schlösing without whom Jacopo and I could never have rented a flat, Caroline Tejedor and Ludmila Namolovan who skillfully kept all administrative nuisances at bay. I thank Raja who let me play the violin in the cafeteria in the

afternoons, and the whole team of CROUS which feeds the students so well.

I also wish to thank my friends at TSE, the bond and lived solidarity among the PhD students which make this place so very special; among them my fantastic PhD cohort Filippo, Hussein, Jacint, Jacopo, Jérôme, Joana, Julia, Kunal, Ling, Luis, Miguel, Miren, Nicolas, Olga, Rodrigo, Stefan, Willy, Yang Yang and Yuting; generations of TSEconomist writers, and the most intense yet gratifying year 2014/15 which Julia, Marica, Joana, Jacint, Jacopo and I enjoyed at its helm; the harmonious and happy time shared with my flatmate and friend Jacopo, with Flaminia and Matteo; and thank you, Alípio, for the music, for the words, for our precious friendship.

Lastly, I wish to say thank you to my family. To my parents Beate and Klaus, the first economists I have ever known, who have always been there for me. If I can feel at home anywhere, I owe it to them. To my brother and sister, Joachim and Theresa, foe and friend in all early adventures. And to my wife, Meg, with whom I share the greatest of all adventures; she is the mother to our two children, Clara and newborn Thomas. I will be forever grateful for our love and the courage that it gave us.

Introduction

This dissertation studies several seemingly disparate problems in economic theory: chapters 2 and 3 (written jointly with Nicolas Bonneton) investigate how (non-stationary, i.e., time-varying) search frictions affect positive assortative matching (PAM), chapter 4 studies higher-order belief uncertainty and their manipulation in games of incomplete information, chapter 1 questions the optimality of second-order price discrimination by a monopolist seller.

Economists view choices—taking the form of individual match acceptance decisions (as in chapter 2 and 3), coordination issues in investment games (as in chapter 4), or consumer purchase and subsequent pricing decisions (as in chapter 1)—as solutions to constrained optimization problems. Comparative properties of payoffs determine whether one agent is more accepting to match with a less desirable partner than another agent, whether investment decisions are mutually encouraging, or to what extent one can decrease quality and price to buyers of lower valuation without sacrificing the demand of high valuation buyers for high quality products.

What binds these disparate essays together is that they identify (as in chapters 1, 2 and 3) or make use of (as in chapter 4) comparative properties of payoffs under which equilibrium behavior is particularly well understood: sorting under random search matching obtains (see chapters 2 and 3); we gain insights into the information structure which guarantees the highest probability of coordination across all equilibria (chapter 4); instead of charging distinct prices for separate qualities or times of purchase, the monopolist posts a single price for the highest level of quality only (chapter 1).

In the context of monopolistic second-degree price discrimination (see chapter 1) I derive a condition under which a single posted price as opposed to sophisticated discriminatory pricing is profit-maximizing. The weakest sufficient condition for this to arise posits a ranking on the concavity of preferences. Under second-degree price discrimination this requires that high valuation buyers' utilities are more concave than low valuation buyers' utilities. Intuitively, this is the case when consuming per-se, not quality is what matters to buyers of greater valuation. This generalizes findings by Anderson and Dana (2009) and Stokey (1979).

In the context of matching absent bargaining (the NTU paradigm), see chapter 2, we identify comparative properties of payoffs under which more highly ranked individuals are choosier with whom they accept to match, thereby stipulating PAM—even in the presence of search frictions. Answers to this question have been provided in the steady state where entry and exit exactly balance each other at every point in time. Our analysis encompasses non-stationary dynamics instead, where the pool of agents searching fluctuates over time. To sustain sorting we find that

more desirable individuals must be less risk-averse than lesser ranked individuals. The intuition is straightforward: the decision to reject a certain match payoff today is a revealed preference for a risky, random match payoff some time in the future. What is then perhaps most surprising is that this ranking on risk-preferences is not required to ensure PAM in the steady state (see ? and ?). The reason is that risk-preferences only come to full bearing in the presence of the risk of future worsening match prospects, inherent to non-stationary dynamics.

In chapter 3, we extend our analysis of non-stationary sorting to an environment where match payoffs are determined via bargaining (the TU paradigm, previously studied by Shimer and Smith (2000) and Atakan (2006)). More fundamentally, chapter 3 broadens the definition of assortative matching to allow for individual heterogeneity. In doing so, chapter 3 identifies a novel force which upsets assortative matching under bargaining – irrespective of match complementarities: frictions disproportionately erode the value of search and hence the bargaining power of more productive agents. At the extreme, when the market is thin so that meetings are rare, agents unanimously exhibit the greatest match surplus when matched with the most productive individuals, and not with individuals of similar characteristics.

In addition, the thesis makes several technical contributions. First, chapter 2 formulates general conditions under which deterministic non-stationary dynamic equilibria of the search and matching economy exist. Here we considered a continuum of agents and infinite time horizon. Motivated by the fact that finite player games are known to be prone to equilibrium multiplicity, i.e., they can coordinate on equilibria in which all agents are more or less accepting of others, chapter 3 then shows that in a framework with aggregate uncertainty and terminal-time, an equilibrium does not only exist but is also unique. The reason is that aggregate uncertainty makes it impossible to coordinate on an equilibrium, when such strategy can only be optimal if one reaches a far-fetched pessimistic (or optimistic) state with few (many) agents in the search pool. More succinctly, aggregate uncertainty impedes coordination.

Finally, chapter 4 proposes a novel approach to information design in games (going beyond the single-receiver case studied by Kamenica and Gentzkow (2011)). An example thereof is the investment game where a sender seeks to encourage several investors to invest amid an unknown state of nature. Central to this novel approach is a new explicit representation of higher-order beliefs which allows for greater tractability than the representation proposed thus far in the literature. As an application, I provide a unified and hopefully clarifying perspective on the oft-invoked infection argument. Here, the sender sends a chain of messages alternating between players such that one message inducing a low action is contagious for another player to likewise choose the low action. I introduce a new distinction based on whether the information structure manipulates first-order beliefs only or creates uncertainty over higher-order beliefs at all orders.

When are posted-price mechanisms profit-maximizing?

1 Introduction

Second-degree price discrimination has long been known as a business strategy to increase profits. Instead of offering a single good for a posted-price (the posted-price mechanism), the seller recognizes that the heterogeneous demand she faces can be divided into distinct segments. Buyers from distinct segments purchase different quantities or qualities according to their taste and willingness to pay. Motivated by information asymmetries between the seller and the buyer, no offer is made contingent on the buyer's identity. Rather, the seller's offer takes the form of a menu. In practice, menus may be of manifold form, such as product line pricing, damaged goods, and intertemporal pricing (i.e., the durable good monopoly).

The theoretical analysis of second-degree price discrimination constitutes one of the most successful applications of the principal-agent model (see for instance Laffont and Martimort (2002) and the original work by Mussa and Rosen (1978) and Maskin and Riley (1984)). A key tenet of the theory is the concept of downward distortions whereby goods provided to buyers of lesser valuation are of exceedingly lower value than those provided in the first-best, i.e., were one to maximize the buyer's utility minus the cost. This derives from the fact that high-valuation buyers will find the option of buying products of lesser value quite tempting, whereas buyers of lesser valuation will find it unattractive to spend more for goods of greater value (as formally expressed through downward-binding incentive constraints). The seller's profit-maximizing pricing rule must resolve a delicate balancing act: by providing products of greater value to buyers of lesser valuation, one also increases the outside option value of high valuation buyers. In effect, an increase in price and value of the good provided to low valuation buyers must go hand in hand with a decrease in price to high valuation buyers. The way the seller resolves these trade-offs at the margin, the literature has shown, hinges on the distribution of buyer valuations.

In this paper I provide a new perspective on the trade-offs inherent to price discrimination. The environment is standard. I operate under the common assumption that preferences are quasi-linear and satisfy single-crossing. A monopolist seller devises a mechanism as to maximize the sum of transfers received minus the cost. Where I differ is the approach: Rather than taking a marginal perspective —by how much ought one increase a given buyer type's provision of value as they become present at greater frequency —I ask the categorical question that

should take precedent: why discriminate in between buyers of lesser or greater valuation in the first place? That price discrimination can increase monopoly profits does not mean that it must always be the case. The answer to this question will be qualitatively different from the analysis of marginal trade-offs. Intuition can be gained by following the preliminary analysis where individual production cost are zero; the seller can provide any and arbitrarily many damaged goods of an initial endowment. Then consider a buyer of intermediate valuation. Suppose he were neither to choose the good of greatest nor least value. The monopolist's trade-off is two-sided: for once, she can increase intermediate value provision, which forces her to lower the price of the greatest value good. She will do so if agents of intermediate type are numerous. Or, she can decrease intermediate value provision, which will allow her to raise the price of the greatest value good. This will be advantageous to her if there are few agents of intermediate type. By studying these trade-off jointly, I am able to derive a condition concerning buyers' preferences under which the resolution of these trade-offs takes a particularly simple form: the seller either increases or decreases the intermediate value provision up to the good of greatest or down to the good of least value. Notably, this result holds irrespective of the distribution over buyer valuations.

The main contribution of this article is to show that if a ranking on the concavity of consumer types' utilities holds, then posted-price mechanisms are profit-maximizing. The ranking on the concavity of buyer types' utilities is most convincing when production cost $c(a)$ per individual purchase a are zero (as is for instance the case when the good in question is software or information). Formally, I require that for any two types $\theta_i < \theta_j$ and allotments $a_1 > a_2 > a_3$ one has

$$\frac{v(a_3; \theta_i) - v(a_2; \theta_i)}{v(a_2; \theta_i) - v(a_1; \theta_i)} \geq \frac{v(a_3; \theta_j) - v(a_2; \theta_j)}{v(a_2; \theta_j) - v(a_1; \theta_j)},$$

where $v(a; \theta)$ denotes consumer type θ 's utility of consuming allotment a . This means that those buyers who are more eager to trade report a consumption utility which is more concave in the value of the good provided. On a more intuitive level, this is the case when consuming per-se, not the value of the good is what matters to buyers of greater valuation. On a formal level, assuming differentiability, the condition is equivalent to log submodularity of $\partial_a v(a; \theta)$.¹ A simple example of preferences satisfying said condition is given by $v(a; \theta) = \frac{a^{1-\theta}}{1-\theta}$ with $a \in [0, 1]$ and $\theta \in (0, 1)$. More generally, observe how this ranking on concavity arises naturally if one imposes that consumption utility is concave and satisfies single-crossing.

I then provide the following result under zero individual production cost: *if preferences satisfy log submodularity in differences and individual production cost are zero, then a posted-price mechanism is profit-maximizing.* I also show, as one might have anticipated from Strausz (2006), that this result is robust if one allows for random mechanisms. Finally, a special case of this result is already known and due to Stokey (1979): when consumption utilities take a separable form $v(a; \theta) = \theta w(a)$ for any (increasing) function $w(a)$, then a posted-price mechanism

¹Bonneton and Sandmann (2019) introduce the terminology of log supermodularity in differences, when studying sorting in the unrelated case of random search matching.

will always be profit maximizing.²

As a second result, corollary to the first, I report that the sufficient condition to render posted-price mechanisms profit-maximizing becomes more stringent when production cost $c(a)$ per individual purchase a are non-zero. For instance, given the utilities $v(a; \theta) = \frac{a^{1-\theta}}{1-\theta}$, posted-price mechanisms will no longer be optimal when individual production cost are non-zero and linear. This is an important economic finding in its own right: the range of goods with zero production cost per individual purchase (also known as non-rival goods) encompasses many goods subsumed in the information and communication technology sector—which accounts for a growing share of output and trade surplus in advanced economies. The theory thus predicts that if second-degree price discrimination were to maintain an important role in those sectors, it is for reasons different from static profit-maximization.

Formally, the analysis of non-zero individual production cost follows from the former. To see this, decompose the price paid for consumption level a into production cost and a markup. Thus define cost-augmented buyer utilities and transfers as $\hat{v}(a; \theta) = v(a; \theta) - c(a)$ and $\hat{t}(a) = t(a) + c(a)$. Under this notation, the seller (as before) ignores individual cost and seeks to maximize the sum of (auxiliary) transfers $\hat{t}(a)$ received, whereas the buyer's utility is given by $\hat{v}(a; \theta) - \hat{t}(a)$. I once more report a sharp condition: *if cost-augmented preferences satisfy log submodularity in differences, then a posted-price mechanism is profit-maximizing.*

Finally, it is important to mention related work by Anderson and Dana (2009). Studying the same question, they identify a related ranking on the curvature of preferences under which posted-price mechanisms are profit-maximizing. However, their result only holds when imposing as a distributional assumption the hazard rate property; in contrast, the herein presented result obtains irrespective of the distribution over types. The main body of the text discusses the connection in greater detail.

The remainder of the chapter is organized as follows: section two pursues the analysis of second-degree price discrimination. Since the notions of quantity and quality are mathematically equivalent, they do not receive separate treatment here. The analysis maintains the assumption that individual production cost are zero, before expanding and showing how the analysis can be mapped into the general case with non-zero individual production cost. In section three I briefly study the durable good monopoly which is shown to be a mirror case of the former: the seller no longer discriminates in quantity but timing of purchase. The analysis is largely equivalent. In fact, both sections can be read independently, and it is up to the reader to decide which point of view they find more intriguing. Lastly, section four concludes. A treatment of random mechanisms is deferred to the appendix.

²Stokey's original result was derived studying the durable good monopoly. It informed the conventional wisdom that posted-price mechanisms are always profit-maximizing. Instead, the binding inequality reveals that multiplicatively separable preferences are the knife-edge case. With the wrong kind of curvature, decreasing prices over time can be profit-maximizing, even when there is commitment.

2 Second-degree price discrimination

There is a finite number of consumers, higher types of which distinguish themselves by higher valuations. A monopolistic seller offers a price-quantity menu, with each item detailing an allotment as well as a transfer. The initial analysis departs from the standard framework in that I focus on zero individual production cost: the monopolistic seller holds a unit endowment of the good and may choose as an allotment any quantity-discounted version of her endowment. I will subsequently relax this assumption.

2.1 Set-up

Consider a single seller and a finite number of vertically differentiated buyer types $\{\theta_1, \dots, \theta_N\}$. The seller holds an exogenously given endowment 1. Given her endowment, the seller devises a menu of trades, taking the form of a direct mechanism. Each trade proposes an allotment $a(\theta_j) \in \mathbb{R}_+$ and a transfer $t(\theta_j) \in \mathbb{R}_+$. None of the allotments may exceed the endowment 1, but arbitrarily many allotments can be sold at no cost. Therefore, $\sup_j a(\theta_j) \leq 1$. This will be referred to as the feasibility constraint. Buyers are vertically differentiated with higher values of $\theta_{j+1} > \theta_j$ representing buyer types of higher valuation. A buyer's preferences consuming trade $(a(\theta_j), t(\theta_j))$ are represented by quasi-linear utility

$$v(a(\theta_j); \theta_j) - t(\theta_j).$$

Throughout I maintain the standard assumption that the buyer's preferences satisfy single-crossing.

Assumption 1 (single-crossing). *Preferences satisfy single-crossing, that is for any allotments $a_2 > a_1$ and types $\theta_j > \theta_i$*

$$v(a_2; \theta_j) - v(a_1; \theta_j) \geq v(a_2; \theta_i) - v(a_1; \theta_i).$$

Lastly, the mass of type θ_j is given by $m(\theta_j)$ which translates cumulatively into $M(\theta_n) = \sum_{i=1}^n m(\theta_i)$. No assumptions regarding the distribution of types are made other than that $m(\theta_j) > 0$ for all θ_j . In standard mechanism design the seller gradually increases the quality provided to higher types as to most profitably extract each consumer type's valuation. I now provide conditions under which this will not be profit-maximizing.

2.2 The optimality of posted-price mechanisms

The following provides a characterization as to when a posted-price mechanism as opposed to menu-pricing is optimal. I begin with a definition.

Definition 1 (log submodularity in differences). *Preferences satisfy log submodularity in dif-*

ferences if for all $a_3 > a_2 > a_1$ and $\theta_2 > \theta_1$

$$\frac{v(a_3; \theta_2) - v(a_2; \theta_2)}{v(a_2; \theta_2) - v(a_1; \theta_2)} \leq \frac{v(a_3; \theta_1) - v(a_2; \theta_1)}{v(a_2; \theta_1) - v(a_1; \theta_1)}.$$

As shown in Chade and Swinkels (2019) and Bonneton and Sandmann (2019) this condition is equivalent to saying that $\partial_a v(a; \theta)$ is log submodular (refer to proposition 9 in chapter two of this thesis). Log submodularity in differences stipulates that a higher type's consumption preferences $v(a; \theta_j)$ are a concave transformation of lower types' consumption preferences, i.e. are more concave. Log submodularity in differences also affords an immediate interpretation in terms of indifference curves. Indifference curves are described by an implicitly defined function $t(a; \theta_j)$ where $v(a; \theta_j) - t(a; \theta_j)$ is equal to some fixed constant for all a and θ_j . Then in the (a, t) -plane log submodularity in differences stipulates that higher types' indifference curves are concave transformations of lower types' indifference curve. Lastly, a consumer's relevant type θ_j is summarized by its inverse demand $p = D(a; \theta_j)$, where a and p are related by $\max_a v(a; \theta_j) - a p$. Supposing that $a \mapsto v(a; \theta_j)$ is quasi-concave, it follows that $D(a; \theta_j) = \partial_a v(a; \theta_j)$. Thus log submodularity in differences of the consumer's preferences implies that inverse demand curves must be log submodular in consumption level and type.

Proposition 1. *Suppose that preferences satisfy log submodularity in differences. Then the monopolist's optimal mechanism is a posted-price mechanism.*

Proof. Take as given an endowment e and a profit-maximizing direct mechanism $(a(\theta_j), t(\theta_j))_{j=1}^N$ selling the endowment. Let $u(\theta_i) = \max_{j \in \{1, \dots, N\}} v(a(\theta_j); \theta_i) - t(\theta_j)$ buyer type θ_i 's indirect utility. Clearly, incentive constraints are downwards binding (see Rochet (1987)). Then $u(\theta_{i+1}) - u(\theta_i) = v(a(\theta_i); \theta_{i+1}) - v(a(\theta_i); \theta_i)$. Or, using the fact that participation constraints are binding for the lowest type, $u(\theta_1) = 0$, this reads as

$$u(\theta_i) = \sum_{j=1}^{i-1} [v(a(\theta_j); \theta_{j+1}) - v(a(\theta_j); \theta_j)].$$

Taking the weighted sum of utilities gives

$$\begin{aligned} \sum_{i=1}^N u(\theta_i) m(\theta_i) &= \sum_{i=1}^N \sum_{j=1}^{i-1} [v(a(\theta_j); \theta_{j+1}) - v(a(\theta_j); \theta_j)] m(\theta_i) \\ &= \sum_{i=1}^N [v(a(\theta_i); \theta_{i+1}) - v(a(\theta_i); \theta_i)] (1 - M(\theta_i)). \end{aligned} \quad (1)$$

This allows to formulate the monopolist seller's objective as follows

$$\begin{aligned} \sum_{i=1}^N t(\theta_i) m(\theta_i) &= \sum_{i=1}^N [v(a(\theta_i); \theta_i) - u(\theta_i)] m(\theta_i) \\ &= \sum_{i=1}^N [v(a(\theta_i); \theta_i) - v(a(\theta_{i-1}); \theta_i)] (1 - M(\theta_{i-1})). \end{aligned} \quad (2)$$

The seller chooses optimal quantities $0 \equiv a(\theta_0) \leq a(\theta_1) \leq \dots \leq a(\theta_N) = e$. If menu-pricing were to yield a strictly greater profit than a posted-price mechanism there exist types θ_i, θ_j with $j > i$ such that $a(\theta_j) > a(\theta_{j-1}) = a(\theta_i) > a(\theta_{i-1}) \geq 0$. The chosen mechanism must yield a greater profit than the two deviations where the seller either provides $a(\theta_j)$ or $a(\theta_i)$ to all types $\{\theta_i, \dots, \theta_{j-1}\}$. Or,

$$\begin{aligned} & [v(a(\theta_i); \theta_i) - v(a(\theta_{i-1}); \theta_i)](1 - M(\theta_{i-1})) + [v(a(\theta_j); \theta_j) - v(a(\theta_{j-1}); \theta_j)](1 - M(\theta_{j-1})) \\ & > \max \left\{ [v(a(\theta_j); \theta_i) - v(a(\theta_{i-1}); \theta_i)](1 - M(\theta_{i-1})); [v(a(\theta_j); \theta_j) - v(a(\theta_{i-1}); \theta_j)](1 - M(\theta_{j-1})) \right\}. \end{aligned}$$

This is equivalent to

$$\begin{aligned} & [v(a(\theta_j); \theta_i) - v(a(\theta_i); \theta_i)](1 - M(\theta_{i-1})) < [v(a(\theta_j); \theta_j) - v(a(\theta_{j-1}); \theta_j)](1 - M(\theta_{j-1})) \\ & [v(a(\theta_{j-1}); \theta_j) - v(a(\theta_{i-1}); \theta_j)](1 - M(\theta_{j-1})) < [v(a(\theta_i); \theta_i) - v(a(\theta_{i-1}); \theta_i)](1 - M(\theta_{i-1})). \end{aligned}$$

Dividing one by another and noting that by construction $a(\theta_{j-1}) = a(\theta_i)$ one obtains

$$\frac{v(a(\theta_i); \theta_i) - v(a(\theta_{i-1}); \theta_i)}{v(a(\theta_j); \theta_i) - v(a(\theta_i); \theta_i)} > \frac{v(a(\theta_i); \theta_j) - v(a(\theta_{i-1}); \theta_j)}{v(a(\theta_j); \theta_j) - v(a(\theta_i); \theta_j)}.$$

If preferences satisfy log submodularity in differences this can never be the case, establishing the desired contradiction. \square

Proposition 1 can be further strengthened. First, another definition is in place.³

Definition 2 (log submodularity). *Preferences satisfy log submodularity if for all $a_2 > a_1$ and $\theta_2 > \theta_1$*

$$\frac{v(a_2; \theta_2)}{v(a_1; \theta_1)} \leq \frac{v(a_2; \theta_1)}{v(a_1; \theta_2)}.$$

Log submodularity and log submodularity in differences are quite different from one another. Log submodularity in differences, or equivalently $\partial_a v(a; \theta)$ being log submodular, is a property regarding the curvature of consumption utility only. In contrast, log submodularity is a property which simultaneously affects the curvature and the level of consumption utility. In the special case where all agent types derive zero utility from zero consumption, i.e. $v(0; \theta) = 0$ for all θ , one easily observes that log submodularity in differences implies log submodularity. (The converse does not hold as illustrated in example 1 though such examples are tedious to construct.) More generally, consider the case where $v(0; \theta) \geq 0$ is non-increasing in type θ . Then proposition 1 can be strengthened as follows.

Corollary 1. *Suppose that preferences satisfy log submodularity. Suppose further that zero consumption values are non-negative and non-increasing, i.e. for all $\theta_i < \theta_j$ it follows that $v(a; \theta_i) \geq v(a; \theta_j) \geq 0$. Then the monopolist's optimal mechanism is a posted-price mechanism.*

Proof. If menu-pricing were to yield a strictly greater profit than a posted-price mechanism, there exist types θ_i, θ_j with $j > i$ such that $a(\theta_j) > a(\theta_{j-1}) = a(\theta_i)$. Without loss of generality let

³Equivalently, $v(a; \theta)$ is log submodular if $\frac{\partial_a v(a; \theta)}{v(a; \theta)}$ is increasing in θ .

θ_i the smallest such type. Then $a(\theta_{i-1}) \in \{0, a(\theta_i)\}$. As in the proof of the main result, it follows that

$$\frac{v(a(\theta_i); \theta_i) - v(a(\theta_{i-1}); \theta_i)}{v(a(\theta_j); \theta_i) - v(a(\theta_i); \theta_i)} > \frac{v(a(\theta_i); \theta_j) - v(a(\theta_{i-1}); \theta_j)}{v(a(\theta_j); \theta_j) - v(a(\theta_i); \theta_j)}.$$

If $a(\theta_{i-1}) = a(\theta_i)$ this inequality can impossibly hold. If $a(\theta_{i-1}) = 0$ substituting $v(0; \theta_i) \geq v(0; \theta_j)$ into the inequality gives

$$\frac{v(a(\theta_i); \theta_i) - v(0; \theta_j)}{v(a(\theta_j); \theta_i) - v(a(\theta_i); \theta_i)} > \frac{v(a(\theta_i); \theta_j) - v(0; \theta_j)}{v(a(\theta_j); \theta_j) - v(a(\theta_i); \theta_j)}.$$

Or, rearranging terms yields

$$\begin{aligned} v(a(\theta_i); \theta_i)v(a(\theta_j); \theta_j) - v(a(\theta_j); \theta_i)v(a(\theta_i); \theta_j) \\ > [v(a(\theta_j); \theta_j) - v(a(\theta_i); \theta_j) - v(a(\theta_j); \theta_i) + v(a(\theta_i); \theta_i)]v(0; \theta_j) \geq 0. \end{aligned}$$

Here the latter inequality stems from the assumption that $v(0; \theta_j) \geq 0$ and single-crossing. This implies that preferences are not log submodular for $\theta_j > \theta_i$ and $a(\theta_j) > a(\theta_i)$, establishing the desired contradiction. \square

In particular, this establishes that if $v(0; \theta) = 0$ for all θ as one typically assumes, log submodularity is sufficient to guarantee the optimality of posted-price mechanisms. This result is not the first to relate log submodularity of preferences to the absence of menu-pricing. Anderson and Dana (2009), proposition 3, derive the same property in the context of rival goods using standard techniques from mechanism design. What follows is a discussion of how corollary 1 strengthens the initial result by Anderson and Dana (2009). The busy reader is invited to skip this part and proceed with the optimality of posted-price mechanisms

Virtual surplus

As is standard in mechanism design, Anderson and Dana rely on the notion of virtual surplus, $\Lambda(a(\theta_i); \theta)$. Adapted to discrete types and non-rival goods this reads as

$$\Lambda(a(\theta_i); \theta_i) = v(a(\theta_i); \theta_i) - [v(a(\theta_i); \theta_{i+1}) - v(a(\theta_i); \theta_i)] \frac{1 - M(\theta_i)}{m(\theta_i)}.$$

A rewriting of equation (1) gives

$$\sum_{i=1}^N t(\theta_i)m(\theta_i) = \sum_{i=1}^N \Lambda(a(\theta_i); \theta_i)m(\theta_i).$$

Thus maximizing the weighted sum of transfers is equivalent to maximizing the weighted sum of virtual surplus. Under appropriate assumptions on Λ , i.e. quasi-concavity in a and supermodularity, maximizing the weighted sum is equivalent to maximizing $\Lambda(a; \theta_i)$ pointwise for each θ_i . The most general version of Anderson and Dana's proposition, adapted to non-rival

goods, makes use of those assumptions and can be stated as follows⁴: *Suppose that virtual surplus is quasi-concave in a and supermodular. Suppose further that preferences satisfy log submodularity. Then the monopolist's optimal pricing mechanism takes the form of a posted-price mechanism.* It becomes clear from here why corollary 1 is an appealing strengthening of proposition 3 by Anderson and Dana (2009). It remains true regardless of the distribution of buyer types and does not make use of technical assumption that are tedious to check. Indeed, one can construct examples (see example 2) where preferences and distribution functions are well-behaved, exhibiting fairly standard properties, yet virtual surplus is not quasi-concave. Here corollary 1 allows one to conclude that a posted-price mechanism is profit-maximizing even when the first-order approach maximizing virtual surplus is inconclusive.

Example 1 (Log submodular but not log submodular in differences). *I provide here an example of a utility $v(a; \theta)$ exhibiting standard properties ($v(a; \theta)$ is increasing and concave in a and exhibits single-crossing) that is log submodular, yet does not satisfy log submodularity in differences. Begin with the latter condition: if $v(a; \theta)$ is log submodular in differences then equivalently (see Bonneton and Sandmann (2019)) $\partial_a v(a; \theta)$ is log submodular. Or, let*

$$\frac{\partial_a^2 v(a; \theta)}{\partial_a v(a; \theta)} = \theta \left[\frac{2}{3} a \theta^2 - \frac{1}{2} \theta \right]$$

for $a \in [0, 1]$ and the extension of θ to the continuum $[\frac{1}{3}, \frac{2}{3}]$. Observe that this ratio is not monotone in θ . It follows that $v(a; \theta)$ is not log submodular in differences for a and θ sufficiently large. Integrating, and imposing that $v(0; \theta) = 0$ for all θ one obtains

$$v(a; \theta) = A(\theta) \int_0^a \exp \left\{ \frac{2x^2 \theta^2 - 3x\theta}{6} \right\} dx.$$

Then set $A(\theta) = -\frac{1}{3} + \theta$. I can verify that $v(a; \theta)$ is concave in a , exhibits single-crossing and, in particular, satisfies log submodularity.

Example 2 (Posted-price mechanism with virtual surplus that is not quasi-concave). *The utility from example 1 is zero for zero consumption, i.e. $v(0; \theta) = 0$ for all θ and log submodular. Thus corollary 1 allows one to conclude that no matter what the distribution of types a monopolist will always find a posted-price mechanism profit-maximizing. In the meantime one need not be able to derive such conclusion from pointwise maximizing virtual surplus. Indeed, consider the uniform distribution on $[\frac{1}{3}, \frac{2}{3}]$. As the number of discrete types grows large, virtual surplus converges to*

$$\Lambda(a; \theta) = v(a; \theta) - \partial_\theta v(a; \theta) \frac{1 - F(\theta)}{f(\theta)}$$

where $F(\theta) = 3\theta - 1$ is the continuous uniform distribution function (with $F(\theta) = M(\theta)$ in the

⁴Anderson and Dana (2009) posit that the distribution over types admits the non-increasing hazard rate property and assume that signs of first-, second- and third-order partial derivatives of utility are such that virtual surplus is concave, thus in particular quasi-concave, and supermodular.

support of M) and f its derivative. Or,

$$\Lambda(a; \theta) = \int_0^a \left\{ \left(\theta - \frac{1}{3} \right) - \left[1 + \left(\theta - \frac{1}{3} \right) (x^2 \theta^2 x \theta) \right] \frac{2 - 3\theta}{3} \right\} \exp \left\{ \frac{2x^2 \theta^3 - 3x\theta^2}{6} \right\} dx.$$

Simple inspection of virtual surplus at $\hat{\theta} = 0.498$ shows that virtual surplus is neither concave nor monotone.




Figure 1: Virtual surplus that is not quasi-concave with an interior minimum: the first-order-approach is inadmissible. This, despite the fact that preferences are increasing, concave, exhibit single-crossing and the distribution admits a non-increasing hazard rate. In contrast, corollary 1 establishes that the profit-maximizing mechanism takes the form of a posted-price mechanism.

2.3 Production Cost

So far, I had focused on zero individual production cost: the monopolistic seller held a unit endowment of the good (where production of the endowment may be costly) and chose as an allotment any quantity-discounted version of her endowment. I will now relax this assumption and adopt the standard assumption that cost are proportional to the mass of buyers served. I will show that a simple accounting argument will extend the preceding analysis to this environment.

Let $c(a)$ the cost of serving allotment a . The seller's profits for given a direct, incentive compatible mechanism $\{(a(\theta_i), t(\theta_i))\}_{i=1}^N$ is given by the sum of transfers minus the cost:

$$\sum_{i=1}^N [t(\theta_i) - c(a(\theta_i))]m(\theta_i).$$

In effect, transfers received must first recoup production cost before the seller can earn a profit. Accordingly, define mark-ups for a given mechanism as $\hat{t}(\theta_i) \equiv t(\theta_i) - c(a(\theta_i))$. From the buyer's point of view, it makes no difference to include the cost born by the seller in his consumption utility, if those cost are subsequently deduced from his transfer. Thus define buyer type θ_j 's cost-augmented utility $\hat{v}(a; \theta_j) \equiv v(a; \theta_j) + c(a)$. Clearly, by virtue of mere accounting, buyer

type θ_i 's reporting problem is equally well represented by

$$\max_{j \in \{1, \dots, N\}} \hat{v}(a(\theta_j); \theta_i) - \hat{t}(\theta_j),$$

whereas the seller's profits are equivalently given by

$$\sum_{i=1}^N \hat{t}(\theta_i) m(\theta_i).$$

This shows that the seller's problem with individual production cost is embedded in the problem where there are none: it suffices to re-write the problem with cost-augmented consumption utility $\hat{v}(a; \theta_j)$. The following result is immediate:

Corollary 2. *Suppose that cost-augmented preferences satisfy log submodularity in differences. Then the monopolists's optimal mechanism is a posted-price mechanism.*

The key economic insight won here is to observe that the sufficient condition to render posted-price mechanisms profit-maximizing becomes more stringent when production cost $c(a)$ per individual purchase a are non-zero.

Formally, one must now satisfy

$$\frac{v(a_3; \theta_2) + c(a_3) - v(a_2; \theta_2) - c(a_2)}{v(a_2; \theta_2) + c(a_2) - v(a_1; \theta_2) - c(a_1)} \leq \frac{v(a_3; \theta_1) + c(a_3) - v(a_2; \theta_1) - c(a_2)}{v(a_2; \theta_1) + c(a_2) - v(a_1; \theta_1) - c(a_1)} \quad (3)$$

for all $a_3 > a_2 > a_1$ and $\theta_2 > \theta_1$. Even when preferences $v(a; \theta)$ do satisfy log submodularity in differences, this is hardly sufficient. Provided that cost are increasing, simple manipulation of 3 reveals that the difference ratio of cost must now be bounded by a ratio of increasing differences of consumption utilities:

$$\frac{c(a_3) - c(a_2)}{c(a_2) - c(a_1)} \leq \frac{[v(a_3; \theta_2) - v(a_2; \theta_2)] - [(a_3; \theta_1) - v(a_2; \theta_1)]}{[v(a_2; \theta_2) - v(a_1; \theta_2)] - [(a_2; \theta_1) - v(a_1; \theta_1)]}$$

Or, more intuitively, production cost cannot be too convex. That this is restrictive is shown by the following example:

Example 3 (linear cost). *Suppose that consumption utilities are given by $v(a; \theta) = \frac{a^{1-\theta}}{1-\theta}$ where $a \in [0, 1]$ and $\theta \in (0, 1)$. One verifies immediately that $v(a; \theta)$ satisfies log submodularity in differences; i.e., $\partial_{a\theta}^2 \ln\{\partial_a v(a; \theta)\} \leq 0$. If one were to introduce linear cost however, $c(a) = c a$, one can easily verify that the augmented consumption utility $\hat{v}(a; \theta) = \frac{a^{1-\theta}}{1-\theta} + c a$ does not satisfy log submodularity in differences; i.e., $\partial_{a\theta}^2 \ln\{\partial_a v(a; \theta) + c\} > 0$ when a is small.*

3 Durable good monopoly

I now show that an analogous result ensures that a posted-price mechanism is profit-maximizing in the durable good monopoly.

Set-up

A monopolist seller sells a unit good over time. There are N buyer types $\{x_1, \dots, x_N\}$, distributed according to frequencies $m(x_j)$ with $M(x_j) = \sum_{k=1}^j m(x_k)$. Production is costless. Each consumer decides whether and if so when to purchase the good. The seller cannot observe the buyers' types and proposes a direct mechanism $(\tau(x_j), q(x_j))$, detailing the time of trade $\tau(x_j)$ as well as the price $q(x_j)$. Buyer type x_j 's benefit of trade at time $\tau(x_j)$ and price $q(x_j)$ is given by

$$U(\tau(x_j), x_j) - e^{-r\tau(x_j)} q(x_j),$$

where r denotes the discount rate and $U(\tau, x)$ is decreasing in τ . The seller's price schedule q evolves over time as to maximize the discounted sum of payments made by consumers,

$$\sum_{j=1}^N e^{-r\tau(x_j)} q(x_j) m(x_j),$$

subject to incentive compatibility and participation constraints. First, assume single-crossing: for all $\tau_2 > \tau_1$ and $j > i$:

$$U(\tau_1, x_i) - U(\tau_2, x_i) \geq U(\tau_1, x_j) - U(\tau_2, x_j).$$

This means that the incremental gains from trading earlier are greater for lower types, and consequently the seller wishes to trade with low types first.

Preliminaries

According to standard results the time of trade $\tau(x_j)$ is implementable by some pricing rule $q(x_j)$ if and only if τ is non-decreasing in type. Following standard arguments define indirect utility

$$v(x_j) = \max_{j \in \{1, \dots, N\}} U(\tau(x_j), x_i) - e^{-r\tau(x_j)} q(x_j),$$

and a mechanism $(\tau(x_j), q(x_j))$ is said to be incentive compatible if

$$v(x_j) - v(x_i) \geq U(\tau(x_i), x_j) - U(\tau(x_i), x_i)$$

for all $i, j \in \{1, \dots, N\}$. One immediately observes that one can re-write the seller's profit as

$$\sum_{i=1}^N e^{-r\tau(x_i)} q(x_i) m(x_i) = \sum_{i=1}^N [U(\tau(x_i), x_i) - v(x_i)] m(x_i).$$

From here it follows that the profit-maximizing mechanism must assign minimal indirect utility to the buyer, leaving incentive constraints to be upward binding⁵:

$$v(x_{k+1}) - v(x_k) = U(\tau(x_k), x_{k+1}) - U(\tau(x_k), x_k) \quad \forall k \in \{1, \dots, N\}.$$

Noting that $v(x_N) = 0$ due to binding participation constraints at the top, this allows to rewrite indirect utility as

$$v(x_k) = \sum_{j=k}^N U(\tau(x_j), x_j) - U(\tau(x_j), x(j+1))$$

(where x_{N+1} is an accounting zero type for which $U(\tau, x_{N+1})$ for all τ). Thus the weighted sum of indirect utilities can be expressed as

$$\sum_{i=1}^N v(x_i)m(x_i) = \sum_{i=1}^N \sum_{j=i}^N [U(\tau(x_j), x_j) - U(\tau(x_j), x_{j+1})]m(x_i) = \sum_{i=1}^N [U(\tau(x_i), x_i) - U(\tau(x_i), x_{i+1})]M(x_i).$$

And the monopolist seller's profit for a given trading time τ (and the associated profit-maximizing pricing rule q) can be written as

$$\begin{aligned} \Pi(\tau) &\equiv \sum_{i=1}^N U(\tau(x_i), x_i)m(x_i) - \sum_{i=1}^N [U(\tau(x_i), x_i) - U(\tau(x_i), x_{i+1})]M(x_i) \\ &= \sum_{i=1}^N [U(\tau(x_i), x_{i+1}) - U(\tau(x_{i+1}), x_{i+1})]M(x_i). \end{aligned}$$

Key argument

Suppose that the seller's profit-maximizing mechanism τ^* were not a posted price. It is known that $\tau^*(x_1) = 0$. Then there exist $m < n$, $m, n \in \{2, \dots, N+1\}$ such that

$$0 = \tau^*(x_1) = \dots = \tau^*(x_{m-1}) < \tau^*(x_m) = \dots = \tau^*(x_{n-1}) < \tau^*(x_n)$$

⁵Note that incentive constraints are still satisfied. Fix arbitrary $j, i \in \{1, \dots, N\}$, then

$$\begin{aligned} &v(x_j) - v(x_i) - [U(\tau(x_i), x_j) - U(\tau(x_i), x_i)] \\ &= \sum_{k=i}^{j-1} \left\{ v(x_{k+1}) - v(x_k) - [U(\tau(x_i), x_{k+1}) - U(\tau(x_i), x_k)] \right\} \\ &= \sum_{k=i}^{j-1} \left\{ (U(\tau(x_k), x_{k+1}) - U(\tau(x_i), x_{k+1})) - (U(\tau(x_k), x_k) - U(\tau(x_i), x_k)) \right\} \geq 0 \end{aligned}$$

due to single-crossing.




Figure 2: optimal mechanism τ^* and deviations $\bar{\tau}$ and $\underline{\tau}$.

(Note that the accounting zero type's trading time $\tau^*(x_{N+1})$ can be chosen freely as long as it exceeds $\tau^*(x_N)$.) Consider two deviations:

$$\bar{\tau}(x_j) = \begin{cases} \tau^*(x_j) & \text{if } j \in \{1, \dots, m-1\} \cup \{n, \dots\} \\ \tau^*(x_n) & \text{if } j \in \{m, \dots, n-1\} \end{cases}$$

and

$$\underline{\tau}(x_j) = \begin{cases} \tau^*(x_j) & \text{if } j \in \{1, \dots, m-1\} \cup \{n, \dots\} \\ \tau^*(x_{m-1}) & \text{if } j \in \{m, \dots, n-1\}. \end{cases}$$

Optimality of τ^* requires that $\Pi(\tau^*) \geq \max\{\Pi(\bar{\tau}), \Pi(\underline{\tau})\}$. Therefore

$$\begin{aligned} & [U(\tau^*(x_{m-1}), x_m) - U(\tau^*(x_m), x_m)]M(x_{m-1}) + [U(\tau^*(x_{n-1}), x_n) - U(\tau^*(x_n), x_n)]M(x_{n-1}) \\ & \geq \max \{ [U(\tau^*(x_{m-1}), x_m) - U(\tau^*(x_n), x_m)]M(x_{m-1}) + [U(\tau^*(x_{m-1}), x_n) - U(\tau^*(x_n), x_n)]M(x_{n-1}). \end{aligned}$$

This is equivalent to

$$\begin{aligned} & [U(\tau^*(x_n), x_m) - U(\tau^*(x_m), x_m)]M(x_{m-1}) \geq [U(\tau^*(x_n), x_n) - U(\tau^*(x_{n-1}), x_n)]M(x_{n-1}) \\ & [U(\tau^*(x_m), x_m) - U(\tau^*(x_{m-1}), x_m)]M(x_{m-1}) \leq [U(\tau^*(x_{n-1}), x_n) - U(\tau^*(x_{m-1}), x_n)]M(x_{n-1}). \end{aligned}$$

Dividing one by the other and noting that by construction $\tau^*(x_{n-1}) = \tau^*(x_m)$, one obtains

$$\frac{U(\tau^*(x_n), x_m) - U(\tau^*(x_m), x_m)}{U(\tau^*(x_m), x_m) - U(\tau^*(x_{m-1}), x_m)} \geq \frac{U(\tau^*(x_n), x_n) - U(\tau^*(x_m), x_n)}{U(\tau^*(x_m), x_n) - U(\tau^*(x_{m-1}), x_n)}.$$

This proves the following proposition:

Proposition 2. *If, for any $\tau_3 > \tau_2 > \tau_1$, the ratio*

$$\frac{U(\tau_3, x_j) - U(\tau_2, x_j)}{U(\tau_2, x_j) - U(\tau_1, x_j)}$$

is non-decreasing in j , then a posted-price mechanism will be profit-maximizing. If, moreover, the ratio is strictly increasing in j , then every profit-maximizing mechanism is a posted-price mechanism.

One can rephrase the assertion as follows: *If high valuation (low type) consumers' utilities over the timing of their purchase are a concave transformation of low valuation (high type) consumer's utilities, then the optimal pricing rule is a posted-price mechanism.* That this is necessary follows from identical arguments as in the preceding section.

Lastly, note that the result encompasses a finding by Stokey (1979): *Suppose there exist non-decreasing functions \hat{v} and \tilde{v} such that $v(\tau; x) = \hat{v}(\tau)\tilde{v}(x)$, i.e. v is separable. Then the monopolist's optimal pricing rule takes the form of a posted-price mechanism.* This is for instance the case when $U(\tau, x) = e^{-\tau r}x$.

4 Conclusion

When are posted-price mechanisms profit-maximizing? I provided a unified answer to this question within two distinct quasi-linear environments: monopolistic second-degree price discrimination and the durable good monopoly with commitment.

The analysis singled out a plausible ranking on the concavity of buyer utilities whereas distributional assumptions were shown to be irrelevant. When said ranking obtains, the analysis revealed that posted-price mechanisms, not sophisticated nonlinear pricing rules are profit-maximizing absent individual production cost. In contrast, when individual production cost are large, the sufficiency condition becomes more stringent.

In the case of the durable good monopoly I showed that abstaining from reducing the price in the future is a plausible strategy to maximize profits. Commitment alone—contrary to received intuition—was insufficient however to guarantee that this was the case.

To the extent that monopolists do price-discriminate even when individual production cost are negligible, these findings invite renewed inquiry into the trade-offs of price discrimination. Where it occurs, static profit maximization appears to be an unsatisfactory explanation of said practice.

Appendix

A Random Mechanisms

I here investigate random mechanisms in the second-degree price discrimination framework. A direct random mechanism assigns to each agent type a lottery over allotments with support in $[0, e]$ as well as a transfer. In line with Strausz (2006), I show that proposition 1 equally obtains if one considers the larger class of random mechanisms. More specifically, I show that no random mechanism can outperform a deterministic mechanism when preferences satisfy either log supermodularity in differences, or log submodularity in differences.

To begin with, note that the use of lotteries can only be prevalent for intermediate types. Due to single-crossing the highest type will always consume the total endowment. This means that some intermediate types are to face greater risk than the highest type. This can never be a property of the seller's optimal mechanism if preferences satisfy log supermodularity in differences. For log supermodularity in differences, interpreted in terms of risk-preferences, asserts that lower types are more risk-averse. Thus every incentive compatible random mechanism admits a profitable deviation to the seller: propose to the highest type subjected to a lottery her certainty equivalent instead. Since this type is more risk-averse than any superior type, this relaxes superior types' incentive constraints and raises the scope for the seller to uniformly increase transfers on all types superior. In conclusion, under log supermodularity in differences the seller's preferred mechanism is deterministic.

Consider then the case where preferences satisfy log submodularity in differences. Interpreted in terms of risk-preferences, this means that higher types are more risk-averse. Potentially, lotteries may play a role here. As the following reasoning shows, the seller's preferred lotteries are Bernoulli assigning positive probability to the two states $\{0, e\}$. Indeed, for any lottery assigning positive probability to some interval $[\underline{a}, \bar{a}] \subseteq]0, e[$ one can propose another lottery yielding a given type θ_k 's certainty equivalent with support in $\{0, e\}$. This construction relies on a sequence of certainty-equivalent-preserving spreads, implying that types greater than θ_k will prefer the initial lottery over the Bernoulli lottery.

Each trade of an optimal random mechanisms thus stipulates a trading probability $p(\theta_j)$ as well as a transfer $t(\theta_j)$ (where as before I restrict attention to direct mechanisms). Clearly, under single-crossing a mechanism is incentive compatible if and only if $p(\theta_j)$ is non-decreasing in θ_j . Incentive compatibility rewrites as

$$u(\theta_{j+1}) - u(\theta_j) = p(\theta_j)[v(e; \theta_{j+1}) - v(e; \theta_j)],$$

where as before $u(\theta_j)$ denotes indirect utility of type θ_j and satisfies $u(\theta_j) = p(\theta_j)v(e; \theta_j) - t(\theta_j)$. For notational convenience define $\Delta p(\theta_j) = p(\theta_j) - p(\theta_{j-1})$. Then iterating yields

$$u(\theta_{j+1}) = p(\theta_j)v(e; \theta_{j+1}) - \sum_{k=1}^j \Delta p(\theta_k)v(e; \theta_k),$$

whence

$$\begin{aligned}
\sum_{j=1}^N u(\theta_j)m(\theta_j) &= \sum_{j=1}^N \left[p(\theta_{j-1})v(e; \theta_j) - \sum_{k=1}^{j-1} \Delta p(\theta_k)v(e; \theta_k) \right] m(\theta_j) \\
&= \sum_{j=1}^N p(\theta_{j-1})v(e; \theta_j)m(\theta_j) - \sum_{j=1}^N \Delta p(\theta_j)v(e; \theta_j)(1 - M(\theta_j)) \\
&= \sum_{j=1}^N p(\theta_j)[v(e; \theta_{j+1}) - v(e; \theta_j)](1 - M(\theta_j)).
\end{aligned}$$

One may thus express the sum over transfers, the seller's objective, as

$$\begin{aligned}
\sum_{j=1}^N t(\theta_j)m(\theta_j) &= \sum_{j=1}^N [p(\theta_j)v(e; \theta_j) - u(\theta_j)]m(\theta_j) \\
&= \sum_{j=1}^N \Delta p(\theta_j)v(e; \theta_j)(1 - M(\theta_{j-1})).
\end{aligned}$$

Then, if price discrimination were optimal, there exist types $\theta_j > \theta_i$ such that $a(\theta_j) > a(\theta_{j-1}) = a(\theta_i) > a(\theta_{i-1}) \geq 0$. And by revealed preferences the seller must strictly prefer such regime over pooling types $\{\theta_i, \dots, \theta_{j-1}\}$ at an allotment $a(\theta_{i-1})$ or $a(\theta_{j-1})$. Therefore

$$\begin{aligned}
&\Delta p(\theta_j)v(e; \theta_j)(1 - M(\theta_{j-1})) + \Delta p(\theta_i)v(e; \theta_i)(1 - M(\theta_{i-1})) \\
&> \max \{ (\Delta p(\theta_j) + \Delta p(\theta_i))v(e; \theta_j)(1 - M(\theta_{j-1})); (\Delta p(\theta_j) + \Delta p(\theta_i))v(e; \theta_i)(1 - M(\theta_{i-1})) \}.
\end{aligned}$$

This is equivalent to

$$\begin{aligned}
v(e; \theta_i)(1 - M(\theta_{i-1})) &> v(e; \theta_j)(1 - M(\theta_{j-1})) \\
v(e; \theta_j)(1 - M(\theta_{j-1})) &> v(e; \theta_i)(1 - M(\theta_{i-1})),
\end{aligned}$$

thus establishing a contradiction.

Non-stationary Search and Assortative Matching

Joint with Nicolas Bonneton

1

1 Introduction

Homer (Odyssey XVII, 218) claims that 'Gods join like things with like things'. This is arguably the oldest affirmation of what economists call *positive assortative matching* (PAM). In more contemporary language, PAM means that individuals that share similar characteristics tend to be matched with one another. A vivid research agenda, that goes back to Becker (1973), has since explored the foundations of PAM in decentralized matching markets. We follow this line of inquiry within a dynamic random search matching framework where finding a potential partner is haphazard and time-consuming (refer to Chade et al. (2017) for a review).

Our understanding of sorting in dynamic environments is limited. Existing results (see Smith (2006), Morgan (1994), Shimer and Smith (2000), and Atakan (2006)) derive sufficient conditions for PAM solely in the steady state, where entry and exit exactly balance each other at every point in time.² Confining the theory to the steady state is, in many instances, a stretch on its empirical plausibility. Some markets display cycling behavior such as seasonal rental markets. Other markets gradually clear over time. Cases in point are the labor market for freshly minted graduates, and, more colorfully, the "market" for prom dates during the last year of high school.

This paper is the first to derive sufficient conditions for PAM away from the steady state. Those conditions ensure that more highly ranked agent types maintain a greater acceptance threshold throughout—even when the search pool population fluctuates over time. We do so in a paradigm where payoffs are non-transferable, i.e., exogenously given, as studied in Morgan (1994) and Smith (2006).³

¹This is a revised version of my job market paper.

²Such a restrictive focus can partly be justified by the considerable technical challenges associated with non-stationary analysis. As pointed out by Smith (2011), "even the simplest non-stationary models can be notoriously intractable."

³Examples include the marriage market, rent-controlled rental housing markets, or labor markets bound by

In line with the literature we develop a continuous-time, infinite-horizon matching model in which pairs of heterogeneous agents randomly meet one another. If both agents give their consent they form a permanent match and derive a non-transferable payoff. They subsequently exit the search pool. Otherwise they keep waiting for a more suitable partner. Crucially, match payoffs are increasing in the partner's type. We consider the two canonical ways of incorporating search frictions: time discounting and explicit search cost. In contrast to the existing literature search is non-stationary. This means that entry and exit from the search pool need not balance each other, as when there is no entry. Accordingly, optimal strategies change endogenously over time as to take into account the time-varying size and composition of the search pool.

To date, the literature has derived equilibrium sorting conditions by drawing on an explicit characterization of the value of search in the steady state. Non-stationary analysis forecloses this avenue as the time-varying value of search is a complicated object to handle. It is characterized by an integral over an infinite time horizon taking as its argument the population dynamics, which are itself a solution to an infinite-dimensional system of integral equations. We circumvent the ensuing tractability issues by using a revealed preferences argument: superior types, being more desirable, can exploit their superior match offerings and replicate match outcomes of any inferior type. Such deviations must be weakly dominated by the actual value of search—establishing a lower bound on superior types' value of search. This lower bound serves as the keystone of all of our equilibrium sorting results. In particular, such lower bound is used to derive novel and short proofs of two existing sorting results (refer to theorems 2 and 2') that hold in the steady state: *if payoffs are log supermodular, then there is PAM when search is costly due to time discounting* as established by Smith (2006); *if payoffs are supermodular, then there is PAM when search entails an explicit time-invariant flow cost* as established by Morgan (1994).

In a non-stationary environment, aforementioned steady state sufficient conditions are insufficient to guarantee PAM. Here, unlike in the steady state, the lowest type accepted today need not be the worst possible match outcome for all future times. As the search pool evolves over time, agents may face a less favorable selection of types to match with in the future. Then an agent who rejects a given type initially, may accept an inferior type at a later stage. This requires an agent to weigh the current acceptance decision against both the upside risk of matching with a superior and the downside risk of ending up with an inferior type in the future. Supermodularity and log supermodularity are mute on this trade-off. In particular, payoff log supermodularity implies that higher types relatively better like to be matched with higher types. On the other hand, it stipulates that higher types stand more to lose from matching with a lower type. Depending on which effect dominates, higher or lower types are choosier. This leaves open the possibility that higher type's fear of the worst outcome upsets PAM, even though steady conditions are satisfied.

The first main contribution of this article (refer to theorems 3 and 3') is to establish that if the respective steady state sufficient condition holds and if payoffs satisfy *log supermodularity in differences*, then there is positive assortative matching across all equilibria. Log supermodularity in differences is a terminology that we introduce here. It is satisfied if for all $y_1 < y_2 < y_3$

binding minimum wages or collective bargaining agreements. In Bonneton and Sandmann (2021) we investigate non-stationary search and matching under Nash bargaining, i.e., allow for time-varying transfers instead.

and $x_1 < x_2$ we have

$$\frac{\pi^X(y_3|x_2) - \pi^X(y_2|x_2)}{\pi^X(y_2|x_2) - \pi^X(y_1|x_2)} \geq \frac{\pi^X(y_3|x_1) - \pi^X(y_2|x_1)}{\pi^X(y_2|x_1) - \pi^X(y_1|x_1)},$$

where $\pi^X(y|x)$ represents agent type x 's payoff from population X if matched with an agent from the opposite population Y of type y (A symmetrical construction applies for the opposite population Y .) Assuming differentiability of payoffs, this condition is equivalent to log supermodularity of $\partial_y \pi(x, y)$. Provided that the respective steady state sufficient condition holds, we also show that log supermodularity in differences is the weakest sufficient condition for PAM: if the inequality holds with opposite sign for an interval of types there exist environments for which PAM fails to obtain. Observe that, unlike steady state sufficient conditions, log supermodularity in differences arises irrespective of how search cost is modeled.

Our analysis reveals an intimate link between the time-variant nature of search frictions and risk preferences. Indeed, rather than speaking about match payoffs we can interpret type x 's payoff over partners y as a utility function. Then log supermodularity in differences defines a ranking over risk preferences in the sense of Arrow (1965)-Pratt (1964). Accordingly, if the respective steady state sufficient condition holds, our main contribution states that *the weakest sufficient conditions for positive assortative matching is that more desirable individuals are less risk-averse*.

There is mounting empirical evidence that characteristics commonly attributed to desirability such as cognitive skills, education, health or income strongly correlate with risk preferences.⁴ For instance, Dohmen et al. (2010) conduct a real-stakes experiment with a large, representative sample of the German population to test whether risk-aversion and patience are related with cognitive ability. The authors find that individuals with higher cognitive ability are significantly more willing to take financial risks and are more patient. Moreover, Dohmen et al. (2011) find significant correlations between financial and non-financial measures of risk-aversion. This suggests that those individuals to which society attributes the greatest desirability are also the greatest risk-takers in matching markets. Our work, considered together with the empirical evidence, provides a theoretical foundation as to why positive assortative matching may arise.

The second main contribution we make concerns the existence of a non-stationary equilibrium (refer to theorem 1). We show that a non-stationary equilibrium exists where meeting rates, entry rates and payoffs are defined in utmost generality. This constitutes a leap forward in a literature that had largely confined itself to establishing the existence of stationary equilibrium amid rigid, closed-form assumptions regarding the economy's fundamentals.

Given the preoccupation with existence one may wonder: how can a non-stationary equilibrium fail to exist? This question admits two answers, one technical and one economical: First, the economic environment must be well-defined so that the agent's optimization problem is meaningful. One must thus propose regularity conditions such that population dynamics are non-explosive and well-defined for any given strategy profile. Assumptions on payoffs further ensure that the value of search is well-behaved. Slight technical issues come to the fore if pay-

⁴Refer to Dohmen et al. (2010) and Dohmen et al. (2011), as well as Guiso and Paiella (2004), Frederick (2005), Benjamin et al. (2013) and Noussair et al. (2013) for evidence.

offs are monotone in the matched partner's type, but not in one's own type; then more highly ranked agent types may have a lower value of search (which is not a problem as long as the value of search is of bounded variation which will be the case with differentiable match payoffs for instance, but a weaker assumption will suffice.) Secondly, and perhaps more relevant from an economic point of view, an equilibrium solves a non-trivial feedback loop between beliefs and strategies: beliefs over future match opportunities must be generated by optimal acceptance strategies. To sketch but one example where this fails, suppose that the meeting rate exhibits a discontinuity in the size of the search pool: assume there are few meetings if the population size were above some threshold, yet many for a population size below some threshold; in effect crowdedness crowds out meetings. What are the optimal strategies in such an environment? Intuitively it is individually rational to be very accepting of other agents when the meeting rate is small, but to be choosy when the meeting rate is large. Observe how this gives rise to a self-defeating feedback loop: if players anticipate few future meetings they will be very accepting, so that the future population size will be small. In effect, their behavior defies their own expectations in that the small population size gives rise to a large meeting rate. What arises in this sketch is mis-coordination between beliefs over future match opportunities and (individually optimal) strategy profiles. Whatever players believe to happen can impossibly arise given the behavior that is optimal for such a belief. We propose (relatively weak) regularity conditions which prevent this from happening.

Our proof revolves around a non-trivial application of Schauder's fixed point theorem. Since the value function is known to exhibit discontinuities in type and the strategy space fails to be compact in infinite time, we construct a sequence of approximate best-response functions that are amenable to Schauder's fixed-point arguments. To show the convergence of such a sequence we rely on a multidimensional generalization of Helly's selection theorem (see Idczak and Walczak (1994), Leonov (1996) and Chistyakov and Tretyachenko (2010b)).

Certain features of the considered framework are noteworthy. Unlike the existing literature, we do not require functional specifications for the *meeting rate* and the *entry rate*. Instead they depend, in a general fashion, on the size and the composition of the search pool. In particular, our model encompasses the two most studied meeting rates: linear or constant returns to scale (e.g. Burdett and Coles (1997), Mortensen and Pissarides (1994)) and quadratic search technologies (e.g. Shimer and Smith (2000) and Smith (2006)). Likewise, we allow the population dynamics to be induced by commonly used entry rates such as no entry, constant flows of entry (e.g. Burdett and Coles (1997)), and exogenous match destruction (e.g. Shimer and Smith (2000) and Smith (2006)).

The rest of this section discusses related literature. Section 2 lays out the model. Section 3 establishes positive assortative matching. Section 4 considers explicit search cost. Section 5 concludes. Proofs, if not found in the text, are relegated to the appendix.

Related literature

The following table summarizes the existing conditions on payoffs that ensure positive assortative matching in the steady state. We denote as π and f the individual payoff and total (sym-

metric) match output, respectively. Both objects take one's own type as the first argument and others' type as the second argument. Subscripts stand for partial derivatives. Observe that one of the conditions of Shimer and Smith (2000), $(\log f_1)_{12} > 0$, is mathematically equivalent to log supermodularity in differences (assuming differentiability of payoffs). The interpretation in terms of risk preferences is different in their context as here f represents joint match output, and not an individual payoff.⁵

	Non-transferable payoff	Transferable payoff
Frictionless	$\pi_2 > 0$ Becker (1973)	$f_{12} > 0$ Becker (1973)
Search friction i: time cost	$\pi_2 > 0, (\log \pi)_{12} > 0$ Smith (2006)	$\partial_2 f(0, x) \leq 0, \partial_2 f(1, x) \geq 1$ $f_{12} > 0, (\log f_1)_{12} > 0, (\log f_{12})_{12} > 0$ Shimer and Smith (2000)
Search friction ii: fixed cost	$\pi_2 > 0, \pi_{12} > 0$ Morgan (1994)	$f_{12} > 0$ Atakan (2006)

Interest in assortative matching extends beyond matching with random search. We mention here two prominent papers in related frameworks that derive conditions for PAM. Legros and Newman (2007) show that a generalized increasing differences condition extends Becker (1973) to not fully transferable payoffs. Studying directed search, Eeckhout and Kircher (2010) show that positive assortative matching obtains if and only if the joint match surplus is root supermodular.

Existence proofs feature prominently in the literature on matching with random search, with the exclusive focus being on the existence of steady state equilibria. Shimer and Smith (2000) prove existence of a steady state equilibrium under Nash bargaining and quadratic search. Smith (2006) derives an identical result when payoffs are not transferable. Lauermaun and Nöldeke (2015) generalize this literature when there are finitely many types only. They dispense with the quadratic meeting rate assumption and consider a general sharing rule, encompassing both transferable and non-transferable match payoffs. In contrast, our focus is on general existence, regardless of whether the economy is in the steady state or not. Here, the time-varying nature of non-stationary strategies increases the dimension of the strategy space. This requires new tools from analysis, dealing with the compactness of higher-dimensional function spaces.

We also wish to mention a related literature on matching with random search and heterogeneous agents that explicitly solves for non-stationary dynamics under parametric assumptions; see for instance Boldrin et al. (1993) and Burdett and Coles (1998). Recently, there has been a more applied focus; see Eeckhout and Lindenlaub (2019) (and the references therein) for an application to on-the-job search. This literature identifies a sorting externality: by rejecting less desirable types' match offers, a greater number of agents stay in the search pool. This renders the prospect of future meetings with more desirable types more likely. Oftentimes, the sorting

⁵Whereas the present article is concerned with non-transferable payoffs, in Bonneton and Sandmann (2021) we generalize and discuss the limits of Shimer and Smith (2000)'s steady state result in non-stationary environments.

externality self-sustains multiple equilibria for identical initial conditions.⁶ We stress that the results derived in the present article hold true for all equilibria.

Relatedly, various authors have worked on the empirical identification of sorting of heterogeneous agents in dynamic matching models with random search. Hagedorn et al. (2017) demonstrate that, in principle, sorting of heterogeneous agents as studied in Shimer and Smith (2000) and Smith (2006) can be identified in the steady state. No such identification result exists for non-stationary dynamics. Lise and Robin (2017) propose and estimate a non-stationary random search matching model of workers and firms with on-the-job search when dynamics are driven by exogenous shocks. A key finding is that the estimated matching sets during lower productivity states are smaller. This leads to more assortative matching in hiring (from unemployment) during recessions.

It is noteworthy that the link between risk preferences and assortative matching has previously been studied in contexts in which the purpose of matching is to share risk with no regard for search frictions (see Chiappori and Reny (2016), Schulhofer-Wohl (2006) and Legros and Newman (2007)). A key insight from these articles is that negative assortative matching arises in the sense that relatively more risk-averse agents are matched with relatively less risk-averse agents. In doing so, the least risk-averse individuals can take on risk previously born by the most risk-averse individuals. The link between risk preferences and assortative matching established in the present paper is of a different nature as it derives from search frictions and risk-sharing is not considered.

Log supermodularity in differences has been found to drive sorting results in unrelated areas such as moral hazard and price discrimination. Chade and Swinkels (2020) prove that log supermodularity in differences of the distribution function of signals is sufficient to guarantee that the first-order approach to hidden action problems with continuous action spaces is valid.⁷ Sandmann (2019) shows that under mild distributional assumptions log supermodularity in differences gives rise to monopoly price discrimination under zero-marginal cost.

2 The model

We develop a continuous-time, infinite-horizon matching model in which heterogeneous agents engage in time-consuming and haphazard search for one another. When two agents meet they observe each other's type. If both agents give their consent they form a permanent match and derive a time-invariant, non-transferable payoff. They subsequently exit the search pool. Otherwise they continue waiting for a more suitable partner. Each agent maximizes her expected present value of payoffs, discounted at rate $\rho > 0$.

⁶Interestingly, under the congestive properties of linear search, cyclical equilibria with endogenous separations may arise (see Burdett and Coles (1998)). Shimer and Smith (2001) find that cycling behavior may also be a feature of the planner's solution.

⁷Chade and Swinkels (2020) refer to log supermodularity in differences as the no-upward-crossing condition.

Set-up

There are two distinct populations X and Y . Each population comprises a continuum of heterogeneous agents that inhabit the economy's search pool. Agents of both populations are characterized by their type which encode the agent's attractiveness to others. We (usually) adopt the convention that type x denotes an agent from population X and type y an agent from population Y . Types are normalized and belong to the unit interval $[0, 1]$.

The distribution of types x from population X that are searching at time t is characterized by a measure $\mu_t^X(x)dx$. This means that the mass of agents with types in $[0, z]$ is given by $\int_0^z \mu_t^X(x)dx$. A symmetric construction applies for agent types y from population Y . We will refer to $\mu_t = (\mu_t^X, \mu_t^Y)$ as the economy's time t state variable. The economy is non-stationary in that μ_t , starting from initial state $\mu_0 = (\mu_0^X, \mu_0^Y)$, changes over time. The admissible state space for fixed time t is the space M^2 of those $\mu_t = (\mu_t^X, \mu_t^Y)$ for which $x \mapsto \mu_t^X(x)$ and $y \mapsto \mu_t^Y(y)$ are non-negative, measurable and bounded. Accordingly, denote $\bar{\mu}_0$ the initial state's uniform upper bound.

Payoffs and Strategies

Agents derive a time-independent one-time payoff if matched with another agent and zero if unmatched: denote $\pi^X(y|x) > 0$ the lump-sum payoff of agent type x from population X when matched with agent type y from population Y , and, symmetrically, $\pi^Y(x|y)$ agent type y 's payoff. Types are vertically differentiated; that is, matching with higher types is always more desirable.

Assumption 1 (Monotonicity of payoffs). *Match payoffs are uniformly bounded by some $\bar{\pi}$ and strictly increasing in the partner's type; that is, $\pi^X(y_2|x) > \pi^X(y_1|x)$ for all $y_2 > y_1$ and $\pi^Y(x_2|y) > \pi^Y(x_1|y)$ for all $x_2 > x_1$.*

Since payoffs are monotone, acceptance decisions have a very simple structure. Each agent accepts everyone above a certain cutoff type, for any other strategy would be weakly dominated by a cutoff strategy. In the aggregate, this defines a population strategy S^X as a time-moving cutoff: agent type $x \in [0, 1]$ from population X accepts agent type $y \in [0, 1]$ from population Y at time t if $S_t^X(x) \leq y$, and rejects such a type otherwise. A symmetric construction applies for agents from population Y . A strategy profile $S = (S^X, S^Y)$ gives rise to the match indicator function $m_t[S](x, y) \in \{0, 1\}$. It is equal to one if agent types x and y are willing to form a match, that is, if $y \geq S_t^X(x)$ and $x \geq S_t^Y(y)$, and zero otherwise. Denote \mathcal{S} the space of jointly (in type x and time t) Lebesgue measurable strategy profiles of a given population. The space of strategy profiles (of both populations) $S = (S^X, S^Y)$ is \mathcal{S}^2 .

It is implicit in this reasoning that any two agents of identical type play identical strategies. This is without loss of generality: as each agent type is atomless, no type can alter the evolution of the search pool in a payoff-relevant manner.

Meeting Rates

Over time agents randomly meet each other. Meetings follow an (inhomogeneous) Poisson point process. Such a process is characterized by the time-variant (Poisson) meeting rate $\lambda = (\lambda^X, \lambda^Y)$. We take λ to be a functional of the underlying state variable μ_t and time t . Then $\lambda_t^X(y|x) \equiv \lambda^X(t, \mu_t)(y|x)$ is agent type x 's time t meeting rate with an agent type y given the prevailing time t and state μ_t .⁸ The meeting rates λ_t^X and λ_t^Y are not arbitrary but intricately linked. Coherence of the model demands that the number of meetings of agent types x with agent types y must be equal to the number of meetings of agent types y with agent types x :

$$\lambda_t^X(y|x)\mu_t^X(x) = \lambda_t^Y(x|y)\mu_t^Y(y).$$

Our results on positive assortative matching further require that the meeting rate is hierarchical in the sense that higher types meet other agents at a weakly faster rate.

Assumption 2 (hierarchical search). *Higher types meet other agents at a faster rate; that is, $\lambda_t^X(y|x_2) \geq \lambda_t^X(y|x_1)$ for $x_2 > x_1$ and $\lambda_t^Y(x|y_2) \geq \lambda_t^Y(x|y_1)$ for $y_2 > y_1$.*

To better understand this assumption, we introduce two functionals $\beta_t^X(y|x) \equiv \beta^X(t, \mu_t)(y|x)$ and $\beta_t^Y(x|y) \equiv \beta^Y(t, \mu_t)(x|y)$ so that $\lambda_t^X(y|x) = \beta_t^X(y|x)\mu_t^Y(y)$ and $\lambda_t^Y(x|y) = \beta_t^Y(x|y)\mu_t^X(x)$. Then coherence demands that $\beta_t^X(y|x) = \beta_t^Y(x|y)$. And search is hierarchical if and only if $\beta_t^X(y|x) = \beta_t^Y(x|y)$ is non-decreasing in both x and y . It is easy to see that with two distinct populations hierarchical search allows for rich heterogeneity in meeting rates within a population. If instead the search pool population is symmetric, i.e., $\lambda_t^X(y|x) = \lambda_t^Y(y|x)$ and $\mu_t^X(x) = \mu_t^Y(x)$ for all $x, y \in [0, 1]$, then hierarchical search stipulates that $\beta_t^X(y|x) = \beta_t^Y(y|x) = \beta_t$ (dependent on time t and state μ_t but not on x and y). This is tantamount to the more restrictive assumption of anonymous search:

Example 1 (anonymity). $\lambda_t^X(y|x) \equiv \lambda_t^X(y)$ and $\lambda_t^Y(x|y) \equiv \lambda_t^Y(x)$ for all x and y .

Anonymity stands for the idea that the probability of meeting other agents does not depend on one's own type. Or, the meeting rate per se does not favor assortative matching. Anonymous search, albeit less general than hierarchical search, encompasses the most commonly studied meeting rates found in the literature: linear (e.g., Mortensen and Pissarides (1994), Burdett and Coles (1997)) and quadratic search technologies (e.g., Shimer and Smith (2000) and Smith (2006)). In the former case, the meeting rate is determined by the distribution of types in the search pool so that $\beta_t = 1 / \int_0^1 (\mu_t^Y(z) + \mu_t^X(z)) dz$. In the latter case the meeting rate is determined by the absolute size of the search pool so that $\beta_t = 1$. Those meeting rates are further discussed in Maskin and Diamond (1979). Importantly, our existence and sorting results do not depend on the functional form of the meeting rate. Rather, we require that search be hierarchical and satisfy some regularity conditions, namely Lipschitz continuity, as shall be made precise in the subsequent section on regularity conditions.⁹

⁸As is well known the (Poisson) meeting rate $\lambda_t^X(y|x)$ is equal to the flow expectation of future meeting.

⁹That sorting does not require anonymity speaks to its robustness. Even if search effort was exogenous, our results on assortative matching obtain as long as higher types choose to search more.

Population Dynamics

Initially, the search pool is characterized by a state $(\mu_0^X, \mu_0^Y) \in M^2$. Over time, agents' strategy profile $S = (S^X, S^Y)$ governs the evolution of the search pool through entry and exit. Any two agents $x, y \in [0, 1]$ of opposite populations that meet and mutually consent to form a match, i.e., $m_t[S](x, y) = 1$, exit the search pool. Furthermore, there is a flow of entry of type x and y agents at time t , denoted $\eta_t^X(x) \geq 0$ and $\eta_t^Y(y) \geq 0$. As usual, subscript t subsumes the dependence on time t and the state μ_t (so that $\eta_t^X(x) \equiv \eta^X(t, \mu_t)(x)$ and $\eta_t^Y(y) \equiv \eta^Y(t, \mu_t)(y)$). Taken together, entry and exit give rise to the population dynamics.

Definition 1 (Population dynamics). *Given a strategy profile $S = (S^X, S^Y) \in \mathcal{S}^2$, the evolution of the search pool is described by a system of integral equations with initial state $\mu_0 = (\mu_0^X, \mu_0^Y) \in M^2$:*

$$\mu_{t+h}^X(x) = \mu_t^X(x) + \int_t^{t+h} \left\{ -\mu_\tau^X(x) \int_0^1 \lambda_\tau^X(y|x) m_\tau[S](x, y) dy + \eta_\tau^X(x) \right\} d\tau \quad \forall h > 0, \quad (1)$$

symmetrically so for population Y, for all $x \in [0, 1]$ and $t \geq 0$.

Our formulation is that of an integral equation rather than a differential equation. This is due to the fact that $t \mapsto \mu_t^X(x)$ and $t \mapsto \mu_t^Y(y)$ will always be absolutely continuous, but not necessarily differentiable due to the discontinuities inherent in the match indicator function.

Assumption 3 (well-posed population dynamics). *The meeting and entry rate are (linearly) bounded and (uniform) Lipschitz continuous:*

- (i) $\eta_t^X(x) \leq L^\eta(1 + \mu_t^X(x))$ and $\eta_t^Y(y) \leq L^\eta(1 + \mu_t^Y(y))$;
- (ii) $\beta^X(y|x)$ and $\beta^Y(x|y)$ are uniformly bounded by some positive constant $\bar{\beta}$;
- (iii) $|\eta^X(t, \mu_t'')(x) - \eta^X(t, \mu_t')(x)| \leq L^\eta \max \{ \|\mu_t''^X - \mu_t'^X\|_\infty, \|\mu_t''^Y - \mu_t'^Y\|_\infty \}$, likewise for population Y;
- (iv) $|\beta^X(t, \mu_t'')(y|x) - \beta^X(t, \mu_t')(y|x)| \leq L^\beta \max \{ \|\mu_t''^X - \mu_t'^X\|_\infty, \|\mu_t''^Y - \mu_t'^Y\|_\infty \}$, likewise for population Y.

Assumption 3 encompasses several natural entry rates such as no entry and constant flows of entry (as in Burdett and Coles (1997)).

Proposition 1. *Suppose assumption 3 is in place. Then for every $S = (S^X, S^Y) \in \mathcal{S}^2$ and initial condition $\mu_0 = (\mu_0^X, \mu_0^Y) \in M^2$ the population dynamics (1) admit a unique solution.*

The proof (see appendix A.1) adapts the finite-dimensional Picard-Lindelöf (or Cauchy-Lipschitz) existence and uniqueness theorem for ordinary differential equations to the (infinite-dimensional) continuum. In light of this proposition, denote $\mu_t[S]$ the time t distribution of agents in the search pool as uniquely characterized by (1) for given strategy profile S .

The value of search

Any given agent's experience in the search pool is characterized by random encounters with other agents. Presented with a match opportunity, an agent must weigh the immediate match payoff, $\pi^X(y|x)$ or $\pi^Y(x|y)$, against the value of search, $V_t^X[S](x)$ or $V_t^Y[S](y)$, the discounted expected future match payoff were one to continue one's search. Naturally, the optimal strategy is to accept to match with another agent whenever the payoff exceeds the option value of search.

Remark 1. *Agent type x accepts to match with agent type y if and only if $\pi^X(y|x) \geq V_t^X[S](x)$, symmetrically so for population Y .*

However intuitive, the value of search is tedious to define; it requires to form expectations over future match offers which depend on the strategy profile as well as individual acceptance decisions. The dependence on the strategy profile $S = (S^X, S^Y)$ is twofold. First, the strategy profile determines—through the population dynamics μ as characterized by (1)—the meeting rates $\lambda_t^X(y|x)$ and $\lambda_t^Y(x|y)$. Secondly, the strategy profile of the opposite population, say Y , determines agents' match opportunities conditional on meeting, i.e, those y for which $x \leq S_t^Y(x)$. Finally, the agent controls the matching rate through his own acceptance decision. Focusing on weakly undominated strategies, this takes once more the form of a time-varying (and Lebesgue measurable) threshold rule $t \mapsto s_t$ according to which one accepts y if $y \geq s_t$.

Then fix an individual strategy and a strategy profile (s, S) . The matching rate is given by

$$\Lambda_t^X[s|S](y|x) \equiv \mathbb{1}\{s_t \leq y\} \mathbb{1}\{S_t^Y(y) \leq x\} \lambda_t^X(y|x).$$

Following the definition of the Poisson point process, the probability law characterizing future matches admits a density,

$$p_{t,T}^X[s|S](y|x) \equiv \Lambda_T^X[s|S](y|x) \exp\left\{-\int_t^T \int_0^1 \Lambda_\tau^X[s|S](y|x) dy d\tau\right\}$$

so that for any subset of types, $A \subset [0, 1]$, $\int_A \int_t^T p_{t,\tau}^X[s|S](y|x) d\tau dy$ is the probability that for a given strategy profile S and individual matching threshold s , agent type x , unmatched at time t matches with some $y \in A$ until time T . A symmetric construction applies for population Y .

The value of search $V_t^X[S](x)$ is the discounted expected match payoff when the agent chooses an optimal individual strategy s :

$$V_t^X[S](x) = \sup_s W_t^X[s|S](x) = \sum_s \int_t^\infty e^{-\rho(\tau-t)} \int_0^1 \pi^X(y|x) p_{t,\tau}^X[s|S](y|x) dy d\tau. \quad (2)$$

What follows from here is a representation result. The term in the squared brackets in the preceding can be subsumed as a discounted match probability over the whole time interval. Then the time-invariance of payoffs allows us to represent the value of search as an integral over types only:

Remark 2. *There exists a non-negative discounted match probability $Q_t^X[S](y|x)dy$ such that*

$$V_t^X[S](x) = \int_0^1 \pi^X(y|x)Q_t^X[S](y|x)dy \quad \text{and} \quad \int_0^1 Q_t^X[s|S](y|x)dy < 1.$$

This characterization highlights what will be key throughout the analysis: it will be properties of abstract measures, not explicit strategy profiles which will allow us to make progress on equilibrium existence and sorting results.

Regularity

We conclude this section with a result on the regularity of the value of search. While superfluous to the economics pertaining to sorting, the following will play an important role to ensure the existence of an equilibrium—arguably necessary for any positive analysis. First, we make an assumption.

Assumption 4. *$x \mapsto \pi^X(y|x)$ and $y \mapsto \pi^Y(x|y)$ admit a uniform bound L^π on total variation.*

Needless to say, this assumption is weaker than requiring that $\pi^X(y|x)$ and $\pi^Y(x|y)$ are continuously differentiable (as in Smith (2006)). We will make use of this assumption to prove proposition 2 (ii):

Proposition 2 (bounded variation of the value of search). *The value of search is*

(i) *Lipschitz continuous in time; i.e., for all moments in time $T : 0 < T < \infty$ there exists $C > 0$ such that for all $0 \leq t_1 < t_2 \leq T$ and $x \in [0, 1]$*

$$|V_{t_2}^X[S](x) - V_{t_1}^X[S](x)| \leq C |t_2 - t_1| \quad \text{for all } S \in \mathcal{S}^2,$$

symmetrically so for population Y.

(ii) *of uniform bounded variation in type; i.e., for all time indices $t \geq 0$, there exists $C > 0$ such that for all partitions of the type interval $[0, 1]$*

$$\sum_{i=0}^m |V_t^X[S](x_{i+1}) - V_t^X[S](x_i)| \leq C \quad \text{for all } S \in \mathcal{S}^2,$$

symmetrically so for population Y.

The proof is relegated to the appendix A.2. Of course, Lipschitz continuity as in (i) implies uniform bounded variation as in (ii). (i) derives from dynamic programming and requires bounded match payoffs; further that under the assumptions made the search pool population and in turn the meeting rate grow at most exponentially in time. (ii) is due to a mimicking argument (which makes use of assumptions 1 and 2), that shall be explored in depth when we discuss assortative matching.

3 Equilibrium existence

Having introduced the random search model, it remains to analyze the outcomes it predicts. Those correspond to the equilibria of the model. Here we shall concern ourselves with the (natural) definition of equilibrium and show (less trivial) that it exists.

Definition

An equilibrium is, first of all, a description of an outcome. Who is when in the search pool? How do agents resolve the trade-off between continuing their search in the hope of meeting a superior (and accepting) partner, and the immediate gains from matching?

Formally speaking, an outcome is a tuple (μ, V, S) , which describes the value of search, population dynamics and a strategy profile at all moments in time. In equilibrium, those three objects are linked with one another. Knowledge of the strategy profile S alone suffices to deduce the population dynamics (through (1)) and the value of search (through (2)). It is thus without loss of generality to posit that an equilibrium is a strategy profile S only. Optimality is addressed by remark 1: agent type x is prepared to match with agent type y if and only if his value of search—induced by the strategy profile—is smaller than the match payoff promised by agent type y . As is common in the literature, we discard equilibria in which agent types x and y mutually reject one another, $S_t^X(x) > y$ and $S_t^Y(y) > x$, even when both agents would be better off by accepting the proposed match. Such a selection would, for instance, be implied by restricting attention to either weakly undominated or trembling-hand perfect equilibria.

Definition 2 (Equilibrium). *An equilibrium (in weakly undominated strategies) is a strategy profile $S = (S^X, S^Y) \in \mathcal{S}^2$ such that at all times $t \geq 0$*

$$S_t^X(x) \leq y \quad \Leftrightarrow \quad V_t^X[S](x) \leq \pi^X(y|x),$$

symmetrically so for population Y .

Equivalently we can say that agent type y is the smallest acceptable type to agent type x , if the greater two of his value of search and the smallest possible match payoff are equal to the match payoff promised by agent type y . In effect, we can think of equilibria as fixed points of best-response mappings:

Remark 3. *S is an equilibrium (in weakly undominated strategies) if and only if it is a fixed point of the best-response mapping $f : \mathcal{S}^2 \rightarrow \mathcal{S}^2$, where $f = (f^X, f^Y)$ and population X 's best-response strategy profile is given by*

$$f^X[S](t, x) \mapsto (\pi^X(\cdot|x))^{-1}(\max\{V_t^X[S](x), \pi^X(0|x)\}),$$

symmetrically so for population Y .

In words: $f^X[S](x, t)$ denotes the acceptance threshold agent type x ought to maintain at time t given the prevailing strategy profile S .

Further assumptions

To establish our equilibrium existence result, we require two additional assumptions. First, not just the value of search (as shown in proposition 2), but best-responses must be of uniform bounded variation. In light of the construction of the fixed-point mapping, this motivates the following assumption:

Assumption 5. *The inverse of match payoffs $z \mapsto (\pi^X(\cdot|x))^{-1}(z)$ and $z \mapsto (\pi^Y(\cdot|y))^{-1}(z)$ is Lipschitz continuous.*

We must also strengthen assumption 3. In particular, we require that the meeting and entry rate be Lipschitz continuous in the L^1 norm, not just the sup norm. Therefore define

$$N(\mu'_\tau, \mu''_\tau) \equiv \max \left\{ \int_0^1 |\mu'^X_t(x) - \mu''^X_t(x)| dx, \int_0^1 |\mu'^Y_t(y) - \mu''^Y_t(y)| dy \right\}.$$

This distinction will be important when proving that no slight perturbation of the strategy profile could dramatically alter the value of search of a non-negligible subset of agents types (refer to proposition 6).

Assumption 6. *The meeting and entry rate are (pairwise) Lipschitz continuous:*

(iv) $|\eta^X(t, \mu'_t)(x) - \eta^X(t, \mu''_t)(x)| \leq L^\eta N(\mu'_\tau, \mu''_\tau)$, likewise for population Y;

(v) $|\beta^X(t, \mu'_t)(y|x) - \beta^X(t, \mu''_t)(y|x)| \leq L^\beta N(\mu'_\tau, \mu''_\tau)$, likewise for population Y.

Result

With assumptions 1-6 in place we can state our first main result.

Theorem 1 (existence). *An equilibrium (in weakly undominated strategies) exists.*

Equilibrium existence results are oftentimes unremarkable; existence can be ascertained by explicit construction. This is not the case here. Indeed, the preceding assumptions did not specify functional forms for economic fundamentals. Accordingly, we wish to emphasize the scope and generality of this result. It allows us to do positive equilibrium analysis, but also provides the theoretical foundation for numerical and (parametric) econometric analysis alike.

Proof outline

We relegate the formal proof of theorem 1 to the appendix. In what follows we provide an outline of the key arguments.

An equilibrium is a fixed point of the best-response mapping $f : \mathcal{S}^2 \rightarrow \mathcal{S}^2$. Our proof of equilibrium existence revolves around Schauder's topological fixed point theorem. To apply it,

we endow the space of strategy profiles with a topology, induced by the discounted supremum metric:

$$\mathbf{d}^2(S', S'') = \max \left\{ \sup_{(x,t)} e^{-t} |S'_t{}^X(x) - S''_t{}^X(x)|; \sup_{(x,t)} e^{-t} |S'_t{}^Y(y) - S''_t{}^Y(y)| \right\} \quad \text{for all } S', S'' \in \mathcal{S}^2.$$

Two strategy profiles are close if all agents at all moments in time maintain a similar acceptance threshold. We use discounting to deal with the unbounded time horizon; the reason why this works is that differences in strategy profiles far removed into the future lead to small contemporaneous differences in the value of search (see proposition 6).

Schauder's theorem asserts that a continuous mapping in a Hausdorff topological vector space whose image is compact admits a fixed point. We must thus prove two properties of the best-response mapping. First, f must be continuous. Secondly, the image of the best-response mapping, $f[\mathcal{S}]$, must be compact. As it turns out this is not straightforward. Under the usual supremum metric at least, neither continuity of the fixed-point mapping nor compactness of the strategy space is satisfied.

The problems are twofold. First, it has been widely noted that if an interval of agent types maintains an identical acceptance cutoff, those agent types of the opposite population above the cutoff will see far superior match opportunities in comparison to those below the cutoff. In effect, the value of search will exhibit a discontinuity at the cutoff type as we vary the prevailing strategy profile. This is notably the case under block segregation (refer to McNamara and Collins (1990) for the original result; also see Smith (2006) and references therein). Secondly, for the strategy space to be compact one must be able to construct a finite subcover. This means that, for every $\epsilon > 0$, one can identify a finite number of base strategy profiles; and every strategy profile is less than ϵ away from one of the base strategy profiles. Since strategy profiles can oscillate in time and type it is always possible to construct a strategy profile whose distance from any of the finite base strategy profiles is arbitrary close to one, in particular greater than ϵ . Or, rapid fluctuations in the strategy profile are problematic and the reason why the space of merely measurable strategy profiles fails to be compact.

Owing to the lack of continuity and compactness, our proof strategy is not to apply Schauder's fixed point theorem to the best-response mapping $f : \mathcal{S}^2 \rightarrow \mathcal{S}^2$ directly. Rather, we consider approximate best-response mappings $f_k : \mathcal{S}^2 \rightarrow \mathcal{S}^2$ which satisfy Schauder's conditions. Each approximate best-response mapping f_k is the average of individual agents' best-responses across neighboring types and moments in time. If f_k satisfies the laid out conditions, there exists a fixed point—an approximate equilibrium—of the approximate best-response mapping f_k . We will then be interested in possible limit points as the approximating error becomes small.

Our proof can be divided into two parts. In part one, we show that there exists a strategy profile which is a fixed point of the approximate best-response mapping. Or, we show that (i) f_k is continuous (lemma 6), and (ii) that its image, $f_k[\mathcal{S}^2]$, is the space of k -Lipschitz continuous strategy profiles, denoted $(\mathcal{S}_k)^2$ (lemma 5). And we observe that $(\mathcal{S}_k)^2$ is compact (proposition 5), since a function which is k -Lipschitz has a bounded slope, whence cannot oscillate too much. The existence of a fixed point of f_k , denoted $S_{<k>}$ is then guaranteed by Schauder's theorem.

If $S \rightarrow f[S]$ fails to be continuous, why is its approximation $S \rightarrow f_k[S]$ continuous? Recall




Figure 1: $f_k[S](t, x)$ is the average best-response in the square-neighborhood of (t, x) .

that there can only be a discontinuity if there exists (t, x) such that $S \mapsto V_t^X[S](x)$ is discontinuous. And the discontinuity in the value function at some strategy profile arises because there exists an interval of time in which almost all y of an interval of types maintain x as their cutoff type. That is to say, $\int_0^1 \int_t^T \mathbb{1}\{S_t^Y(y) = x\} d\tau dy > 0$ for T sufficiently large. Yet for all fixed T there can only be countably many (t, x) for which the integral measure is non-zero. Therefore there are only countably many (t, x) for which $S \mapsto f^X[S](t, x)$ is discontinuous, symmetrically so for population Y . And when averaging (as we do to construct the approximate best-response mapping) those discontinuities vanish.

$$\left\{ \begin{array}{l} ((\mathcal{S}_k)^2, \mathbf{d}^2) \text{ is compact} \\ f_k : (\mathcal{S}^2, \mathbf{d}^2) \rightarrow (\mathcal{S}^2, \mathbf{d}^2) \text{ is continuous} \\ f_k[\mathcal{S}^2] \subseteq (\mathcal{S}_k)^2 \end{array} \right. \quad \begin{array}{c} \text{Schauder} \\ \implies \end{array} \quad f_k \text{ admits a fixed point } S_{\langle k \rangle}$$

Figure 2: Part one of the existence proof.

In part two of the proof, we consider the sequence of $S_{\langle k \rangle}$ of fixed points of f_k . Each $S_{\langle k \rangle}$ is approximately a best-response to the population dynamics, value of search and acceptance decisions $(\mu[S_{\langle k \rangle}], V[S_{\langle k \rangle}], S_{\langle k \rangle})$ to which it gives rise. As $k \rightarrow \infty$, the approximating error becomes small. We show that this sequence admits a pointwise convergent subsequence with limit point $S \in \mathcal{S}^2$. To do so, we draw on a Helly-type selection theorem.

The original Helly's selection theorem asserts that a space of functions defined on an interval that are of uniform bounded total variation is sequentially compact, i.e., sequences in said space admit convergent subsequences.¹⁰ In our non-stationary environment strategy profiles are not one-dimensional. They depend both on type and time. Accordingly, we must draw on a lesser-known multidimensional Helly-type selection theorem. Such result has been proven by Idczak and Walczak (1994), Leonov (1996), Chistyakov and Tretyachenko (2010b) (and references

¹⁰Smith (2006), to prove existence of stationary equilibria, makes use of the original Helly's selection theorem and shows bounded variation of best-response strategy profiles in the steady state (refer to his appendix D). Our proposition 2 (i) generalizes this result to all non-stationary environments.

therein; refer to theorem 4 in the appendix for the relevant version of their result in dimension two). As it turns out, their measure of total variation in dimension two is uniformly bounded if the function admits compact support, is Lipschitz in one argument and of uniform bounded total variation in the other —as we had established in proposition 2 for the value of search. What follows from here is sequential compactness of the best-responses as long as those best-responses are truncated, i.e., defined on a finite time interval only. In particular, there exists a subsequence of $(S_{<k>})_{k \in \mathbb{N}}$ which converges pointwise on $[0, T] \times [0, 1]^2$. Induction over T then establishes pointwise convergence over an infinite time horizon.

To conclude the proof we then show that the best-response to the accumulation point of approximate equilibria, $\bar{S} \equiv f[S]$, is an equilibrium. That is, we show that $f[\bar{S}] = \bar{S}$.

4 Positive Assortative Matching

In this section we derive sufficient conditions for positive assortative matching across all equilibria, encompassing those embedded in non-stationary environments. This completes the analysis by Smith (2006), whose focus is on stationary equilibrium only.

We begin with a definition of positive assortative matching.

Definition 3 (PAM). *There is positive assortative matching (PAM) in equilibrium S at time t , if both $x \mapsto S_t^X(x)$ and $y \mapsto S_t^Y(y)$ are strictly increasing.*

In what follows we briefly provide some context for this definition.

4.1 Definition of positive assortative matching

Positive assortative matching is a sorting condition on equilibrium play: types of similar characteristics are more likely to match with one another. Since matching is random, its outcome can only be ascertained in expectation. It follows that a sensitive definition of positive assortative matching must necessarily be over strategies, not outcomes. Shimer and Smith (2000) propose such a definition. We will refer to their definition as weak PAM and detail the connection hereafter.

Definition 4 (weak PAM, Shimer and Smith (2000)). *There is weak positive assortative matching in equilibrium S at time t , if for types $x_1 < x_2$ and $y_1 < y_2$ we have*

$$m_t[S](x_1, y_2) = 1 \text{ and } m_t[S](x_2, y_1) = 1 \quad \text{imply} \quad m_t[S](x_1, y_1) = 1 \text{ and } m_t[S](x_2, y_2) = 1.$$

In other words, Shimer and Smith (2000) say that matching is positively assortative if, when any two agreeable matches are severed, both the greater two and the lesser two types can be agreeably rematched. This notion is closely related to monotone acceptance thresholds:

Proposition 3 (choosiness \Rightarrow weak PAM). *Suppose that both $x \mapsto S_t^X(x)$ and $y \mapsto S_t^Y(y)$ are non-decreasing. Then there is weak positive assortative matching at time t .*

Proof. Fix $x_1 < x_2$ and $y_1 < y_2$. Suppose that $S \in \mathcal{S}^2$ is such that $m_t[S](x_1, y_2) = m_t[S](x_2, y_1)$ at time t . Then $S_t^X(x_2) \leq y_1$ and $S_t^Y(y_2) \leq x_1$, whence $m_t[S](x_2, y_2) = 1$. And due to weak monotonicity, $S_t^X(x_1) \leq S_t^X(x_2) \leq y_1$ and $S_t^Y(y_1) \leq S_t^Y(y_2) \leq x_1$, whence $m_t[S](x_1, y_1) = 1$. \square

Accordingly, there is weak PAM if the marginal partner is non-decreasing in one's type, a condition commonly referred to as choosiness. Moreover, this makes apparent that PAM implies weak PAM. When matching sets are non-empty the preceding becomes an equivalence however.

Remark 4. *Suppose that matching sets $M_t^X[S](x) \equiv \{y \in [0, 1] : m_t[S](x, y) = 1\}$ and $M_t^Y[S](y) \equiv \{x \in [0, 1] : m_t[S](x, y) = 1\}$ are non-empty throughout. Then the reverse is also true.*

Finally, when populations are symmetric (as in Shimer and Smith (2000) and Smith (2006)), matching sets are non-empty in any equilibrium where there is weak PAM, since identical agent types will always accept one another: $m_t^S(x, x) = 1$. As a consequence, monotone acceptance strategies do not only imply but characterize weak PAM.

Remark 5. *With symmetric populations, there is weak positive assortative matching in equilibrium at time t if and only if both $x \mapsto S_t^X(x)$ and $y \mapsto S_t^Y(y)$ are non-decreasing.*

Our definition of PAM (not weak PAM) can then be best understood in light of remarks 4 and 5. In what follows we will derive sufficient and necessary conditions for PAM, i.e., for $x \mapsto S_t^X(x)$ and $y \mapsto S_t^Y(y)$ to be strictly increasing

4.2 The mimicking argument

To derive equilibrium sorting properties, one needs to compare the value of search across types. Such a comparison is challenging. This stems from the fact that the law of motion is intractable in non-stationary environments, rendering it impossible to characterize the value of search in closed form. To circumvent this problem, we introduce what we call the *mimicking argument*. This relies on two observations.

First, recall that the value of search admits an integral representation over payoffs which subsumes the time dimension (remark 2):

$$V_t^X[S](x) = \sup_s \int_0^1 \pi^X(y|x) Q_t^X[s|S](y|x) dy \quad \text{and} \quad \int_0^1 Q_t^X[s|S](y|x) dy < 1. \quad (3)$$

Here $Q_t^X[s|S](y|x) dy$ is the discounted probability measure of x forming a match with some agent type y anytime in the future, when agent type x 's individual time-variant acceptance threshold is $s : t \mapsto s_t$ and all other agents accept other agents according to strategy profile S .

The second observation we make establishes a lower bound on the value of search. As payoffs are monotone, an agent that is willing to match with a lower agent type x_1 is also

willing to match with a higher agent type x_2 . Thus, agent type x_2 can match with all the agent types (and possibly even other, more attractive ones) that agent type x_1 is matching with. It follows that

$$\sup_s \int_0^1 \pi^X(y|x_2) Q_t^X[s|S](y|x_2) dy \geq \sup_s \int_0^1 \pi^X(y|x_2) Q_t^X[s|S](y|x_1) dy.$$

In particular, this establishes the desired bound on the value of search:

$$V_t^X[S](x_2) \geq \int_0^1 \pi^X(y|x_2) Q_t^X[S](y|x_1) dy \quad \text{for all } x_2 > x_1 \in [0, 1], \quad (4)$$

symmetrically so for population Y .

Another way of thinking about this is to consider a richer strategy space. Suppose strategies take the form of indicator functions, issuing a rejection or acceptance decision for each pair of types. Thus, agents are no longer confined to playing cutoff strategies. Not playing a cutoff strategy is obviously a weakly dominated strategy. However, by considering this richer strategy space, superior types can exactly replicate inferior types' match indicator functions. Such deviations must be weakly dominated by the actual value of search.

Consider then a situation in which PAM is upset at some time t . That is, for some population, say X , there exist some types $x_2 > x_1$ and $z \in [0, 1]$ such that $S_t^X(x_1) \geq z \geq S_t^X(x_2)$. In effect,

$$V_t^X[S](x_2) \leq \pi^X(z|x_2) \quad \text{and} \quad V_t^X[S](x_1) \geq \pi^X(z|x_1). \quad (5)$$

Then, combining (3), (3) and (5), we arrive at the following lemma:

Lemma 1 (mimicking argument). *Suppose that $S_t^X(x_2) \leq S_t^X(x_1)$ in some equilibrium S for some types $x_1 < x_2$ and time t . Then for all $z \in [S_t^X(x_2), S_t^X(x_1)]$*

$$\int_0^1 \pi^X(y|x_1) Q_t^X[S](y|x_1) dy \geq \pi^X(z|x_1) \quad \text{and} \quad \int_0^1 \pi^X(y|x_2) Q_t^X[S](y|x_1) dy \leq \pi^X(z|x_2),$$

symmetrically so for population Y .

Lemma 1 is the keystone of our proofs for the sorting results in theorems 2 and 3.

4.3 Stationary environment

We now introduce a condition that is sufficient for PAM in stationary environments. This condition is log supermodularity. It states that higher types stand relatively more to gain from matching with higher types.

Definition 5 (Log supermodularity). *Payoffs are said to satisfy log supermodularity for popu-*

lation X if for all $y_1 < y_2$ and $x_1 < x_2$

$$\frac{\pi^X(y_2|x_2)}{\pi^X(y_1|x_2)} > \frac{\pi^X(y_2|x_1)}{\pi^X(y_1|x_1)},$$

symmetrically so for population Y .

If the inequality holds with reverse sign we say that payoffs satisfy log submodularity. The following result is due to Smith (2006).

Theorem 2 (stationary PAM, Smith (2006)). *Suppose that payoffs are log supermodular for both populations. Then there is positive assortative matching (PAM) at all times in any stationary equilibrium.*

Smith's original proof, motivated by the analysis of block segregation, proceeds recursively from the highest type to the lowest type. We present here a shorter proof, based on lemma 1, that makes explicit the link between stationarity and log supermodularity.

Proof. The proof proceeds by contraposition, as suggested by the preceding lemma 1. Let $x_2 > x_1$ such that $S_t^X(x_2) \leq S_t^X(x_1)$. (A symmetric reasoning applies for population Y .) By virtue of stationarity, agents always match with a weakly better type than the one rejected initially. Indeed, if an agent currently rejects another type in a stationary environment, she must reject such a type forever. Formally, $Q_t^X[S](y|x_1) = 0$ for all $y < S_t^X(x_1)$. Then, for any $z \in [S_t^X(x_2), S_t^X(x_1)]$, the preceding lemma implies

$$\int_z^1 \frac{\pi^X(y|x_1)}{\pi^X(z|x_1)} Q_t^X[S](y|x_1) dy \geq \int_z^1 \frac{\pi^X(y|x_2)}{\pi^X(z|x_2)} Q_t^X[S](y|x_1) dy. \quad (6)$$

This contradicts the fact that payoffs are log supermodular. \square

The proof equally admits a partial interpretation: if payoffs were log supermodular for population X only, then $x \mapsto S_t^X(x)$ is strictly increasing, symmetrically so for population Y .

It is also worthwhile to note that theorem 2 extends beyond stationary environments. The proof only requires that agents unanimously perceive the economy as being on a weak upward trajectory. That is, at no time in the future do acceptance thresholds fall below the highest type currently rejected. More generally, this implies that *log supermodularity is sufficient to establish PAM whenever strategies are weakly increasing in time*. The steady state emerges as the knife-edge case of weakly increasing strategies: here, acceptance thresholds do not change at all.

Even more can be said when populations are symmetric and the lowest type is concerned.

Corollary 1. *Suppose that populations are symmetric. When payoffs are log supermodular, the lowest type $x = 0$ will indiscriminately accept everyone, even when the economy is non-stationary.*

Proof. Omit superscripts. Suppose by contraposition that $S_t(0) > 0$ for all t in some time interval $[0, t_0)$. Without loss of generality denote $[0, t_0)$ the maximal interval, that is either

$t_0 = \infty$ or $V_{t_0}[S](0) = \pi(0|0)$ and $S_{t_0}(0) = 0$. Then there exists some $t \geq 0$ and some $x_2 > 0$ such that $S_t(x_2) = 0$. (For otherwise the lowest type would never get to match with anyone but her own (atomless) type. This would leave her value of search $V_t[S](0)$ equal to zero, in spite of her initial rejection of the lowest type implying $V_t[S](0) > \pi(0, 0) > 0$.) Thus denote t_1 the minimal time $t \geq 0$ at which $S_{t_1}(x_2) = 0$ for some $x_2 > 0$. If $t_1 \geq t_0$, then the lowest type does not match with anyone during time interval $[0, t_0)$, and thus $V_0[S](0) = e^{-\rho t_0} V_{t_0}[S](0) < \pi(0, 0)$, in contradiction to $S_0(0) > 0$. Thus, suppose that $t_1 < t_0$. In this case we can apply identical arguments as in the proof of theorem 2, noting that $x_1 = z = 0$ to arrive at

$$\int_0^1 \frac{\pi(x'|0)}{\pi(0,0)} Q_t[S](x'|0) dx' \geq \int_0^1 \frac{\pi(x'|x_2)}{\pi(0|x_2)} Q_t[S](x'|0) dx'.$$

As in the main theorem, this contradicts the fact that payoffs are log supermodular. \square

4.4 Non-stationary environments

In a non-stationary environment, log supermodularity is insufficient to guarantee PAM. Here, unlike in the steady state, the lowest type accepted today need not be the worst possible match outcome for all future times. As the search pool evolves over time, agents may face a less favorable selection of types to match with in the future; an agent who rejects a given type initially may accept an inferior type at a later stage. This requires an agent to weigh the current acceptance decision against both the upside risk of matching with a superior and the downside risk of ending up with an inferior type in the future. Log supermodularity is mute on this trade-off. On the one hand, payoff log supermodularity implies that higher types relatively better like to be matched with higher types. On the other hand, it stipulates that higher types stand more to lose from matching with a lower type. Depending on which effect dominates, higher or lower types are choosier. In particular, higher type's fear of the worst outcome may upset PAM, even though payoffs are log supermodular (as illustrated in the following example).

Example 2 (PAM fails without LSD in a gradually clearing search pool). *We construct here an example where PAM is upset despite log supermodular payoffs. Populations are symmetric. The market gradually clears with no entrants joining the search pool. Assuming quadratic search, meetings are less and less likely to occur over time. For simplicity we focus on three types, $x_3 > x_2 > x_1$ such that the lowest type x_1 is accepting of all agent types throughout. Further payoffs are such that the high type x_3 is highly averse of matching with the lowest type x_1 . The mediocre type, by contrast, is almost indifferent between the lesser two types. The example is illustrated in Figure 3, with time on the x-axis and the value of search on the y-axis. Both x_2 and x_3 are willing to match with a mediocre type x_2 as soon as their value of search falls below the horizontal line.¹¹*

Since the market gradually clears over time, the high type x_3 eventually accepts to match with a mediocre type x_2 . Anticipating the possibility of matching with the highest type x_3 , the

¹¹The equilibrium is constructed backwards, starting with an almost empty search pool far into the future. Payoffs are normalized, $\pi(x_2|x_2) = \pi(x_2|x_3) = 1$, and satisfy log supermodularity. Further parameter values are $\pi(x_1|x_2) = 97/100$, $\pi(x_3|x_2) = 10$, $\pi(x_1|x_3) = 1/100$, $\pi(x_3|x_3) = 10 + 1/1000$, and $\rho = 0.9999$




Figure 3: PAM is upset with log supermodular payoffs only.

mediocre type x_2 experiences a surge in her value of search (shortly beginning before time t_0). In the specific example considered, this leaves her not only to reject the lowest, but also her own type. In particular, from time t_1 to t_2 , the highest type finds the mediocre type acceptable, whereas the mediocre type does not. This upsets PAM.

The main contribution of this paper is to establish sufficient conditions for which PAM obtains away from the steady state. First, a definition is in place.

Definition 6. *Payoffs are said to satisfy log supermodularity in differences for population X if for all $y_1 < y_2 < y_3$ and $x_1 < x_2$*

$$\frac{\pi^X(y_3|x_2) - \pi^X(y_2|x_2)}{\pi^X(y_2|x_2) - \pi^X(y_1|x_2)} \geq \frac{\pi^X(y_3|x_1) - \pi^X(y_2|x_1)}{\pi^X(y_2|x_1) - \pi^X(y_1|x_1)},$$

symmetrically so for population Y .

If the inequality holds with the reverse sign, we say that payoffs satisfy log submodularity in differences. The notion of *log supermodularity in differences*, a terminology that we introduce here, postulates that higher types stand relatively more to gain from matching with a high type than they stand to lose from matching with a low type. Log supermodularity in differences is equivalent to $\partial_y \pi^X(y|x)$ being log supermodular, insofar as such a derivative exists.¹²

Strict log supermodularity in differences is typically a stronger condition than log supermodularity. In particular, if payoffs are normalized to zero, i.e. $\pi^X(0|x) = 0$ for all $x \in [0, 1]$, then log supermodularity in differences implies log supermodularity. Similarly, if $\pi^X(0|x)$ is weakly decreasing in x (or constant for that matter) and preferences are supermodular, then again log supermodularity in differences implies log supermodularity.

We can interpret the match payoff $\pi^X(\cdot|x) \equiv u_x(\cdot)$ as agent type x 's utility function. This affords us an interpretation of log supermodularity in differences in terms of risk preferences. More specifically, Pratt (1964) shows that *given arbitrary $x_2 > x_1$ the following statements are equivalent:*

¹²This is a well known result, refer to proposition 7 in the textbook by Gollier (2004). Our proof (see proposition 9) relaxes differentiability assumptions and is included for pedagogical purposes.

1. Agent type x_1 is more risk-averse than agent type x_2 ; that is, x_1 does not accept a lottery that is rejected by x_2 .
2. For any $y_3 > y_2 > y_1$ we have

$$\frac{u_{x_2}(y_3) - u_{x_2}(y_2)}{u_{x_2}(y_2) - u_{x_2}(y_1)} \geq \frac{u_{x_1}(y_3) - u_{x_1}(y_2)}{u_{x_1}(y_2) - u_{x_1}(y_1)}.$$

The use of this result is twofold. First, it features prominently in the proof of theorem 3. Secondly, it provides a simple interpretation of log supermodularity in differences. It amounts to saying that lesser ranked agent types are also more risk-averse. Here we are dealing with payoffs of course, not utilities. This is why we caution against viewing log supermodularity in differences solely in the guise of risk-aversion. The curvature of π is implied by the specific model in mind. It may consequently be derived from economic fundamentals rather than risk preferences.

Having established the terminology we can now state the main result.

Theorem 3 (non-stationary PAM). *Suppose that payoffs satisfy, for both populations, log supermodularity and log supermodularity in differences. Then there is positive assortative matching (PAM) at all times in any (non-stationary) equilibrium.*

Proof. We proceed, as in the stationary case, by assuming that matching is not strictly assortative. Then there exist two types $x_2 > x_1$ and some population, X say, such that at some time t : $S_t^X(x_2) \leq S_t^X(x_1)$. Let $z \in [S_t^X(x_2), S_t^X(x_1)]$. According to lemma 1, we have

$$\int_0^1 \pi^X(y|x_1) Q_t^X[S](y|x_1) dy \geq \pi^X(z|x_1) \quad \text{and} \quad \int_0^1 \pi^X(y|x_2) Q_t^X[S](y|x_1) dy \leq \pi^X(z|x_2). \quad (\star)$$

Define $w > z$ such that $\int_0^1 Q_t^X[S](y|x_1) dy \pi^X(w|x_1) = \pi^X(z|x_1)$. Such $w \in [0, 1]$ does exist. (Suppose the contrary: then for $w = 1$ we have $\int_0^1 Q_t^X[S](y|x_1) dy \pi^X(1|x_1) < \pi^X(z|x_1)$). However, then $\int_0^1 [\pi^X(y|x_1) - \pi^X(1|x_1)] Q_t^X[S](y|x_1) dy > 0$ due to (\star) . Absurd.) Due to log supermodularity

$$\frac{\pi^X(w|x_2)}{\pi^X(z|x_2)} > \frac{\pi^X(w|x_1)}{\pi^X(z|x_1)} \equiv 1 / \int_0^1 Q_t^X[S](y|x_1) dy \quad \Leftrightarrow \quad \pi^X(z|x_2) < \pi^X(w|x_2) \int_0^1 Q_t^X[S](y|x_1) dy. \quad (\star\star)$$

Then let F be a cdf over $[0, 1]$ defined by $F(y) = \int_0^y Q_t^X[S](y'|x_1) dy' / \int_0^1 Q_t^X[S](y'|x_1) dy'$. F represents a lottery. It follows from (\star) and $(\star\star)$ that

$$\int_0^1 \pi^X(y|x_1) F(y) dy \geq \pi^X(w|x_1) \quad \text{and} \quad \int_0^1 \pi^X(y|x_2) F(y) dy < \pi^X(w|x_2).$$

This means that, when offered the choice between the lottery represented by F and the sure outcome w , type x_1 accepts the lottery that is rejected by agent type x_2 . This runs counter to the characterization of log supermodularity in differences in terms of risk preferences and establishes a contradiction. \square

This proof also admits a partial interpretation: if payoffs satisfy LS and LSD for population X only, then $x \mapsto S_t^X(x)$ is strictly increasing, symmetrically so for population Y .

To the best of our knowledge, this is the first sorting result in the literature on matching with search where the search pool is non-stationary. It is established in utmost generality. In particular, the theorem does not rely on initial conditions, differentiability of payoffs, or functional specifications of the meeting and entry rates.

Discussion. It may come as a surprise that risk preferences do not play as prominent a role in the steady state. After all, the decision to reject a certain match payoff today is a revealed preference for a risky, random match payoff sometime in the future—regardless of whether the environment is stationary or not. Our analysis shows that the randomness of search translates into less risk in the steady state. Indeed, in a stationary world, the lowest type accepted initially constitutes a bound on the worst possible match outcome for all future dates; the prospect of future matches below one's current acceptance threshold does not arise. This renders downside risk a feature of non-stationary dynamics only. In consequence, sorting in the steady state solely relies on a preference ranking over upside risk. Non-stationarity in contrast requires a preference ranking over any kind of lottery, entailing both upside and downside risk.

4.5 Necessity

In theorem 3, PAM obtained due to conditions on payoff curvature. Those are not strictly necessary. As higher types are more likely to be accepted by others, higher types can be choosier, regardless of payoff curvature, as in Becker (1973). Thus, PAM can obtain when payoffs are neither log supermodular nor log supermodular in differences.

In this section, we show that said conditions are sharp nevertheless: if either of the sufficient conditions reverses locally for some interval of types, then there exist circumstances under which PAM is upset. We show that this is particularly true when there is no entry into the search pool, arguably the simplest instance of non-stationary dynamics.

Proposition 4 (weak sufficiency). *Consider an economy with symmetric populations and zero entry and suppose that for all types on some interval $[\underline{x}, \bar{x}] \in [0, 1]$ payoffs satisfy either of the following:*

1. *payoffs are log submodular, or*
2. *payoffs are log submodular in differences;*

then there exist meeting rates λ and an initial search pool μ_0 such that PAM does not obtain for some time preceding the (empty) steady state.

The proof of proposition 4 revolves around two counterexamples. We sketch their construction here.

Counterexample 1. When there are only two types, both agent types face identical match opportunities once the superior type begins to indefinitely accept the mediocre type. If payoffs are log submodular, the mediocre type relatively better likes matching with the higher type.

Then PAM must be precisely upset the moment the high type changes her mind about the mediocre type.

Counterexample 2. We emphasize the role of risk as opposed to time by letting the expected time spent in the search pool become exceedingly small, all the while maintaining the downside risk of matching with the lowest type. To that end, we construct a meeting rate that allows for very frequent meetings initially, but renders meetings with high types extremely rare once its population has fallen below some threshold. As a result, a fraction of high types inevitably matches with the low type, while discounting plays virtually no role here. When payoffs are log submodular in differences, the lesser, mediocre type is less risk-averse than the high type. In such a context, it is straightforward that log submodularity in differences renders the mediocre type relatively choosier, upsetting PAM.

5 Explicit search cost

So far, we have embedded search cost through time discounting (as espoused by Shimer and Smith (2000) and Smith (2006)). In this section, we re-establish sufficient conditions for PAM adopting the other prominent representation of search cost: explicit search cost (see Morgan (1994) and Atakan (2006)). Here, discounting plays no role ($\rho = 0$), and each agent in the search pool pays a flow cost c . Whereas time discounting captures the opportunity cost of time, explicit search cost elevates the act of search to be the critical cost.

As was the case under discounting, this framework has been exclusively studied in the steady state (see Morgan (1994)). In what follows, we broaden the scope of the analysis to consider all equilibria. We show that log supermodularity in differences is as essential to positive assortative matching under explicit search cost as it was under discounting.

Our analysis under explicit search cost greatly resembles formal arguments presented under discounting. We begin by re-stating an adapted version of the mimicking argument that incorporates explicit search cost. As under time-discounting, the value of search admits an integral representation over payoffs, where a symmetric construction applies for population Y throughout.

$$V_t^X[S](x) = \sup_s \int_0^1 \pi^X(y|x) Q_t^X[s|S](y|x) dy - C_t[s|S](x) \quad \text{for all } x \in [0, 1]. \quad (7)$$

Here, $Q_t^X[s|S](y|x) dy$ is, as in section 2, the (undiscounted) probability of x forming a match with some agent type in y anytime in the future when the strategy profile is S , and individual agent type x plays strategy s . (If no s is given, then let $s = S^X(x)$).

Since there is no discounting, $\int_0^1 Q_t^X[s|S](y|x) dy = 1$. (That is, discarding the trivial case where the meeting rate is zero and the strategy profile is such that x rejects every agent type for all but at most a finite amount of time.) And $C_t^X[s|S](x)$ is the expected time that agent type x

spends in the search pool from time t onward, multiplied by the explicit search cost c :

$$C_t^X[s|S](x) = c \int_t^{\infty} (\tau - t) p_{t,\tau}^X[s|S](y|x) d\tau.$$

Following the same reasoning as in the original mimicking argument, we can formulate lemma 1 as follows:

Lemma 1' (mimicking argument). *Suppose that $S_t^X(x_2) \leq S_t^X(x_1)$ in some equilibrium S for some types $x_1 < x_2$ and time t . Then for all $z \in [S_t^X(x_2), S_t^X(x_1)]$*

$$\int_0^1 \pi^X(y|x_1) Q_t^X[S](y|x_1) dy - C_t^X[S](x_1) \geq \pi^X(z|x_1)$$

and

$$\int_0^1 \pi^X(y|x_2) Q_t^X[S](y|x_1) dy - C_t^X[S](x_1) \leq \pi^X(z|x_2),$$

symmetrically so for population Y .

We now introduce a condition that is sufficient for PAM in stationary environments. This condition is supermodularity.

Definition 7 (supermodularity). *Payoffs are said to satisfy supermodularity for population X if for all $y_1 < y_2$ and $x_1 < x_2$*

$$\pi^X(y_2|x_2) + \pi^X(y_1|x_1) > \pi^X(y_1|x_2) + \pi^X(y_2|x_1),$$

symmetrically so for population Y .

Supermodularity is typically a weaker condition than log supermodularity. In particular, if payoffs are weakly increasing in x then log supermodularity implies supermodularity.

The following result is due to Morgan (1994).

Theorem 2' (stationary PAM, Morgan (1994)). *Suppose that payoffs are supermodular for both populations. Then there is positive assortative matching (PAM) at all times in any stationary equilibrium.*

Proof. The proof proceeds by contraposition, as suggested by the preceding lemma 1. Let $x_2 > x_1$ be two types whose strategies satisfy $S_t^X(x_2) \leq S_t^X(x_1)$ at time t . By virtue of stationarity, agents always match with a weakly better type than the one rejected initially. Indeed, if an agent currently rejects another type in a stationary environment, she must reject such a type forever. Formally, $Q_t^X[S](y|x_1) = 0$ for all $y < S_t(x_1)$. Then, for any $z \in [S_t(x_2), S_t(x_1)]$, the preceding lemma implies

$$\int_z^1 \pi^X(y|x_1) Q_t^X[S](y|x_1) dy - \pi^X(z|x_1) \geq \int_z^1 \pi^X(y|x_2) Q_t^X[S](y|x_1) dy - \pi^X(z|x_2).$$

As $Q_t^X[S](y|x_1)dy$ is a probability, it integrates to one. The preceding inequality thus simplifies to

$$\int_z^1 [\pi^X(y|x_1) + \pi^X(z|x_2) - \pi^X(z|x_1) - \pi^X(y|x_2)] Q_t^X[S](y|x_1) dy \geq 0.$$

This contradicts the fact that payoffs are supermodular. \square

Supermodularity is insufficient to guarantee positive assortative matching in non-stationary environments for reasons enunciated in the analysis of search with discounting. Again, log supermodularity in differences turns out to be the required condition to ensure PAM across all equilibria.

Theorem 3' (non-stationary PAM). *Suppose that payoffs satisfy, for both populations, supermodularity and log supermodularity in differences. Then there is positive assortative matching (PAM) at all times in any (non-stationary) equilibrium.*

Proof. We proceed, as in the stationary case, by assuming that matching is not strictly assortative. Then there exist two types $x_2 > x_1$ and some population, X say, such that strategies satisfy $S_t^X(x_2) \leq S_t^X(x_1)$ at some time t . Let $z \in [S_t^X(x_2), S_t^X(x_1)]$. According to lemma 1 we have

$$\int_0^1 \pi^X(y|x_1) Q_t^X[S](y|x_1) dy - C_t^X[S](x_1) \geq \pi^X(z|x_1)$$

and

$$\int_0^1 \pi^X(y|x_2) Q_t^X[S](y|x_1) dy - C_t^X[S](x_1) \leq \pi^X(z|x_2)$$

Define $w > z$ such that $\pi^X(w|x_1) = \pi^X(z|x_1) + C_t^X[S](x_1)$. Such $w \in [0, 1]$ does exist (for $\pi^X(z|x_1) + C_t^X[S](x_1) \leq V_t^X[S](x_1) \leq \pi^X(1|x_1)$; then conclude using the intermediate value theorem). Due to supermodularity,

$$\pi^X(w|x_2) + \pi^X(z|x_1) > \pi^X(z|x_2) + \pi^X(w|x_1) \quad \Leftrightarrow \quad \pi^X(w|x_2) > \pi^X(z|x_2) + C_t^X[S](x_1).$$

It follows that

$$\int_0^1 \pi^X(y|x_1) Q_t^X[S](y|x_1) dy \geq \pi^X(w|x_1) \quad \text{and} \quad \int_0^1 \pi^X(x_2, y) Q_t^X[S](y|x_1) dy < \pi^X(w|x_2),$$

in spite of the posited ranking of risk-preferences. \square

The argument used under time discounting extends to explicit search cost. Observe that if payoffs of one population only satisfy log supermodularity and in differences, then as before this ensures that that population's equilibrium acceptance thresholds are strictly increasing in type.

The key message is that unlike steady state sufficient conditions which differ in between environments with discounting and explicit search cost, log supermodularity in differences ensures PAM in non-stationary equilibrium irrespective of how search cost is modeled.

6 Conclusion

This article studies sorting of heterogeneous agents in a general non-stationary matching model. It demonstrates that the study of sorting need not confine itself to particular examples or stationary environments. We hope that it will inspire future ventures into the study of non-stationary dynamics in related frameworks.

Our analysis reveals an intimate link between the time-variant nature of search frictions and risk preferences. We find that the weakest sufficient conditions for positive assortative matching entail that more desirable individuals are less risk-averse in the sense of Arrow-Pratt. This result, taken together with the empirical evidence, provides a theoretical foundation as to why positive assortative matching may arise.

Appendix

A Set-up: Proofs

A.1 Population Dynamics

We begin by establishing a (standard) bound on the evolution of the search pool.

Lemma 2. *Any solution $t \mapsto (\mu_t^X, \mu_t^Y)$ to the population dynamics (1) satisfies*

$$0 \leq \mu_t^X(x) \leq \bar{\mu}_t^X(x) \equiv (1 + \mu_{t_0}^X(x))e^{(t-t_0)L^\eta} - 1,$$

where $(\mu_{t_0}^X, \mu_{t_0}^Y) \in M^2$; symmetrically so for population Y .

Proof. That $\mu_t^X(x)$ is bounded from above follows from considering population dynamics where entry attains its upper bound and none of the agents in the search pool exits. Formally, this bound is described by $\frac{d\mu_t^X(x)}{dt} \leq L^\eta(1 + \mu_t^X(x))$ with starting value $\mu_{t_0}^X(x)$; whence $\mu_t^X(x) \leq (1 + \mu_{t_0}^X(x))e^{L^\eta(t-t_0)} - 1$. That $\mu_t^X(x)$ is bounded from below by zero follows from observing that entry is non-negative and exit of agent types x scales linearly in $\mu_t^X(x)$. A symmetric construction applies for population Y . \square

We then proceed with the proof of proposition 1.

Proof. Step 1: We equip the space of possible evolutions of the state μ over a finite time interval with a norm.

Equip M with the sup norm, denoted $\|\cdot\|_\infty$ (recall M is the space of measurable, bounded and non-negative functions $h : [0, 1] \rightarrow \mathbb{R}_+$). And denote N the identical space without the requirement that functions be non-negative. It is well-known that $(N, \|\cdot\|_\infty)$ and its closed subspace $(M, \|\cdot\|_\infty)$ are Banach spaces. Finally, denote $I_\delta(t_0)$ the time interval $[t_0, t_0 + \delta)$.

Now define, for a given initial condition $\mu_{t_0} \in M^2$, $\mathcal{M}_\delta(t_0)$ the vector space of continuous mappings $\mu : I_\delta(t_0) \rightarrow N^2$ such that $\mu_t^X(x) \leq \bar{\mu}_t^X(x)$ and $\mu_t^Y(y) \leq \bar{\mu}_t^Y(y)$ (where $\bar{\mu}_t^X(x)$ and $\bar{\mu}_t^Y(y)$ are as in the preceding lemma). We equip $\mathcal{M}_\delta(t_0)$ with the norm

$$\|\mu\|_{\mathcal{M}_\delta(t_0)} = \sup_{t \in I_\delta(t_0)} \max \{ \|\mu_t^X\|_\infty, \|\mu_t^Y\|_\infty \}.$$

$\mathcal{M}_\delta(t_0)$ equipped with the norm $\|\cdot\|_{\mathcal{M}_\delta(t_0)}$ is a Banach space (see for instance Francis (2013), proposition 5.6.)

Step 2: We define a mapping $T : \mathcal{M}_\delta(t_0) \rightarrow \mathcal{M}_\delta(t_0)$ whose non-negative fixed points $\mu \in \mathcal{M}_\delta(t_0)$ correspond to the solutions of (1) within time interval $I_\delta(t_0)$:

$$(T^X \mu, T^Y \mu)(t) = (\mu_{t_0}^X + \int_{t_0}^t f^X(\tau, [\mu_\tau]_+) d\tau, \mu_{t_0}^Y + \int_{t_0}^t f^Y(\tau, [\mu_\tau]_+) d\tau).$$

For given $\mu_t \in N^2$, $[\mu_t]_+ \in M^2$ means $[\mu_t]_+^X(x) = \max\{\mu_t^X(x), 0\}$ and $[\mu_t]_+^Y(y) = \max\{\mu_t^Y(y), 0\}$; further, $f = (f^X, f^Y) : I_\delta(t_0) \times M^2 \rightarrow N^2$ is given by

$$f^X(t, \mu_t)(x) = -\mu_t^X(x) \int_0^1 \mu_t^Y(y) \beta^X(t, \mu_t)(y|x) m_t[S^X, S^Y](x, y) dy + \eta^X(t, \mu_t)(x),$$

symmetrically and throughout so for f^Y .

That non-negative fixed points of T constitute solutions of (1) is immediate. Further, every fixed point of T is non-negative. Indeed for any fixed point μ , whenever $(T^X \mu)(t)(x) = \mu_t^X(x) = 0$ it must be that $f^X(t, [\mu_t]_+) \geq 0$, so that due to continuity in time the initially non-negative size of the search pool can never fall below zero.

We then show that T is well-defined, i.e., that the image of T is in $\mathcal{M}_\delta(t_0)$. First, one shows that $(T^X \mu)(t)(x)$ is bounded from below. This is due to the fact that $f^X(t, [\mu_t]_+)(x)$ is uniformly bounded from below—consider maximal exit and no entry. Secondly, one shows that $T^X(\mu)(x)$ is bounded from above by $\bar{\mu}_t^X(x)$. This is due to an identical construction—no exit and maximal entry—as in the preceding lemma.

Step 3: We show that T is contraction mapping for δ sufficiently small. Whence by the contraction mapping theorem it admits a unique fixed point. To begin with, consider arbitrary $\mu', \mu'' \in \mathcal{M}_\delta(t_0)$. Then

$$\sup_{t \in I_\delta(t_0)} \|(T^X \mu')_t - (T^X \mu'')_t\|_\infty \leq \delta \sup_{t \in I_\delta(t_0)} \|f^X(t, \mu'_t) - f^X(t, \mu''_t)\|_\infty.$$

Expanding upon the difference in flows gives, for all $x \in [0, 1]$ and $t \in I_\delta(t_0)$,

$$\begin{aligned} & |f^X(t, \mu'_t)(x) - f^X(t, \mu''_t)(x)| \leq \\ & |\mu''_t^X(x) - \mu'_t^X(x)| \int_0^1 \mu''_t^Y(y) \beta^X(t, \mu''_t)(y|x) m_t[S^X, S^Y](x, y) dy \\ & + \mu'_t^X(x) \int_0^1 |\mu''_t^Y(y) - \mu'_t^Y(y)| \beta^X(t, \mu''_t)(y|x) m_t[S^X, S^Y](x, y) dy \\ & + \mu'_t^X(x) \int_0^1 \mu'_t^Y(y) |\beta^X(t, \mu''_t)(y|x) - \beta^X(t, \mu'_t)(y|x)| m_t[S^X, S^Y](x, y) dy \\ & + |\eta(t, \mu'_t)(x) - \eta(t, \mu''_t)(x)|. \end{aligned}$$

Now make use of assumption 3 (i)-(iv). It follows that the preceding is bounded by

$$\begin{aligned} & |\mu''_t^X(x) - \mu'_t^X(x)| \|\bar{\mu}_t^Y\|_\infty \bar{\beta} \\ & + \|\mu''_t^Y - \mu'_t^Y\|_\infty \bar{\mu}_t^X(x) \bar{\beta} \\ & + \bar{\mu}_t^X(x) \|\bar{\mu}_t^Y\|_\infty L^\beta \max\{\|\mu''_t^X - \mu'_t^X\|_\infty, \|\mu''_t^Y - \mu'_t^Y\|_\infty\} \\ & + L^\eta \max\{\|\mu''_t^X - \mu'_t^X\|_\infty, \|\mu''_t^Y - \mu'_t^Y\|_\infty\}. \end{aligned}$$

Finally, taking the supremum across types and recalling the explicit bound of $\bar{\mu}_t$ establishes that

$$\begin{aligned} & \|T\mu' - T\mu''\|_{\mathcal{M}_\delta(t_0)} \\ & \leq \delta \left\{ 2 \left((1 + \max \{ \|\mu_{t_0}^X\|_\infty, \|\mu_{t_0}^Y\|_\infty \}) e^{L\delta} - 1 \right) \bar{\beta} \right. \\ & \quad \left. + \left((1 + \max \{ \|\mu_{t_0}^X\|_\infty, \|\mu_{t_0}^Y\|_\infty \}) e^{L\delta} - 1 \right)^2 L^\beta + L^\eta \right\} \|\mu' - \mu''\|_{\mathcal{M}_\delta(t_0)}. \end{aligned}$$

Then choose δ such that for all $\delta' \leq \delta$ the term in the curly brackets is smaller than $\frac{1}{2} \frac{1}{\delta'}$. It follows that $T : \mathcal{M}_\delta(t_0) \rightarrow \mathcal{M}_\delta(t_0)$ is a contraction mapping, whence the Banach fixed point theorem guarantees the existence of a unique fixed point in $\mathcal{M}_\delta(t_0)$. This fixed point is the solution to (1) in $I_\delta(t_0)$.

Two conclude, we make two observations. First, the solution can not depend on the choice of δ . If there existed two solution, by virtue of the preceding argument they must agree on any interval where they are both defined. Secondly, the solution initially established on $I_\delta(t_0)$ extends in fact to $[t_0, \infty)$. Since t_0 was arbitrary, it follows that the solution exists on \mathbb{R}_+ . If this was not the case, $t \mapsto \max \{ \|\mu_t^X\|_\infty, \|\mu_t^Y\|_\infty \}$ would be unbounded on the maximal interval on which a solution exists. Yet, as ascertained in step 2, $\mu_t^X(x) \leq (1 + \mu_{t_0}^X(x)) e^{L(t-t_0)}$ for all $x \in [0, 1]$, likewise for population Y , and this expression is bounded for t in a finite interval. \square

A.2 The value of search

We then proceed with the proof of proposition 2.

Proof. (i) We make use of the dynamic programming principle:

$$\begin{aligned} V_t^X[S](x) = \sup_s \left\{ \int_0^1 \pi^X(y|x) \left[\int_t^{t+h} e^{-\rho(\tau-t)} p_{t,\tau}^X[s|S](y|x) d\tau \right] dy \right. \\ \left. + e^{-\rho h} \left[1 - \int_0^1 \int_t^{t+h} p_{t,\tau}^X[s|S](y|x) d\tau dy \right] V_{t+h}^X[S](x) \right\}. \end{aligned}$$

Equivalently we can write

$$\begin{aligned} \frac{V_{t+h}^X[S](x) - V_t^X[S](x)}{h} = \sup_s \left\{ - \int_0^1 \pi^X(y|x) \frac{1}{h} \left[\int_t^{t+h} e^{-\rho(\tau-t)} p_{t,\tau}^X[s|S](y|x) d\tau \right] dy \right. \\ \left. + \frac{1 - e^{-\rho h}}{h} V_{t+h}^X[S](x) + e^{-\rho h} \frac{1}{h} \left[\int_0^1 \int_t^{t+h} p_{t,\tau}^X[s|S](y|x) d\tau dy \right] V_{t+h}^X[S](x) \right\}. \end{aligned}$$

Note that the term in the curled brackets is finite, whence $t \mapsto V_t^X[S](x)$ is continuous. We use the little- o Landau notation and write $o(1)$, meaning that $\lim_{h \rightarrow 0} o(1) = 0$. Then $V_{t+h}^X[S](x) =$

$V_t^X[S](x) + o(1)$. Further note that

$$\begin{aligned} \frac{1}{h} \left[\int_t^{t+h} e^{-\rho(\tau-t)} p_{t,\tau}^X [s|S](y|x) d\tau \right] dy &= \frac{1}{h} \left[\int_t^{t+h} e^{-\rho(\tau-t)} \exp \left\{ - \int_t^\tau \int_0^1 \Lambda_r^X [s|S](y|x) dy dr \right\} \Lambda_\tau^X [s|S](y|x) d\tau \right] \\ &= \frac{1}{h} \left[\int_t^{t+h} \Lambda_\tau^X [s|S](y|x) d\tau \right] + o(1) \end{aligned}$$

and

$$\frac{1 - e^{-\rho h}}{h} = \rho + o(1).$$

It follows that

$$\frac{V_{t+h}^X[S](x) - V_t^X[S](x)}{h} = \sup_s \left\{ \int_0^1 (V_{t+h}^X[S](x) - \pi^X(y|x)) \frac{1}{h} \int_t^{t+h} \Lambda_\tau^X [s|S](y|x) d\tau dy + \rho V_{t+h}^X[S](x) \right\} + o(1).$$

To conclude, recall from lemma 2 that $\mu_t^X(x) \leq (1 + \mu_0^X(x))e^{L\eta t}$. Denoting $\bar{\mu}_0$ the maximal initial population size for any agent type, it follows that $\Lambda_\tau^X [s|S](y|x) \leq \bar{\beta}(1 + \bar{\mu}_0)e^{L\eta \tau}$. Therefore

$$\left| \frac{V_{t+h}^X[S](x) - V_t^X[S](x)}{h} \right| \leq (\pi^X(1|x) - \pi^X(0|x))\bar{\beta}(1 + \bar{\mu}_0)e^{L\eta t} + \rho\pi^X(1|x) + o(1).$$

Then, for fixed yet finite T , and $0 \leq t_1 < t_2 \leq T$,

$$\left| V_{t_2}^X[S](x) - V_{t_1}^X[S](x) \right| \leq C |t_2 - t_1| \quad \text{where} \quad C = (\pi^X(1|x) - \pi^X(0|x))\bar{\beta}(1 + \bar{\mu}_0)e^{L\eta T} + \rho\pi^X(1|x)$$

as desired.

(ii) Consider an arbitrary partition of the unit interval $[0, 1]$: $0 = x_0 < x_1 < \dots < x_m = 1$. Recall the mimicking argument, namely equation 3, which can be expressed as follows:

$$V_t^X[S](x_i) - V_t^X[S](x_{i-1}) \geq \int_0^1 (\pi^X(y|x_i) - \pi^X(y|x_{i-1})) Q_t^X(y|x_i) dy.$$

Further recall assumption 1 which posits that match payoffs are bounded. Then

$$\begin{aligned}
& \sum_{i=1}^m |V_t^X[S](x_i) - V_t^X[S](x_{i-1})| \\
&= -2 \sum_{i=1}^m \min \{V_t^X[S](x_i) - V_t^X[S](x_{i-1}), 0\} + \sum_{i=1}^m (V_t^X[S](x_i) - V_t^X[S](x_{i-1})) \\
&\leq -2 \sum_{i=1}^m \min \left\{ \int_0^1 (\pi^X(y|x_i) - \pi^X(y|x_{i-1})) Q_t^X(y|x_i) dy, 0 \right\} + \bar{\pi} \\
&\leq 2 \int_0^1 \sum_{i=1}^m |\pi^X(y|x_i) - \pi^X(y|x_{i-1})| Q_t^X(y|x_i) dy + \bar{\pi} \\
&\leq 2 \sup_y \sum_{i=1}^m |\pi^X(y|x_i) - \pi^X(y|x_{i-1})| + \bar{\pi} \\
&\leq 2L^\pi + \bar{\pi}.
\end{aligned}$$

The last inequality is due to assumption 4 which posits that match payoffs are of uniformly bounded total variation; L^π denotes the uniform bound (for all y) on total variation. \square

B Existence of an equilibrium

B.1 Compactness of strategy subspaces

Recall that \mathcal{S} is the set of (Lebesgue measurable) population strategies for either population X or Y . Throughout, Let $S \in \mathcal{S}$ a population strategy (we omit superscripts in this subsection). We endow \mathcal{S} with a metric \mathbf{d} :

$$\mathbf{d}(S, S') = \sup_{x,t} e^{-t} |S_t(x) - S'_t(x)|.$$

A cutoff strategy S is k -Lipschitz if for any $(x, t), (y, r) \in [0, 1] \times [0, \infty)$ we have

$$|S_t(x) - S_r(y)| \leq k \cdot \max \{|x - y|, |t - r|\}.$$

Denote $\mathcal{S}_k \subset \mathcal{S}$ the (convex) subspace of k -Lipschitz cutoff strategies. Observe that \mathcal{S} is the closure of $\cup_{k \geq 1} \mathcal{S}_k$, the union of k -Lipschitz cutoff strategy subspaces.

Proposition 5. $(\mathcal{S}_k, \mathbf{d})$ is compact.

Proof. We show that $(\mathcal{S}_k, \mathbf{d})$ is complete and totally bounded. This establishes compactness (see for instance Munkres (2015), theorem 45.1, p. 274). We focus on completeness first. Let $(S^n)_{n \in \mathbb{N}}$ with each $S^n \in \mathcal{S}_k$ a Cauchy sequence. Then for each $(x, t) \in [0, 1] \times [0, \infty)$ the sequence $(S_t^n(x))_{n \in \mathbb{N}}$ converges as $n \rightarrow \infty$. Denote $S_t(x)$ its pointwise limit and S the thereby obtained strategy. We first show that $S \in \mathcal{S}_k$. Fix arbitrary $\epsilon > 0$ and $(x, t), (y, r) \in [0, 1] \times [0, \infty)$. Due to

pointwise convergence there exists $N \in \mathbb{N}$ such that for all $n \geq N$ we have

$$\max \left\{ |S_t^n(x) - S_t(x)|, |S_r^n(y) - S_r(y)| \right\} < \frac{\epsilon}{2}.$$

It follows from the triangle inequality and k -Lipschitz continuity of S^n that

$$\begin{aligned} |S_t(x) - S_r(y)| &\leq |S_t(x) - S_t^N(x)| + |S_t^N(x) - S_r^N(y)| + |S_r^N(y) - S_r(y)| \\ &< \epsilon + k \cdot \max \{|x - y|, |t - r|\}. \end{aligned}$$

Since $\epsilon > 0$ was arbitrary this establishes that $S \in \mathcal{S}_k$. We then show that $S^n \rightarrow S$ in the \mathbf{d} -metric. Again fix arbitrary $\epsilon > 0$. If for any given $n \in \mathbb{N}$ the sup were attained for some $t > T$ where $e^{-T} < \epsilon$, clearly $\mathbf{d}(S^n, S) < \epsilon$. Let's then focus our attention on the case $(x, t) \in [0, 1] \times [0, T]$. Define $B_k^\epsilon(x, t) = \{(y, r) : \max \{|x - y|, |t - r|\} < \frac{\epsilon}{4k}\}$ and let $N \in \mathbb{N}$ such that for all $n \geq N$ we have $|S_t^n(x) - S_t(x)| < \frac{\epsilon}{2}$. Then for any $(y, r) \in B_k^\epsilon(x, t)$

$$\begin{aligned} |S_r^n(y) - S_r(y)| &\leq |S_r^n(y) - S_t^n(x)| + |S_t^n(x) - S_t(x)| + |S_t(x) - S_r(y)| \\ &< 2k \max \{|x - y|, |t - r|\} + \frac{\epsilon}{2} < \epsilon. \end{aligned}$$

Finally observe that the set $\{B_k^\epsilon(x, t) : (x, t) \in [0, 1] \times [0, T]\}$ forms an open covering of the compact space $[0, 1] \times [0, T]$. Whence there exists a finite subcovering of that space, $\{B_k^\epsilon(x_j, t_j) : j \in \{1, \dots, M\}\}$. For any $j \in \{1, \dots, M\}$ let N_j such that for all $n \geq N_j$ we have $|S_{t_j}^n(x_j) - S_{t_j}(x_j)| < \frac{\epsilon}{2}$. Then it follows from the preceding arguments that for all $n \geq N \equiv \max \{N_j : j \in \{1, \dots, M\}\}$ we have $\mathbf{d}(S^n, S) < \epsilon$. This establishes completeness.

Let's now focus attention on total boundedness. That is, for every $\epsilon > 0$ there exists a finite number of strategies $S^j \in \mathcal{S}$ such that for all $S \in \mathcal{S}_k$ we have $\mathbf{d}(S^j, S) < \epsilon$ for some $j \in \{1, \dots, M\}$. We achieve this by choosing a partition \mathcal{R}^ϵ of $[0, 1]$ as well as a partition \mathcal{P}^ϵ of $[0, 1] \times [0, T]$ for some $T > 0$ such that $e^{-T} < \epsilon$. In particular, let $\mathcal{R}^\epsilon = \{0, \epsilon, \dots, l^\epsilon \epsilon\}$ where $l^\epsilon \epsilon \leq 1 < (l^\epsilon + 1)\epsilon$ and $\mathcal{P}^\epsilon = \{(\frac{m}{k} \frac{\epsilon}{2}, \frac{n}{k} \frac{\epsilon}{2}) : m, n \in \{0, \dots, m^\epsilon\} \times \{0, \dots, n^\epsilon\}\}$ where $\frac{m^\epsilon}{k} \frac{\epsilon}{2} \leq 1 < \frac{m^\epsilon + 1}{k} \frac{\epsilon}{2}$ and $\frac{n^\epsilon}{k} \frac{\epsilon}{2} \leq T < \frac{n^\epsilon + 1}{k} \frac{\epsilon}{2}$. We then consider the (finite) space of functions $\mathcal{F}^\epsilon = \{h : \mathcal{P}^\epsilon \rightarrow \mathcal{R}^\epsilon\}$. Let h an element in this space. The corresponding strategy S^h is defined pointwise where $S_t^h(x) = h(\frac{m}{k} \frac{\epsilon}{2}, \frac{n}{k} \frac{\epsilon}{2})$ for $(x, t) \in [\frac{m}{k} \frac{\epsilon}{2}, \frac{m+1}{k} \frac{\epsilon}{2}] \times [\frac{n}{k} \frac{\epsilon}{2}, \frac{n+1}{k} \frac{\epsilon}{2}]$. Denote $\mathcal{S}_k^\epsilon \equiv \{S^h \in \mathcal{S} : h \in \mathcal{F}^\epsilon\}$. ϵ -proximity of \mathcal{S}_k to \mathcal{S}_k^ϵ then follows immediately: for arbitrary $S \in \mathcal{S}_k$ there exists $h \in \mathcal{F}^\epsilon$ such that for all $(y, \tau) \in \mathcal{P}^\epsilon$ we have $|S_\tau(y) - h(y, \tau)| \leq \frac{\epsilon}{2}$. Then consider any $(x, t) \in [0, 1] \times [0, T]$. Let (x^ϵ, t^ϵ) the greatest element in \mathcal{P}^ϵ such that $x^\epsilon \leq x$ and $t^\epsilon \leq t$. Then by construction $S_{t^\epsilon}^h(x^\epsilon) = S_t^h(x)$ and $\max \{|x - x^\epsilon|, |t - t^\epsilon|\} \leq \frac{1}{k} \frac{\epsilon}{2}$. Using the triangle inequality and the fact that S is k -Lipschitz continuous we obtain

$$\begin{aligned} |S_t(x) - S_t^h(x)| &\leq |S_t(x) - S_{t^\epsilon}(x^\epsilon)| + |S_{t^\epsilon}(x^\epsilon) - S_t^h(x)| \\ &= |S_t(x) - S_{t^\epsilon}(x^\epsilon)| + |S_{t^\epsilon}(x^\epsilon) - S_{t^\epsilon}^h(x^\epsilon)| \\ &\leq k \max \{|x - x^\epsilon|, |t - t^\epsilon|\} + \frac{\epsilon}{2} \leq \epsilon. \end{aligned}$$

As (x, t) was arbitrary, this bound holds uniformly across $[0, 1] \times [0, T]$. Meanwhile, for $t > T$

ϵ -closeness is satisfied vacuously, whence the result. \square

B.2 Continuity a.e. of the value function

In what follows we will prove continuity a.e. of the value of search in strategy profiles of both populations. Accordingly, define $\mathbf{d}^2 : \mathcal{S}^2 \times \mathcal{S}^2 \rightarrow \mathbb{R}_+$:

$$\mathbf{d}^2((S^X, S^Y), (S'^X, S'^Y)) = \max \{ \mathbf{d}(S^X, S'^X); \mathbf{d}(S^Y, S'^Y) \}.$$

Clearly, $(\mathcal{S}^2, \mathbf{d}^2)$ is a metric space. To prove continuity of the value of search, we require two lemmata.

Lemma 3. Fix $\bar{S} \in \mathcal{S}^2$. Then for all $t \in [0, \infty)$ and almost every $x \in [0, 1]$: for all $\epsilon > 0$ there exists $\delta > 0$ such that

$$\int_0^1 |m_t[S](x, y) - m_t[\bar{S}](x, y)| dy < \epsilon \quad \text{for all } S : \mathbf{d}^2(S, \bar{S}) < \delta.$$

Proof. Fix \bar{S} and arbitrary $t \in [0, \infty)$. Let $x \in [0, 1]$ such that $\{y \in [0, 1] : \bar{S}_t^Y(y) = x\}$ is negligible. Note that almost every $x \in [0, 1]$ is an admissible candidate.

Fix $\epsilon > 0$. And let $\delta_1 > 0$ such that $\{y \in [0, 1] : |\bar{S}_t^Y(y) - x| < \delta_1\}$ has Lebesgue measure smaller than $\frac{\epsilon}{2}$. Then for all $S : \mathbf{d}^2(S, \bar{S}) < e^{-t}\delta_1$

$$\int_0^1 |\mathbb{1}\{x \geq S_t^Y(y)\} - \mathbb{1}\{x \geq \bar{S}_t^Y(y)\}| dy = \int_{\substack{y: \bar{S}_t^Y(y) \in \\ (x-\delta_1, x+\delta_1)}} |\cdot| dy \leq \int_{\substack{y: \bar{S}_t^Y(y) \in \\ (x-\delta_1, x+\delta_1)}} dy \leq \frac{\epsilon}{2}.$$

Next, choose $\delta_2 \equiv \frac{\epsilon}{4}$. Then for all $S : \mathbf{d}^2(S, \bar{S}) < e^{-t}\delta_2$

$$\int_0^1 |\mathbb{1}\{y \geq S_t^X(x)\} - \mathbb{1}\{y \geq \bar{S}_t^X(x)\}| dy = \int_{\bar{S}_t^X(x)-\delta_3}^{\bar{S}_t^X(x)+\delta_3} |\cdot| dy \leq \int_{\bar{S}_t^X(x)-\delta_3}^{\bar{S}_t^X(x)+\delta_3} dy \leq 2\delta_3 = \frac{\epsilon}{2}.$$

Finally, set $\delta \equiv e^{-t} \min\{\delta_1, \delta_2\}$. Since

$$|m_t[S](x, y) - m_t[\bar{S}](x, y)| \leq |\mathbb{1}\{x \geq S_t^Y(y)\} - \mathbb{1}\{x \geq \bar{S}_t^Y(y)\}| + |\mathbb{1}\{y \geq S_t^X(x)\} - \mathbb{1}\{y \geq \bar{S}_t^X(x)\}|,$$

the asserted $\epsilon - \delta$ -argument is satisfied. \square

Lemma 4. Fix $\bar{S} \in \mathcal{S}^2$. Then for all $t \in [0, \infty)$: for all $\epsilon > 0$ there exists $\delta > 0$ such that

$$\int_0^1 |\mu_t^X[S](x) - \mu_t^X[\bar{S}](x)| dx < \epsilon \quad \text{for all } S : \mathbf{d}^2(S, \bar{S}) < \delta;$$

symmetrically so for population Y.

Proof. Fix $t > 0$. Consider arbitrary $\bar{S} \in \mathcal{S}^2$. Manipulating the integral equation (1) governing the evolution of the search pool, we obtain

$$\begin{aligned}
\mu_t^X[S](x) - \mu_t^X[\bar{S}](x) &= \int_0^t \left\{ \mu_\tau^X[\bar{S}](x) \int_0^1 \beta^X(\tau, \mu_\tau[\bar{S}])(y|x) \mu_\tau^Y[\bar{S}](y) m_\tau[\bar{S}](x, y) dy \right. \\
&\quad \left. - \mu_\tau^X[S](x) \int_0^1 \beta^X(\tau, \mu_\tau[S])(y|x) \mu_\tau^Y[S](y) m_\tau[S](x, y) dy + \eta^X(\tau, \mu_\tau[S])(x) - \eta^X(\tau, \mu_\tau[\bar{S}])(x) \right\} d\tau \\
&= \int_0^t \left\{ (\mu_\tau^X[\bar{S}](x) - \mu_\tau^X[S](x)) \int_0^1 \beta^X(\tau, \mu_\tau[\bar{S}])(y|x) \mu_\tau^Y[\bar{S}](y) m_\tau[\bar{S}](x, y) dy \right. \\
&\quad + \mu_\tau^X[S](x) \int_0^1 (\beta^X(\tau, \mu_\tau[\bar{S}])(y|x) - \beta^X(\tau, \mu_\tau[S])(y|x)) \mu_\tau^Y[\bar{S}](y) m_\tau[\bar{S}](x, y) dy \\
&\quad + \mu_\tau^X[S](x) \int_0^1 \beta^X(\tau, \mu_\tau[S])(y|x) (\mu_\tau^Y[\bar{S}](y) - \mu_\tau^Y[S](y)) m_\tau[\bar{S}](x, y) dy \\
&\quad + \mu_\tau^X[S](x) \int_0^1 \beta^X(\tau, \mu_\tau[S])(y|x) \mu_\tau^Y[S](y) (m_\tau[\bar{S}](x, y) - m_\tau[S](x, y)) dy \\
&\quad \left. + \eta^X(\tau, \mu_\tau[S])(x) - \eta^X(\tau, \mu_\tau[\bar{S}])(x) \right\} d\tau.
\end{aligned}$$

Using Lipschitz continuity of η and β , we get

$$\begin{aligned}
\int_0^1 |\mu_t^X[S](x) - \mu_t^X[\bar{S}](x)| dx &\leq (1 + \bar{\mu}_0) e^{tL\bar{\beta}} \int_0^t \int_0^1 |\mu_\tau^X[\bar{S}](x) - \mu_\tau^X[S](x)| dx d\tau \\
&\quad + (1 + \bar{\mu}_0)^2 e^{2tL\bar{\beta}} L^\beta \int_0^t N(\mu_\tau[S], \mu_\tau[\bar{S}]) d\tau + (1 + \bar{\mu}_0) e^{tL\bar{\beta}} \bar{\beta} \int_0^t N(\mu_\tau[S], \mu_\tau[\bar{S}]) d\tau \\
&\quad + (1 + \bar{\mu}_0)^2 e^{2tL\bar{\beta}} \bar{\beta} \int_0^t \int_0^1 \int_0^1 |m_\tau[\bar{S}](x, y) - m_\tau[S](x, y)| dy dx d\tau \\
&\quad + L^\eta \int_0^t N(\mu_\tau[S], \mu_\tau[\bar{S}]) d\tau. \tag{*}
\end{aligned}$$

To prove the lemma it is sufficient to show that for all $\epsilon > 0$ there exists $\delta > 0$ such that

$$N(\mu_t[S], \mu_t[\bar{S}]) < \epsilon \quad \forall S : \mathbf{d}^2(S, \bar{S}) < \delta.$$

To make progress we establish a bound on the difference in match indicator functions: fix some

$\xi > 0$ (to be determined) and consider a sequence $(\sigma_n)_{n \in \mathbb{N}}$ with $\sigma_n > 0$ and $\sigma_n \rightarrow 0$. Define

$$A^X(\sigma_n, \xi) = \left\{ x \in [0, 1] : \int_0^t \int_0^1 |m_\tau[\bar{S}](x, y) - m_\tau[S](x, y)| dy d\tau < \xi \text{ for all } S : \mathbf{d}^2(S, \bar{S}) < \sigma_n \right\}$$

and

$$a^X(\sigma_n, \xi) \equiv \int_{[0,1] \setminus A^X(\sigma_n, \xi)} dy.$$

According to the preceding lemma, $a^X(\sigma_n, \xi) dy \rightarrow 0$ as $n \rightarrow \infty$. Whence there exists N such that $a^X(\sigma_N, \xi) \leq \xi$. This establishes that for all $S : \mathbf{d}^2(S, \bar{S}) < \sigma_N$:

$$\begin{aligned} & \int_0^t \int_0^1 \int_0^1 |m_\tau[\bar{S}](x, y) - m_\tau[S](x, y)| dy dx d\tau \\ & \leq \underbrace{\int_{A^X(\sigma_N, \xi)} \int_0^t \int_0^1 |m_\tau[\bar{S}](x, y) - m_\tau[S](x, y)| dy d\tau dx}_{< \xi} + \underbrace{\int_{[0,1] \setminus A^X(\sigma_N, \xi)} dx}_{= a^X(\sigma_N, \xi) < \xi} < 2\xi. \quad (\star\star) \end{aligned}$$

From here we conclude: inequalities (\star) and $(\star\star)$ jointly imply that

$$N(\mu_t[S], \mu_t[\bar{S}]) \leq \underbrace{(2(1 + \bar{\mu}_0)e^{tL^\eta} \bar{\beta} + (1 + \bar{\mu}_0)^2 e^{2tL^\eta} L^\beta + L^\eta)}_{\equiv K_1} \int_0^t N(\mu_\tau[S], \mu_\tau[\bar{S}]) d\tau + \underbrace{2(1 + \bar{\mu}_0)^2 e^{2tL^\eta} \bar{\beta} 2 \xi}_{\equiv K_2}$$

for all $S : \mathbf{d}^2(S, \bar{S}) < \sigma_N$. And an application of Grönwall's lemma gives

$$N(\mu_t[S], \mu_t[\bar{S}]) \leq K_1 \xi e^{K_2 t}.$$

Then to satisfy the $\epsilon - \delta$ argument, choose $\xi \equiv \frac{\epsilon}{K_1} e^{-K_2 t}$ and set δ equal to the associated σ_N \square

Proposition 6. Fix $\bar{S} \in \mathcal{S}^2$. Then for all $t \in [0, \infty)$ and almost all $x \in [0, 1]$, $S \mapsto V_t^X[S](x)$ is continuous at \bar{S} , symmetrically so for population Y .

Proof. We begin with a preliminary remark. Observe that

$$|V_t^X[S'](x) - V_t^X[S''](x)| \leq \max \{ |W_t^X[s'|S'](x) - W_t^X[s'|S''](x)|, |W_t^X[s''|S'](x) - W_t^X[s''|S''](x)| \},$$

where s', s'' are such that $W_t^X[s'|S'](x) = V_t^X[S'](x)$ and $W_t^X[s''|S''](x) = V_t^X[S''](x)$. So it is sufficient to show that for all s and almost all $(x, t) \in [0, 1] \times [0, \infty)$, $S \mapsto W_t^X[s|S](x)$ is continuous at \bar{S} .

We next determine the subset of admissible types and time for whom continuity holds. Fix $\bar{S} \in \mathcal{S}^2$. Pick arbitrary t_0 . Pick x_0 such that $\int_0^\infty \int_0^1 \mathbb{1}\{\bar{S}_t^Y(y) = x_0\} dy d\tau$ is zero. We remark that

almost every x is an admissible choice. Indeed, denote

$$A_k = \{x \in [0, 1] : \int_0^k \int_0^1 \mathbb{1}\{\bar{S}_\tau(y) = x\} dy d\tau \geq \frac{1}{k}\}.$$

Then any $x \in [0, 1] \setminus \cup_{k \in \mathbb{N}} A_k$ is an admissible choice for x_0 . And since each A_k is finite—containing at most k^2 many types— $\cup_{k \in \mathbb{N}} A_k$ has zero measure. Further note that lemma 1 applies for any admissible x at almost every moment in time τ .

We then show continuity of $S \mapsto W_{t_0}^X[s|S](x_0)$ at \bar{S} . Fix $\epsilon > 0$. Let $t_1 \geq t_0$ such that $e^{t_1 \rho} \bar{\pi} \leq \frac{\epsilon}{4}$. And define for $t \leq t_1$ the time 0-discounted, time t_1 -truncated continuation value of strategy profile $s|S$:

$$\hat{W}_t[s|S](x) = e^{-\rho t} \int_0^1 \pi^X(y|x) \int_t^{t_1} e^{-\rho(\tau-t)} p_{t_0,t}^X[s|S](y|x) d\tau dy.$$

Due to the triangle inequality we have

$$|W_{t_0}[s|S](x_0) - W_{t_0}[s|\bar{S}](x_0)| \leq e^{t_0 \rho} |\hat{W}_{t_0}[s|S](x_0) - \hat{W}_{t_0}[s|\bar{S}](x_0)| + \frac{\epsilon}{2}.$$

Further note that

$$\frac{\hat{W}_{t+h}[s|S](x) - \hat{W}_t[s|S](x)}{h} = \frac{1}{h} \int_t^{t+h} \int_0^1 (\hat{W}_\tau^X[s|S](x) - e^{-\rho\tau} \pi^X(y|x)) \Lambda_\tau^X[s|S](y|x) dy d\tau + o(1).$$

and $\hat{W}_{t_1}[s|S](x) = 0$. To see this, employ identical dynamic programming as in the proof of proposition 2 (i):

$$\begin{aligned} \frac{\hat{W}_{t+h}[s|S](x) - \hat{W}_t[s|S](x)}{h} &= - \int_0^1 \pi^X(y|x) e^{-\rho t} \left[\frac{1}{h} \int_t^{t+h} e^{-\rho(\tau-t)} p_{t,\tau}^X[s|S](y|x) d\tau \right] dy \\ &\quad + \frac{e^{-\rho(t+h)}}{h} \left(W_{t+h}^X[s|S](x) - \left[1 - \frac{1}{h} \int_t^{t+h} p_{t,\tau}^X[s|S](y|x) d\tau \right] W_{t+h}^X[s|S](x) \right). \end{aligned}$$

We next derive a bound for

$$\begin{aligned} &|\hat{W}_{t_0}[s|S](x_0) - \hat{W}_{t_0}[s|\bar{S}](x_0)| \\ &= \left| \int_{t_0}^{t_1} \int_0^1 \left\{ (\hat{W}_\tau^X[s|\bar{S}](x_0) - e^{-\rho\tau} \pi^X(y|x_0)) \Lambda_\tau^X[s|\bar{S}](y|x_0) \right. \right. \\ &\quad \left. \left. - (\hat{W}_\tau^X[s|S'](x_0) - e^{-\rho\tau} \pi^X(y|x_0)) \Lambda_\tau^X[s|S'](y|x_0) \right\} dy d\tau \right|. \end{aligned}$$

First recall that $\Lambda_\tau^X[s|S](y|x_0) = \beta^X(\tau, \mu_\tau[S])(y|x_0) \cdot \mu_\tau^Y[S](y) \cdot m_\tau[s|S](x_0, y)$. Expanding, we can

then bound the preceding by

$$\begin{aligned}
&\leq \int_{t_0}^{t_1} \int_0^1 \left\{ \left| \hat{W}_\tau^X[s|\bar{S}](x_0) - e^{-\rho\tau} \pi^X(y|x_0) \right| \cdot \lambda^X(\tau, \mu_\tau[\bar{S}])(y|x_0) \right. \\
&\quad \left. \cdot \left| m_\tau[s|\bar{S}](x_0, y) - m_\tau[s|S'](x_0, y) \right| \right\} dy d\tau \\
&+ \int_{t_0}^{t_1} \int_0^1 \left\{ \left| \hat{W}_\tau^X[s|\bar{S}](x_0) - e^{-\rho\tau} \pi^X(y|x_0) \right| \cdot \mu_\tau^Y[\bar{S}](y) \right. \\
&\quad \left. \cdot \left| \beta^X(\tau, \mu_\tau[\bar{S}])(y|x_0) - \beta^X(\tau, \mu_\tau[S'])(y|x_0) \right| \cdot m_\tau[s|S'](x_0, y) \right\} dy d\tau \\
&+ \int_{t_0}^{t_1} \int_0^1 \left\{ \left| \hat{W}_\tau^X[s|\bar{S}](x_0) - e^{-\rho\tau} \pi^X(y|x_0) \right| \cdot \left| \mu_\tau^Y[\bar{S}](y) - \mu_\tau^Y[S'](y) \right| \right. \\
&\quad \left. \cdot \beta^X(\tau, \mu_\tau[S'])(y|x_0) \cdot m_\tau[s|S'](x_0, y) \right\} dy d\tau \\
&+ \int_{t_0}^{t_1} \int_0^1 \left\{ \left| \hat{W}_\tau^X[s|\bar{S}](x_0) - \hat{W}_\tau^X[s|S'](x_0) \right| \cdot \Lambda_\tau^X[s|S'](y|x_0) \right\} dy d\tau.
\end{aligned}$$

And, due to lemma 2 and assumptions 1 and 3, this is smaller than

$$\begin{aligned}
&\leq \bar{\pi} \cdot \bar{\beta} \cdot (1 + \bar{\mu}_0) e^{(t_0 + \Delta)L^\eta} \cdot \int_{t_0}^{t_1} \int_0^1 \left| m_\tau[s|\bar{S}](x_0, y) - m_\tau[s|S'](x_0, y) \right| dy d\tau \\
&+ \bar{\pi} \cdot (1 + \bar{\mu}_0) e^{(t_0 + \Delta)L^\eta} \cdot (L^\beta + \bar{\beta}) \cdot \int_{t_0}^{t_1} \int_0^1 \left| \mu_\tau^Y[\bar{S}](y) - \mu_\tau^Y[S'](y) \right| dy d\tau \\
&+ \bar{\Lambda} \cdot \int_{t_0}^{t_1} \left| \hat{W}_\tau^X[s|\bar{S}](x_0) - \hat{W}_\tau^X[s|S'](x_0) \right| d\tau.
\end{aligned}$$

We then make use of the preceding lemmata to establish bounds on the difference in match indicator functions as well as the difference in state. Fix some $\xi > 0$ (to be determined). We know there exists $\sigma > 0$ —uniform in s —such that for all $S : \mathbf{d}^2(S, \bar{S}) < \sigma$

$$\int_{t_0}^{t_1} \int_0^1 \left| m_\tau[s|\bar{S}](x_0, y) - m_\tau[s|S'](x_0, y) \right| dy d\tau < \xi \quad \text{and} \quad \int_{t_0}^{t_1} \int_0^1 \left| \mu_\tau^Y[\bar{S}](y) - \mu_\tau^Y[S'](y) \right| dy d\tau < \xi.$$

It thereby follows that

$$\left| \hat{W}_{t_0}^X[s|S](x_0) - \hat{W}_{t_0}^X[s|\bar{S}](x_0) \right| \leq \bar{\pi}(1 + \bar{\mu}_0) e^{t_1 L^\eta} (2\bar{\beta} + L^\beta) \xi + \bar{\Lambda} \int_{t_0}^{t_1} \left| \hat{W}_\tau^X[s|\bar{S}](x_0) - \hat{W}_\tau^X[s|S](x_0) \right| d\tau.$$

Then according to Grönwall's lemma

$$\left| \widehat{W}_{t_0}^X[s|S](x_0) - \widehat{W}_{t_0}^X[s|\bar{S}](x_0) \right| \leq \underbrace{\xi (\bar{\pi}(1 + \bar{\mu}_0)e^{t_1 L'} (2\bar{\beta} + L^\beta) \exp\{(t_1 - t_0)\bar{\Lambda}\})}_{\equiv K}$$

and K is a constant which does not depend on the choice of ξ . Then pick $\xi = \frac{1}{K} \frac{\epsilon}{2}$ and choose δ equal to the corresponding σ . \square

B.3 The approximate fixed-point mapping

We next construct an approximate best-response to a given strategy profile $S \in \mathcal{S}^2$. Such approximation will be achieved through convolution with the fixed-point mapping. To begin with, extend the fixed-point mapping $f = (f^X, f^Y)$ such that $f^X[S]$ is defined for all $(t, x) \in \mathbb{R}^2$. We define for all $S \in \mathcal{S}^2$

$$\bar{f}^X[S](t, x) = \begin{cases} f^X[S](|t|, |x|) & \text{if } x \in [-1, 0) \\ f^X[S](|t|, x) & \text{if } x \in [0, 1] \\ f^X[S](|t|, 1 + |x - 1|) & \text{if } x \in (1, 2], \end{cases}$$

arbitrary, say one, otherwise and symmetrically so for population Y . Next, consider a sequence $b_k = 4/k$ where $k \in \mathbb{N}$ and define

$$\delta_k(t, x) = \frac{1}{(b_k)^2} \quad \text{if } (t, x) \in B_k(0) = \{(t, x) \in \mathbb{R}^2 : \max\{|t|, |x|\} \leq \frac{b_k}{2}\}, \text{ zero otherwise.}$$

We remark that δ_k is a Dirac sequence, more specifically, a sequence of functions $\delta_k : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ such that for all $k \in \mathbb{N}$

$$\int_{\mathbb{R}^2} \delta_k(t, x) d(t, x) = 1 \quad \text{and} \quad \int_{\mathbb{R}^2 \setminus B_k(0)} \delta_k(t, x) d(t, x) = 0.$$

Now define the approximate best-response $f_k = (f_k^X, f_k^Y) : \mathcal{S}^2 \rightarrow \mathcal{S}^2$ for all $(t_0, x_0) \in [0, \infty) \times [0, 1]$ through convolution:

$$f_k^X[S](t_0, x_0) = (\bar{f}^X[S] * \delta_k)(t_0, x_0) = \int_{\mathbb{R}^2} \bar{f}^X[S](t, x) \delta_k(t_0 - t, x_0 - x) d(t, x).$$

As is well-known, see for instance Königsberger (2004) 10.1 II, the convolution with a Dirac sequence converges to the function itself:

Remark 6.

$$\forall T > 0 : \int_{[0, T] \times [0, 1]} |f^X[S](t, x) - f_k^X[S](t, x)| d(t, x) \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

A symmetric construction applies for population Y . We proceed with two lemmata. The first lemma establishes that the approximate best-responses are Lipschitz-continuous.

Lemma 5. $f_k[\mathcal{S}^2] \subseteq (\mathcal{S}_k)^2$, i.e., for all $S \in \mathcal{S}^2$ both $f_k^X[S] \in \mathcal{S}_k$ and $f_k^Y[S] \in \mathcal{S}_k$.

Proof. Pick arbitrary $S \in \mathcal{S}^2$. Pick arbitrary $(t_1, x_1), (t_0, x_0) \in [0, \infty) \times [0, 1]$. We show that

$$|f_k^X[S](t_1, x_1) - f_k^X[S](t_0, x_0)| \leq k \max\{|t_1 - t_0|; |x_1 - x_0|\} \equiv kC.$$

Or, this is vacuously the case if $kC > 1$. Thus suppose otherwise that $kC \leq 1$. In particular this implies that $C \leq \frac{1}{k} < \frac{2}{k} = b_k/2$. Then

$$|f_k^X[S](t_1, x_1) - f_k^X[S](t_0, x_0)| \leq \frac{1}{(b_k)^2} \int_{B_k(t_1, x_1) \Delta B_k(t_0, x_0)} d(t, x) \leq \frac{1}{(b_k)^2} \int_{B_k(t_1, x_1) \Delta B_k(t_1 + C, x_1 + C)} d(t, x) \leq 2 \frac{C b_k + (b_k - C)C}{(b_k)^2} \leq kC,$$

where $A \Delta B = (A \setminus B) \cup (B \setminus A)$ denotes the symmetric difference.




Figure 4: The shaded area corresponds to the measure of $B_k(t_1, x_1) \Delta B_k(t_1 + C, x_1 + C)$.

□

The second lemma shows that the approximate best-response mapping is continuous in strategies.

Lemma 6. $f_k : (\mathcal{S}^2, \mathbf{d}^2) \rightarrow (\mathcal{S}^2, \mathbf{d}^2)$ is continuous.

Proof. It suffices to show continuity of f_k^X , identical arguments apply for f_k^Y . Fix \bar{S} . Fix $\epsilon > 0$. Let $T : e^{-T} < \epsilon$. Let

$$A_n^X = \left\{ (t, x) \in \cup_{t' \in [0, T], x' \in [0, 1]} B_k(t', x') : |\bar{f}^X[S](t, x) - \bar{f}^X[\bar{S}](t, x)| < \frac{\epsilon}{2} \text{ for all } S : \mathbf{d}^2(S, \bar{S}) < \frac{1}{n} \right\}.$$

Due to proposition 6, $S \mapsto f^X[S](t, x)$ is continuous at \bar{S} for almost all (t, x) . Whence there exists N such that $\int_{A_N^X} d(t, x) < \frac{\epsilon}{2}(b_k)^2$. Denote $\delta = \frac{1}{N}$.

To conclude, pick arbitrary $(t, x) \in [0, \infty) \times [0, 1]$. And notice that for all $S : \mathbf{d}^2(S, \bar{S}) < \delta$ we can bound $e^{-t} |f_k[S](t, x) - f_k[\bar{S}](t, x)|$ by

$$e^{-t} \int_{B_k(t,x) \cap A_N^X} \frac{|\bar{f}[S](t', x') - \bar{f}[\bar{S}](t', x')|}{(b_k)^2} d(t', x') + e^{-t} \int_{B_k(t,x) \setminus A_N^X} \frac{|\bar{f}[S](t', x') - \bar{f}[\bar{S}](t', x')|}{(b_k)^2} d(t', x') \leq \frac{\epsilon}{2} + \frac{\epsilon}{2},$$

whence $\mathbf{d}(f_k^X[S], f_k^X[\bar{S}]) \leq \epsilon$ as desired. □

We conclude by showing that for all $k \in \mathbb{N}$ the approximate best-response mapping f_k admits a fixed-point strategy profile $S \in (\mathcal{S}_k)^2$. In light of remark 6 we can think of a fixed point as an approximate equilibrium.

Proposition 7. $f_k : \mathcal{S}^2 \rightarrow \mathcal{S}^2$ has a fixed point $S_{\langle k \rangle}$ in $(\mathcal{S}_k)^2$.

Proof. This is an immediate consequence of Schauder's fixed-point theorem: this theorem asserts that if a mapping $f_k : \mathcal{S}^2 \rightarrow \mathcal{S}^2$ is such that \mathcal{S}^2 is a nonempty convex closed subset of a Hausdorff topological vector space and f_k is a continuous mapping whose image is contained in a compact subset of \mathcal{S}^2 , then f_k has a fixed point.

Clearly, $(\mathcal{S}^2, \mathbf{d}^2)$, being a metric space, possesses the Hausdorff property and is nonempty, convex, and closed. Lemma 6 asserts that f_k is continuous. Finally, lemma 5 ensures that the image of f_k is $(\mathcal{S}_k)^2$. And proposition 5 establishes that $(\mathcal{S}_k, \mathbf{d})$ is compact, hence also $((\mathcal{S}_k)^2, \mathbf{d}^2)$. Whence the result. □

B.4 Convergence of the approximate fixed points

In this subsection we will prove the following statement:

Proposition 8. *The sequence of approximate equilibria $(S_{\langle k \rangle})_{k \in \mathbb{N}}$ admits a pointwise convergent subsequence with limit in \mathcal{S}^2 .*

We first introduce a generalized version of Helly's selection theorem. This theorem asserts that a sequence of functions (of a single variable) which are uniformly bounded and of uniformly bounded total variation admits a pointwise convergent subsequence. Here we follow and adapt Chistyakov and Tretyachenko (2010a) who extend this theorem to functions of several variables.

Definition 8 (total variation). *The total variation norm in dimension two on a compact subset $[a, b] \times [\underline{t}, \bar{t}]$ for a function $v : [a, b] \times [0, \infty) \rightarrow [0, \infty)$ is given by*

$$T\mathcal{V}(v, [a, b] \times [\underline{t}, \bar{t}]) = \mathcal{V}_0^1(v(\cdot, \underline{t})) + \mathcal{V}_{\underline{t}}^{\bar{t}}(v(0, \cdot)) + \mathcal{V}_2(v, [a, b] \times [\underline{t}, \bar{t}])$$

with

$$\mathcal{V}_0^1(v(\cdot, t_0)) = \sup_{\mathcal{P}} \sum_{i=1}^m |v(x_i, t_0) - v(x_{i-1}, t_0)|$$

where \mathcal{P} is a partition of $[a, b]$, i.e. $a = x_0 < x_1 < \dots < x_m = b$,

$$\mathcal{V}_{\underline{t}}^{\bar{t}}(v(0, \cdot)) = \sup_{\mathcal{P}} \sum_{i=1}^m |v(0, t_i) - v(0, t_{i-1})|$$

where \mathcal{P} is a partition of $[\underline{t}, \bar{t}]$, i.e. $\underline{t} = t_0 < t_1 < \dots < t_m = \bar{t}$,

$$\mathcal{V}_2(v, [a, b] \times [\underline{t}, \bar{t}]) = \sup_{\mathcal{P}} \sum_{i=1}^m |v(x_i, t_i) - v(x_i, t_{i-1})| + |v(x_{i-1}, t_i) - v(x_{i-1}, t_{i-1})|$$

where \mathcal{P} is a discrete path in $[a, b] \times [\underline{t}, \bar{t}]$, s.t. $\underline{t} = t_0 < t_1 < \dots < t_m = \bar{t}$

and $a = x_0 < x_1 < \dots < x_m = b$.

We then present the two-dimensional extension of Helly's selection theorem (which extends to arbitrarily many dimensions).

Theorem 4 (Idczak and Walczak (1994); Leonov (1996); Chistyakov and Tretyachenko (2010b) and references therein). *A sequence of uniformly bounded maps $v_k : [a, b] \times [\underline{t}, \bar{t}] \rightarrow \mathbb{R}$ such that*

$$C \equiv \sup_{k \in \mathbb{N}} T\mathcal{V}(v_k, [a, b] \times [\underline{t}, \bar{t}]) \text{ is finite}$$

contains a subsequence which converges pointwise on $[a, b] \times [\underline{t}, \bar{t}]$ to a map v such that $T\mathcal{V}(v, [a, b] \times [\underline{t}, \bar{t}]) < C$.

This theorem relates to our sequence of approximate equilibria $(S_{\langle k \rangle})_{k \in \mathbb{N}}$, in that each strategy profile $S = (S^X, S^Y)$, restricted to a finite time interval $[0, T]$, can be thought of as a mapping $v : [0, 2] \times [0, T] \rightarrow [0, 1]$, where $v(x, t) = S_t^X(x)$ if $x \in [0, 1]$, and $v(x, t) = S_t^Y(x-1)$ if $x \in (1, 2]$. Further, by virtue of being a fixed point of f_k , each $S_{\langle k \rangle}$ is a composite map involving the value of search (which we can similarly view as a mapping from $[0, 2] \times [0, T]$ into the bounded interval $[0, \bar{\pi}]$).

Note that this construction omits the strategy of type 0 in population Y . This is without loss since a single type's behavior can not be payoff-relevant. Indeed, having fixed a strategy for all types $\neq 0$, any type 0 strategy gives rise to an identical value of search and population dynamics for all agent types $\neq 0$. In turn, this defines a unique best-response for all agent types (including 0) —regardless of agent type 0's initial individual strategy.

We now return our attention to the existence of a convergent subsequence of the approximate equilibria $S_{\langle k \rangle}$. Those being composite maps involving the value of search, we begin by showing that the total variation of the value of search is uniformly bounded across strategy profiles.

Lemma 7. *For all $T > 0$ there exists a finite C such that for all $S \in \mathcal{S}^2$*

$$T\mathcal{V}(V[S], [0, 2] \times [0, T]) < C.$$

Proof. Recall from proposition 2 (i) and (ii) that both $\mathcal{V}_0^1(x \mapsto V_t^X[S](x))$ (whence also $\mathcal{V}_0^2(x \mapsto V_t[S](x))$) and $\mathcal{V}_t^1(t \mapsto V_t^X(x))$ are uniformly bounded for all $S \in \mathcal{S}^2$. Further observe that due to 2 (i) there exists C' which does not depend on the choice of x nor population X such that $|V_{t_2}^X[S](x) - V_{t_1}^X[S](x)| \leq C' |t_2 - t_1|$. Therefore

$$\mathcal{V}_2((x, t) \mapsto V_t[S](x), [0, 2] \times [0, T]) \leq 2 \sup_{\mathcal{P}} \sum_{i=1}^m \max_{Z \in \{X, Y\}} \sup_{x \in [0, 1]} |V_{t_i}^Z[S](x) - V_{t_{i-1}}^Z[S](x)| \leq 2 C' T,$$

where \mathcal{P} is a partition of the time interval $[0, T]$. Then letting $C = 2 C' T$ gives the result. \square

Corollary 2. *For all $T > 0$ there exists a finite C such that for all $k \in \mathbb{N}$ and all $S \in \mathcal{S}^2$*

$$T \mathcal{V}(f_k[S], [0, 2] \times [0, T]) < C.$$

Proof. It suffices to prove the claim for $f[S]$, as the uniform bound on the total variation will be preserved under convolution. Let L^π be the associated Lipschitz constant of $z \mapsto (\pi^X(\cdot|x))^{-1}(z)$ and, for given T, C' the uniform bound on the total variation of $V[S]$ for all S . Then $C \equiv 2 L^\pi C'$ is the uniform bound on the total variation of $f[S]$ for all $S \in \mathcal{S}^2$. \square

We finally prove the main proposition.

Proof of the proposition. By construction $S_{\langle k \rangle} = f_k(S_{\langle k \rangle})$. Whence there exists C such that $T \mathcal{V}(S_{\langle k \rangle}, [0, 2] \times [0, T]) \leq C$ for all k due to the preceding corollary. It then follows from the selection theorem that $(S_{\langle k \rangle})_{k \in \mathbb{N}}$ admits a subsequence that converges pointwise on $[0, 2] \times [0, T]$. Similarly, we can find a subsequence of the subsequence which converges pointwise on $[0, 2] \times [0, T+1]$. Proceeding by induction then establishes pointwise convergence on $[0, 2] \times [0, \infty)$. As noted, we can recover the strategy of agent type 0 from population Y through his best-response. Finally, since each $S_{\langle k \rangle}$ is continuous, the limit is jointly measurable, whence in \mathcal{S}^2 . \square

B.5 Limiting arguments

Proof of theorem 1. Denote S the limit of a subsequence of approximate equilibria $(S_{\langle k \rangle})_{k \in \mathbb{N}}$ which both exist to according propositions 7 and 8. We show that $\bar{S} \equiv f[S]$ is an equilibrium, i.e., that $\bar{S} = f[\bar{S}]$.

Notation: Fix $T > 0$. Let $h \in \mathcal{S}^2$ and define $\|h\|_T \equiv \max \left\{ \int_0^T \int_0^1 |h_t^X(x)| dt dx; \int_0^T \int_0^1 |h_t^Y(y)| dt dy \right\}$.

Preliminary step: We show that $\|S - f[S]\|_T = 0$ for all T .

Due to the triangle inequality

$$\|S - f[S]\|_T \leq \|S - S_{\langle k \rangle}\|_T + \|f_k[S_{\langle k \rangle}] - f_k[S]\|_T + \|f_k[S] - f[S]\|_T,$$

where we have made use of the fact that $S_{\langle k \rangle}$ is a fixed point, i.e., $f_k[S_{\langle k \rangle}] = S_{\langle k \rangle}$. By construction and due to remark 6, both the first and the last term go to zero as $k \rightarrow \infty$. This is also

true for the second term, since

$$\begin{aligned}
& \int_0^T \int_0^1 |f_k^X[S_{<k>}] (t, x) - f_k^X[S] (t, x)| dx dt \\
&= \int_0^T \int_0^1 \left| \frac{1}{(b_k)^2} \int_{B_k(t,x)} \bar{f}^X[S_{<k>}] (\tau, z) - \bar{f}^X[S] (\tau, z) d(\tau, z) \right| dx dt \\
&\leq \int_{-b_k/2}^{T+b_k/2} \int_{-b_k/2}^{1+b_k/2} |\bar{f}^X[S_{<k>}] (t, x) - \bar{f}^X[S] (t, x)| dx dt \\
&\leq 9 \int_0^T \int_0^1 |f^X[S_{<k>}] (t, x) - f^X[S] (t, x)| dx dt
\end{aligned}$$

for k sufficiently large (whence $b_k/2 < \max\{1, T\}$). It follows that

$$\|f_k[S_{<k>}] - f_k[S]\|_T \leq 9 \|f[S_{<k>}] - f[S]\|_T.$$

This converges to zero as $k \rightarrow \infty$, because the fixed-point mapping f inherits pointwise continuity of the value of search in strategy profiles a.e. (refer to proposition 6).

Conclusion: Denoting $\bar{S} \equiv f[S]$, the preliminary step establishes that $\|S - \bar{S}\|_T = 0$ for all $T > 0$. Now recall that two strategy profiles that agree almost everywhere give rise to an identical best-response strategy profile. Or, $f[S] = f[S']$ if $S' : \|S - S'\|_T = 0$ for all $T > 0$. Whence $f[S] = f[\bar{S}]$ as desired.

Conclusion: Recall that two strategy profiles that agree almost everywhere give rise to an identical best-response strategy profile. Or, $f[S] = f[S']$ if $S' : \|S - S'\|_T = 0$ for all $T > 0$. In particular, the preliminary step establishes that for $\bar{S} \equiv f(S)$ we have $\|S - \bar{S}\|_T = 0$ for all $T > 0$. Whence $f[\bar{S}] = f[S] = \bar{S}$ as desired. \square

C PAM: miscellaneous proofs

Proof of 4. Suppose by contraposition that the strategy profile S is such that at time t there is weak PAM, yet $S_t^X(x_1) > S_t^X(x_2)$ for some $x_1 < x_2$.

case one: $\exists \underline{y} \in [S_t^X(x_2), S_t^X(x_1)) \cap M_t^X[S](x_2)$. Then pick arbitrary $\bar{y} \in M_t^X[S](x_1)$. Clearly, $\bar{y} \geq S_t^X(x_1) > \underline{y}$, whence, due to weak PAM, $m_t[S](x_2, \underline{y}) = 1$, $m_t[S](x_1, \bar{y}) = 1 \Rightarrow m_t[S](x_1, \underline{y}) = 1$. This contradicts $\underline{y} < S_t^X(x_1)$.

case two: $[S_t^X(x_2), S_t^X(x_1)) \cap M_t^X[S](x_2) = \emptyset$. Then pick arbitrary $y_2 \in M_t^X[S](x_2)$ and $y_1 \in [S_t^X(x_2), S_t^X(x_1))$. Clearly, $y_2 > y_1$ and $S_t^Y(y_1) > x_2$. Then pick arbitrary $x_3 \in M_t^Y[S](y_1)$. Clearly $x_3 > x_2$, whence, due to weak PAM, $m_t[S](x_2, y_2) = 1$, $m_t[S](x_3, y_1) = 1 \Rightarrow m_t[S](x_2, y_1) = 1$. This contradicts $S_t^Y(y_1) > x_2$. \square

Proof of remark 5. Omit population superscripts. We show that PAM implies that $m_t[S](x, x) = 1$ for all types x and time t . Suppose by contraposition there exists $t \geq 0$ such that $S_t(x) > x$. Then, due to continuity of $t \mapsto S_t(x)$ (as implied by continuity of $t \mapsto V_t[S](x)$; proposition 2) there exists a non-empty interval (t_0, t_1) such that $S_t(x) > x$ for all $t \in (t_0, t_1)$.

Next, note that PAM implies that for all $t \in (t_0, t_1)$ agent type x does not match with anyone: $m_t(x, x') = 0$ for all $x' \in [0, 1]$ throughout. Indeed, $m_t(x, x') = 1$ and $m_t(x', x) = 1$ (drawing on the symmetry of populations) would imply that $m_t(x, x) = 1$. Since (t_0, t_1) is maximal, either t_1 is infinite, in which case agent type x would never match with anyone, whence $V_t[S](x) = 0$; or t_1 is finite and $V_{t_1}[S](x) = \pi(x|x)$, whence $V_t[S](x) = e^{-\rho(t_1-t)}\pi(x|x) < \pi(x|x)$ for all $t \in [t_0, t_1]$. In either case self-rejection cannot be optimal. \square

For pedagogical reasons, we include here a slightly more general statement of a well known result (see Pratt (1964) as well as the textbook analysis by Gollier (2004)).

Proposition 9. *Assume that $y \mapsto \pi^X(y|x)$ is monotone and continuously differentiable. Take $y_1 < y_2$. Then payoffs satisfy strict log-supermodularity in differences (LSD) if and only if $\partial_y \pi^X(y|x)$ is strictly log supermodular.*

In the proof, we denote $f(x) \nearrow x$ if $f(x)$ is strictly increasing in x .

Proof. (\Leftarrow) Suppose that for any $y_1 < y_2$ the fraction $\frac{\partial_y \pi^X(y_2|x)}{\partial_y \pi^X(y_1|x)}$ is strictly increasing in x . This property is preserved under integration, hence

$$\frac{\int_{y_2}^{y_3} \partial_z \pi^X(z|x) dz}{\partial_y \pi^X(y_1|x)} \nearrow \text{in } x \quad \Leftrightarrow \quad \frac{\pi^X(y_3|x) - \pi^X(y_2|x)}{\partial_y \pi^X(y_1|x)} \nearrow \text{in } x$$

for any $y_3 > y_2 > y_1$. Then by the same argument (reversing the nominator and the denominator)

$$\frac{\int_{y_1}^{y_2} \partial_z \pi^X(z|x) dz}{\pi^X(y_3|x) - \pi^X(y_2|x)} \searrow \text{in } x \quad \Leftrightarrow \quad \frac{\pi^X(y_2|x) - \pi^X(y_1|x)}{\pi^X(y_3|x) - \pi^X(y_2|x)} \searrow \text{in } x.$$

Then once more reversing the nominator and the denominator gives log supermodularity in differences as was to be shown.

(\Rightarrow) Suppose that payoffs satisfy LSD. Define

$$f_x(y_2, y_1) = \pi^X(y_2|x) - \pi^X(y_1|x).$$

Then LSD rewrites as

$$\frac{f_x(y_3, y_2)}{f_x(y_2, y_1)} \nearrow \text{in } x$$

for any $y_3 > y_2 > y_1$. This implies

$$\lim_{y_3 \rightarrow y_2} \frac{f_x(y_3, y_2)}{f_x(y_2, y_1)} \text{ weakly } \nearrow \text{in } x \quad \Leftrightarrow \quad \frac{\partial_y \pi^X(y_2|x)}{\pi^X(y_2|x) - \pi^X(y_1|x)} \text{ weakly } \nearrow \text{in } x.$$

Then for any $x_2 > x_1$ and $y_2 > y_1$

$$\frac{\partial_y \pi^X(y_2|x_2)}{\partial_y \pi^X(y_2|x_1)} \geq \frac{\pi^X(y_2|x_2) - \pi^X(y_1|x_2)}{\pi^X(y_2|x_1) - \pi^X(y_1|x_1)}.$$

Now for fixed $x_2 > x_1$ and supposing that $y_2 > y_1$, LSD amounts to saying that $\frac{\pi^X(y_2|x_2) - \pi^X(y_1|x_2)}{\pi^X(y_2|x_1) - \pi^X(y_1|x_1)}$ is strictly increasing in y_2 . Therefore

$$\frac{\partial_y \pi^X(y_2|x_2)}{\partial_y \pi^X(y_2|x_1)} > \lim_{y_2 \rightarrow y_1} \frac{\frac{\pi^X(y_2|x_2) - \pi^X(y_1|x_2)}{y_2 - y_1}}{\frac{\pi^X(y_2|x_1) - \pi^X(y_1|x_1)}{y_2 - y_1}},$$

where the latter expression is equal to $\frac{\partial_y \pi^X(y_1|x_2)}{\partial_y \pi^X(y_1|x_1)}$ due to $\partial_y \pi^X(y|x)$ being monotone in y and continuous. This concludes the proof. \square

D Necessity

Proof of necessity, proposition 4.

Counterexample A. There are two types, $x_2 > x_1$, whose payoffs are strictly log submodular. That is

$$\frac{\pi(x_2|x_2)}{\pi(x_1|x_2)} < \frac{\pi(x_2|x_1)}{\pi(x_1|x_1)}.$$

Search is quadratic, i.e. $\lambda(t, \mu_t) = \mu_t$ and there is no entry.

As match prospects are bleakening over time, there exists a time t^* beyond which the high type will always accept the low type and $V_{t^*}(x_2) = \pi(x_1|x_2)$. Drawing on the integral representation of the value of search we can express $V_{t^*}(x_2)$ as

$$V_{t^*}(x_2) = \sum_{j \in \{1,2\}} \pi(x_j|x_2) Q_{t^*}(x_j)$$

where $Q_{t^*}(x_j)$ is the probability of type x_2 matching with x_j - discounted by the time at which such event materializes (see also the subsequent proof). Now observe that if the low type found it desirable, she could always exactly replicate discounted match probabilities of the high type, that is

$$V_{t^*}(x_1) \geq \sum_{j \in \{1,2\}} \pi(x_j|x_1) Q_{t^*}(x_j).$$

Then $V_{t^*}(x_1) > \pi(x_1|x_1)$ and the low type rejects other low types at time t^* . For otherwise the integral representation of the value of search combined with the inequalities implies that

$$\sum_{j \in \{1,2\}} \frac{\pi(x_j|x_2)}{\pi(x_1|x_2)} Q_{t^*}(x_j) \geq \sum_{j \in \{1,2\}} \frac{\pi(x_j|x_1)}{\pi(x_1|x_1)} Q_{t^*}(x_j) \Leftrightarrow \frac{\pi(x_2|x_2)}{\pi(x_1|x_2)} \geq \frac{\pi(x_2|x_1)}{\pi(x_1|x_1)}$$

in spite of strict log submodularity.

Counterexample B.

Consider symmetric populations consisting of three types x_1, x_2, x_3 . Omit superscripts. Suppose that x_3 is strictly more risk-averse than x_2 . We construct a number of equilibria indexed by n in which there exists a moment in time such that x_3 accepts x_2 whereas x_2 rejects a fellow x_2 . Specifically, consider two distinct moments in time, t_0^n and 0 where t_0^n precedes 0: at time t_0^n the high type x_3 begins accepting the intermediate type x_2 and at time 0 the high type begins accepting the low type x_1 ; PAM will be upset because type x_2 will reject another type x_2 at time t_0^n .

The construction makes apparent that the failure of PAM at time t_0^n arises due to a reversal of risk preferences. As n grows large both (i) $t_0^n \rightarrow 0$ and (ii) the probability of matching after time 0 will go to zero. As a consequence agent type x_3 's future match outcomes at time t_0^n converge towards a lottery assigning positive probability to both the event that x_3 match with another x_3 and to the event that x_3 match with an agent type x_1 . Crucially, at time t_0^n agent types x_2 are accepted by agent types x_3 . They thus face identical match opportunities. Like agent types x_3 they may either choose to play the lottery—or accept x_2 . Note that since agent type x_3 is indifferent between playing the lottery, i.e., waiting, or accepting x_2 , by virtue of being less risk-averse agent type x_2 must strictly prefer the lottery and therefore reject another type x_2 .

To construct the failure of PAM analytically, we consider the simplest non-stationary matching environment conceivable. There is zero entry. Agent type x_2 is present in zero proportion and solely of hypothetical interest. Due to log supermodularity agent type x_1 will accept any agent he meets. Proceed then to define the (anonymous) meeting rate: it becomes stationary eventually and is piecewise constant over time. We set

$$\lambda_t(x_1) = n(1 - h(n)) \quad \text{if } t \geq 0 \quad \text{and} \quad \lambda_t(x_3) = \begin{cases} nh(n) & \text{if } t \geq 0 \\ n & \text{if } t < 0. \end{cases}$$

$h(n)$ is determined as to ensure indifference of agent type x_3 between accepting and rejecting agent types x_1 :

$$\rho V_0^n(x_3) = n[h(n)\pi(x_3|x_3) + (1 - h(n))\pi(x_1|x_3) - V_0^n(x_3)] \quad \text{and} \quad V_0^n(x_3) = \pi(x_1|x_3).$$

Then

$$h(n) = \frac{\rho}{n} \frac{\pi(x_1|x_3)}{\pi(x_3|x_3) - \pi(x_1|x_3)}.$$

(Note that log supermodularity implies that at time 0 agent type x_2 must likewise be accepting agent types x_1 .) Finally, choose as time 0 'starting values' $(\mu_0(x_1), \mu_0(x_2), \mu_0(x_3))$ such that $\mu_0(x_2) = 0$ and $\frac{\mu_t(x_3)}{\mu_t(x_1)} = \frac{\lambda_t(x_3)}{\lambda_t(x_1)}$. Note that this construction does not run counter the demand that the search technology be anonymous: following time 0 there is common acceptance of all types so that under any anonymous search technology the ratio $\frac{\mu_t(x_3)}{\mu_t(x_1)}$ is constant for all $t \geq 0$. $\lambda_t(x_1)$ for $t < 0$ will be uniquely pinned down by $\lambda_t(x_3)$ and μ_t (in particular $\lambda_t(x_1) = \lambda_t(x_3) \frac{\mu_t(x_1)}{\mu_t(x_3)}$), but this is inconsequential as agent type x_3 rejects agent types x_1 at $t < 0$.

Time t_0^n , the moment in time at which agent type x_3 is indifferent between accepting and

rejecting agent type x_2 likewise admits a closed-form representation:

$$\begin{aligned} V_{t_0^n}^n(x_3) &= \int_{t_0^n}^0 e^{-\rho(\tau-t_0^n)} \pi(x_3|x_3) n e^{-n(\tau-t_0^n)} d\tau + e^{\rho t_0^n} e^{n t_0^n} \pi(x_1|x_3) \\ &= \frac{n}{\rho+n} [1 - e^{t_0^n(\rho+n)}] \pi(x_3|x_3) + e^{t_0^n(\rho+n)} \pi(x_1|x_3) \end{aligned}$$

and $V_{t_0^n}^n(x_2) = \pi(x_2|x_3)$.

Then

$$t_0^n = \frac{1}{\rho+n} \ln \frac{\frac{n}{\rho+n} \pi(x_3|x_3) - \pi(x_2|x_3)}{\frac{n}{\rho+n} \pi(x_3|x_3) - \pi(x_1|x_3)}$$

Clearly, $t_0^n < 0$ and $t_0^n \rightarrow 0$ as n goes to infinity.

Agent type x_3 's discounted match probabilities of matching with agent types x_1 and x_3 are given by $Q_{t_0^n}^n(x_1)$ and $Q_{t_0^n}^n(x_3)$ respectively, where

$$\begin{aligned} Q_{t_0^n}^n(x_1) &= e^{t_0^n(\rho+n)} \int_0^\infty e^{-\rho\tau} n(1-h(n)) e^{-n\tau} d\tau \\ &= e^{t_0^n(\rho+n)} \frac{n(1-h(n))}{\rho+n} \\ &= \frac{\frac{n}{\rho+n} \pi(x_3|x_3) - \pi(x_2|x_3)}{\frac{n}{\rho+n} \pi(x_3|x_3) - \pi(x_1|x_3)} \frac{n(1-h(n))}{\rho+n} = \frac{\pi(x_3|x_3) - \pi(x_2|x_3)}{\pi(x_3|x_3) - \pi(x_1|x_3)} + o(1) \\ &\equiv q + o(1) \end{aligned}$$

$$\begin{aligned} Q_{t_0^n}^n(x_3) &= \int_{t_0^n}^0 e^{-\rho(\tau-t_0^n)} n e^{-n(\tau-t_0^n)} d\tau + e^{t_0^n(\rho+n)} \int_0^\infty e^{-\rho\tau} n h(n) e^{-n\tau} d\tau \\ &= \frac{n}{\rho+n} [1 - e^{(\rho+n)t_0^n}] + e^{(\rho+n)t_0^n} \frac{n h(n)}{\rho+n} \\ &= \frac{n}{\rho+n} - \frac{n(1-h(n))}{\rho+n} \frac{\frac{n}{\rho+n} \pi(x_3|x_3) - \pi(x_2|x_3)}{\frac{n}{\rho+n} \pi(x_3|x_3) - \pi(x_1|x_3)} \\ &= 1 - \frac{\pi(x_3|x_3) - \pi(x_2|x_3)}{\pi(x_3|x_3) - \pi(x_1|x_3)} + o(1) = (1-q) + o(1). \end{aligned}$$

and $Q_{t_0^n}^n(x_1) + Q_{t_0^n}^n(x_3) = 1 + o(1)$. (Here $o(1)$ denotes the Landau notation: $\lim_{n \rightarrow \infty} o(1) = 0$.)

Now observe that beginning from time t_0^n agent type x_2 is accepted by agent type x_3 , and thus faces identical match opportunities as an agent type x_3 . Accordingly she can mimic the higher type's match probabilities so that

$$V_{t_0^n}^n(x_2) \geq \pi(x_1|x_2)q + \pi(x_3|x_2)(1-q) + o(1).$$

(Recall by construction that $\pi(x_2|x_3) = V_{t_0^n}^n(x_3) = \pi(x_1|x_3)q + \pi(x_3|x_3)(1-q) + o(1)$.) We then claim that $V_{t_0^n}^n(x_2) > \pi(x_2|x_2)$ for n sufficiently large, so that PAM is upset at time t_0^n . Indeed, this

follows from the characterization of risk-preferences. Suppose by contradiction that $V_{t_0}^n(x_2) \leq \pi(x_2|x_2)$ for all $n \in \mathbb{N}$. Letting $n \rightarrow \infty$ gives

$$\pi(x_2|x_2) \geq \pi(x_1|x_2)q + \pi(x_3|x_2)(1 - q) \quad \text{and} \quad \pi(x_2|x_3) = \pi(x_1|x_3)q + \pi(x_3|x_3)(1 - q)$$

This means that (i) agent type x_3 is indifferent between the lottery assigning probability q to x_1 and $1 - q$ to x_2 and the sure outcome x_2 , whereas (ii) agent type x_2 weakly prefers the sure outcome x_2 . This contradicts the assertion that agent type x_2 is less risk-averse than agent type x_3 . \square

Probabilistic Assortative Matching under Nash Bargaining

Joint with Nicolas Bonneton

1

1 Introduction

This paper provides a novel perspective on what is perhaps the most studied question in the theory of matching: when is matching assortative? Interest in this question is due to both the efficiency and distributive implications of assortative matching on the economy at large. For instance, when human and physical capital are complementary, a reallocation of more productive workers to capital-rich firms raises aggregate output. Regarding distributive concerns, a greater number of marriages between individuals of similar socio-economic status is known to exacerbate existing income inequality at the household level.

The theory of decentralized matching has been spurred by Becker (1973), who formulates the first key insight within a frictionless competitive market: assortative matching arises due to complementarities in output between similar individuals. However, frictionless search often-times presents an unattainable idealization of matching markets; unemployment spells are the most obvious case in point that search frictions can be sizeable. In this paper, we study sorting in matching markets where random search frictions impede the instantaneous clearing of supply and demand. Thereby, a given individual is willing to match with a range of individuals as opposed to a single ideal partner. As in the Diamond-Mortensen-Pissarides paradigm of the labor market, we replace the Walrasian auctioneer by Nash bargaining.²

Our understanding of sorting in this framework is due to Shimer and Smith (2000), and complementary work by Atakan (2006).³ Their main contribution is to identify sufficient conditions on output for which (a set based notion of) positive assortative matching arises despite

¹This is Nicolas' job market paper.

²The Nash bargaining solution in search and matching models is also referred to as the transferable utility (TU) paradigm in the literature; in contrast, a model where match payoffs are exogenously given (analyzed in Bonneton and Sandmann (2019)) is referred to as the non-transferable utility (NTU) paradigm.

³Refer to Chade et al. (2017) for an overview of the literature.




Figure 1: Deterministic and probabilistic matching patterns.

the presence of search frictions: higher types match with intervals of higher types.⁴ In contrast, we argue that positive assortative matching is not a robust feature of matching markets—even when match output exhibits a high degree of complementarity. In this article we identify economic forces that prevent positive assortative matching from occurring when search frictions are sizeable. To make this point, we propose a framework in which match outcomes are probabilistic: no match is inconceivable, but those pairs who produce a greater match surplus are more likely to match. By comparison, in the deterministic framework of Shimer and Smith (2000) and Atakan (2006) meeting-contingent match probabilities are binary, 0 – 1.

The difference between deterministic and probabilistic matching patterns can be best appreciated visually. Figure 1 (left) illustrates the set of matching pairs in a deterministic environment where Shimer and Smith (2000)’s set based notion of assortative matching obtains: matching sets are increasing in type. Conversely, figure 1 (center) illustrates meeting-contingent match probabilities in a probabilistic environment, where warmer colors correspond to higher probabilities of matching. The shaded region in figure 1 (right) depicts the pairs of types for whom the (herein introduced) general notion of *probabilistic assortative matching* (PAM) fails: PAM requires that probabilistic level lines, depicting pairs of types that match with equal probability, are non-decreasing in either type. This is visibly not the case here. We call this a failure of assortative matching, for it implies that for all pairs of types $x_1 < x_2$ and $y_1 < y_2$ in this region, one of the assortative pairs (x_1, y_1) or (x_2, y_2) has a lower probability of matching than both non-assortative pairs (x_1, y_2) and (x_2, y_1) . This failure can even arise even when (x_1, y_1) and (x_2, y_2) is the assortative assignment absent search frictions, so that pairs of perfect complements need not be the ones matching with the greatest frequency.

The failure of PAM in the presence of search frictions is not an artefact of specific modeling choices other than Nash bargaining; it is common to all random-search matching models in which meeting-contingent match probabilities are increasing in the match surplus. Under Nash bargaining, surplus rises in the output generated, yet falls in both agents’ option values of search. Only the latter vary with the amount of search frictions present in the economy. Search frictions impede assortative matching from occurring because frictions disproportionately erode

⁴Common to both is the well-known condition from Becker (1973): supermodularity. In addition, Shimer and Smith (2000) require more stringent complementarity conditions on the curvature of $f(x, y)$, as well as boundary conditions that translate into weak self-preference among the lowest and the highest types; $f(x, y) = xy$ satisfies all of their conditions (as first studied by ?).

the value of search and hence the bargaining power of more productive agents. This is because more productive agents have greater opportunity costs of time due to discounting. Additionally, the initial increase in the opportunity cost of time further feeds back into payoffs: being more adversely affected by search frictions translates into a comparatively lower bargaining power vis-a-vis other agents. As a result, unproductive agents prioritize matching with productive agents. At the extreme when the market is thin, agents unanimously exhibit the greatest match surplus when matched with the most productive individuals. In our model, this translates one-to-one into greater meeting-contingent match probabilities with the most highly ranked individuals. In line with our visual representation, this implies that the assortative pair (x_1, y_1) exhibits a lower meeting-contingent match probability than the less assortative pairs (x_1, y_2) and (x_2, y_1) .

If not assortatively, how do agents sort into pairs? Our first main result (theorem 3) provides a general characterization of matching patterns under known complementarity conditions—regardless of the level of search frictions present in the economy. We show that sorting takes the qualitative features depicted in figure 1 (center and right). First, meeting-contingent match probabilities are single-peaked: each agent type has an ideal partner's type, and meeting-contingent match probabilities are greater for those agents more closely resembling the ideal partner's type. Furthermore, more highly-ranked types have more highly-ranked ideal partners, an implication of single-crossing. In figure 1 (right) there are two curves enclosing the pairs of types for whom PAM fails; the more gradual curve depicts ideal partners' types of agents types x , whereas the steeper curve depicts ideal partners' types of agents types y . In the region enclosed by both curves, hereafter the enclosure, level lines are decreasing so that PAM fails, whereas in the outer region level lines are increasing so that PAM obtains.⁵ The pairs of types enclosed by both curves are well-assorted, in the sense that either they comprise only low or high types, or their meeting-contingent match probability is large.

Our modeling choices emphasize two elements of the theory, that have been abstracted from in much of the preceding empirical and theoretical literature alike, but that we view as essential to correctly identify sorting. Our objective is to provide a theory simple enough to allow for theoretical insights into sorting, yet rich enough to be plausibly at the origin of empirically observed matching patterns.

First, we view match output as random, not deterministic. Specifically, as in Choo and Siow (2006), we take output to be the sum of a type-specific deterministic component $f(x, y)$ and a pair-specific production shock ξ .⁶ When two agents meet, they match if the realization of their production shock ξ is large enough, i.e., if $f(x, y) + \xi$ exceeds their joint option value of search. Thereby, match outcomes upon meeting are *random*. In particular, pairs of types that produce a larger surplus match with strictly greater probabilities.

Secondly, the economy is non-stationary. This means that the size and composition of the

⁵In the paper we think of (and establish an equivalence of) meeting-contingent match probabilities in terms of intuitive and empirically relevant ordinal properties of individual preferences over meetings: agent type x prefers to meet y_2 over y_1 if it gives x a greater expected match payoff. In our model x prefers to meet y_2 over y_1 if and only if (x, y_2) have a greater meeting-contingent match probability than (x, y_1) . Thus refer to ideal partner's types also as preferred partner's types.

⁶In the literature review we provide a detailed empirical justification as to why we need production shocks to understand matching patterns.

search pool fluctuates endogenously over time; equilibrium match acceptance decisions and individual payoffs adjust accordingly. Non-stationarity is arguably a first-order concern to sorting. As has been well-documented, wages vary over the business cycle and are higher during periods of increased match creation, i.e., booms.⁷ Here we report results that make no distinction between non-stationary and stationary environments thus far studied in the literature.⁸

Studying sorting at this level of generality is non-trivial. As pointed out by Smith (2011), “even the simplest non-stationary models can be notoriously intractable.” We circumvent the tractability issues that come with non-stationary dynamics by constructing tight bounds on the difference in the value of search between two agents. In lemma 1, following Bonneton and Sandmann (2019), we apply a revealed preference argument. The underlying idea is to let one agent replicate someone else’s matching decisions; such “mimicking” strategy must be weakly dominated due to the intratemporal efficiency of Nash bargaining. As a second step, lemma 3 uses a novel inductive reasoning over the revealed preference argument, giving rise to perhaps the most beautiful proof in this paper, *the matryochka dolls*.

Our second main contribution is to prove the existence of a unique non-stationary equilibrium (theorem 1) in a framework where economic fundamentals are defined in utmost generality.

What is a non-stationary equilibrium? The notion of equilibrium always entails an optimization problem on the side of the agents. Under Nash bargaining, this is easily resolved: two agents match whenever ex-post output $f(x, y) + \xi$ exceeds the joint option value of search. However, this solution requires knowledge of agents’ option values of search. The difficulty in establishing an equilibrium arises from solving an intricate feedback loop between the value of search and population dynamics. Agents’ expectations and decisions must conform with the population dynamics to which they give rise. As the population evolves, so do agents’ expectations. Accordingly, we define an equilibrium as a coupled system characterizing jointly forward-looking population dynamics and backward-looking agents’ option values of search under the Nash bargaining solution. We call such coupled system and the mathematical techniques that pertain to it a *mean field game*.⁹

Here we find it suitable to consider a stochastic economy: the evolution of the search pool, while dependant on agents’ matching decisions, is in addition subject to random entry. Random entry allows us to draw on techniques from stochastic calculus, notably the martingale representation theorem and Ito’s lemma. The key result (corollary 1) is to show that our equi-

⁷Lise and Robin (2017) propose and estimate a non-stationary random search matching model of workers and firms with on-the-job search when dynamics are driven by exogenous shocks. A key finding is that the estimated matching sets during lower productivity states are smaller. This leads to more assortative matching in hiring (from unemployment) during recessions.

⁸We do not suggest that non-stationarity plays a limited role under bargaining. Differences in sorting between stationary and non-stationary environments remain to be explored. It may well be the case that complementarity conditions are needlessly strong in the steady state as is the case in the NTU paradigm: as shown by Bonneton and Sandmann (2019), the well-known (steady state) results on assortative matching from Smith (2006) and Morgan (1994) only obtain under stronger complementarity conditions when the economy is non-stationary. This derives from the risk of worsening match prospects, inherent to non-stationary dynamics.

⁹The name comes from the “Mean Field theory” in physics; it is an analogy to the continuum limit taken in which one approximates large systems of interacting particles by assuming that these interact only with the statistical mean of other particles.

librium can be represented as a system of *forward backward stochastic differential equations (FBSDEs)*. The theory of FBSDEs is a powerful tool in environments where laws governing the population dynamics and the value of search, as described by the Kolmogorov and Hamilton-Jacobi-Bellman partial differential equations, are not as smooth as one would like them to be.¹⁰ Our representation of the stochastic environment as a system of FBSDEs allows us to (subject to identifying the relevant economic regularity conditions such as Lipschitz continuity and boundedness of parameters) resolve existence and uniqueness with a well-posedness theorem due to ?.

The remainder of this paper is organized as follows. Section 2 lays out the model. Section 3 establishes existence and uniqueness. Section 4 studies properties of agents' preference. Section 5 characterizes equilibrium sorting. Section 6 concludes. Proofs, if not found in the text, as well as the related literature are relegated to the appendix.

2 The model

We develop a continuous-time, finite-horizon $[0, T]$ matching model in which continua of ex-ante heterogeneous agents engage in time-consuming and haphazard search for one another. Any two agents that meet may decide whether to form a match, whereupon they exit the search pool and divide match output according to Nash bargaining. Otherwise they continue waiting for a more suitable partner.

2.1 Set-up

There are two distinct populations, employment-seeking workers and firms with vacancies say, each comprising a continuum of ex-ante heterogeneous agents. Agents of type $x \in X$ seek to match with agents of type $y \in Y$. We take X and Y to be finite, ordered and disjoint sets, i.e., $X = \{1, \dots, N^X\}$ and $Y = \{N^X + 1, \dots, N\}$, where $N = N^X + N^Y$. In what follows we usually take the viewpoint of an agent type $x \in X$. It goes without saying that symmetric constructions apply for agent types $y \in Y$.

The mass of agent types x in the search pool at time t is given by a positive number $\mu_t(x) \in \mathbb{R}_+$. The environment we study is non-stationary in the sense that the mass and composition of the search pool $\mu_t = ((\mu_t(x))_{x \in X}, (\mu_t(y))_{y \in Y}) \in \mathbb{R}_+^N$ changes over time. The search pool population μ_t and agents' continuation values of search, V_t , introduced at a later stage, are the economy's Markovian state at time t .

Output and matching decisions

Two agents that match with one another produce a lump sum match output. Match output is the sum of deterministic $f(x, y)$ and random ξ :

$$f(x, y) + \xi \quad \text{where} \quad \xi \sim \Xi_t.$$

¹⁰In the deterministic environment studied in Bonneton and Sandmann (2019) we went to great lengths to circumvent discontinuities inherent to random search matching.

We refer to $f(x, y)$ as the *ex-ante match output*, because it is manifest before a meeting has taken place. Symmetrically, we refer to $f(x, y) + \xi$ as *ex-post match output*, because it is only manifest once a meeting has taken place. ξ is distributed independently and identically across all agent pairs according to a state- and time-variant cumulative distribution function Ξ_t . The subscript t subsumes throughout the dependence on both time and the economy's Markovian state variable.

As will become apparent soon, individual match payoffs are increasing in the realization of ξ . It follows that in equilibrium joint match acceptance decisions conditional on meeting are given by a time- and state-varying threshold rule $\theta_t(x, y)$. The probability that ξ exceeds $\theta_t(x, y)$ determines the probability $m_t^\theta(x, y)$ that two agent types x and y that meet at time t match with one another:

$$m_t^\theta(x, y) = \int_{\theta_t(x, y)}^{+\infty} \Xi_t(d\xi). \quad (1)$$

We call this the *probabilistic matching function*. Observe that absent production shocks ξ , match probabilities are binary, either zero or one.

Meetings

Meetings follow a time- and state-varying Poisson point process where $\lambda_t(y)$ is the rate at which any agent type x meets an agent type y at time t . As is common in the literature, the meeting technology is anonymous. This means that one's own type x has no bearing on the probability of meeting other agents. Concordantly, we do not impose technological assumptions that favor assortative matching.

Remark 1. Any anonymous meeting rate can be represented as

$$\lambda_t(y) = \beta_t \mu_t(y)$$

for some time- and state-dependent β_t , common to all types.

For instance, search is quadratic when $\beta_t = 1$, search is linear when $\beta_t = 1 / \sum_z \mu_t(z)$.¹¹ The representation derives from internal consistency of the model, which requires that the flow number of meetings between two agent types $x \in X$ and $y \in Y$ can be symmetrically computed taking either agent type x 's or y 's perspective: $\lambda_t(y)\mu_t(x) = \lambda_t(x)\mu_t(y)$.

Entry, exit and population dynamics

Population dynamics are governed by entry and exit:

- The rate at which agent types x exit the search pool (for a given threshold rule θ_t) is given

¹¹Our existence result requires that $\lambda_t(y)$ be bounded, i.e., $\beta_t \leq K^\beta / \sum_z \mu_t(z)$ for some $K^\beta > 0$. This formally rules out quadratic search, but only so for type distributions which realize with arbitrarily small probabilities.

by x 's individual hazard rate times the total mass of types x in the search pool, $\mu_t(x)$:

$$\sum_{y \in Y} \lambda_t(y) m_t^\theta(x, y) \mu_t(x).$$

- The rate at which agent types x enter into the search pool is described by a geometric Brownian motion with endogenous parameters,

$$\mu_t(x) \eta_t(x) dt + \mu_t(x) \sigma_t(x) dB_t(x).$$

Here $\eta_t(x)$ denotes the time-and state-varying drift and $\sigma_t(x)$ the time-and state-varying volatility; $B = ((B(x)_t), (B(y)_t))_{t \geq 0}$ is an N -dimensional \mathcal{F}_t -adapted standard Brownian motion in the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$ satisfying the usual conditions.

Entry and exit aggregate into a forward stochastic differential equation where the initial population in the search pool is $\mu_0 \in \mathbb{R}_+^N$:

$$\mu_t(x) = \mu_0(x) + \int_0^t \left[- \sum_y \lambda_\tau(y) \mu_\tau(x) m_\tau^\theta(x, y) + \mu_\tau(x) \eta_\tau(x) \right] d\tau + \int_0^t \mu_\tau(x) \sigma_\tau(x) dB_\tau(x). \quad (2)$$

Nash bargaining and the value of search

If two agent types x and y decide to match, they share ex-post output according to Nash bargaining: agents receive as payoff their respective continuation value of search $V_t(x)$ or $V_t(y)$ plus a share of the remaining output $S_t(x, y) + \xi$. Here

$$S_t(x, y) \equiv f(x, y) - V_t(x) - V_t(y)$$

denotes ex-ante surplus. Under this interpretation, standard in the search and matching literature, the option value of search is used as the outside option in the bargaining game. Agent type x 's expected match payoff for given threshold rule θ_t is then given by

$$\pi_t^\theta(y|x) = V_t(x) + \frac{\alpha_t^X}{m_t^\theta(x, y)} \int_{\theta_t(x, y)}^{+\infty} [S_t(x, y) + \xi] \Xi_t(d\xi), \quad (3)$$

where α_t^X denotes the share of surplus accruing to types $x \in X$ with $\alpha_t^X + \alpha_t^Y \leq 1$ ¹² (and $m_t^\theta(x, y) > 0$ by assumption). It is easy to see that said payoff is maximal if matches are consummated whenever ex-post match surplus is non-negative, i.e.

$$\theta_t^*(x, y) = -S_t(x, y).$$

In equilibrium, agents are forward-looking and weigh the expected discounted match payoff against the immediate match payoff at hand. The remainder of this section asserts that

¹²This encompasses typically considered bargaining weights $\alpha_t^X = \alpha_t^Y = \frac{1}{2}$.

$\theta_t^*(x, y) = -S_t(x, y)$ also maximizes time t individual expected discounted match payoffs (where the meeting rate λ_τ and value of search V_τ for $\tau > t$ is exogenously given), and defines equilibrium matching thresholds.

To formalize this claim, fix $(\lambda_\tau)_{\tau>t}$ and $(V_\tau)_{\tau>t}$ (objects which are beyond the control of an individual agent), and define $r_t^\theta(y|x)$ the earliest time $r \geq t$ such that type x meets a type y and draws $\xi \geq \theta_r(x, y)$:

$$r_t^\theta(y|x) = \min \{r \geq t : \text{a given } x \text{ meets some } y \text{ s.t. } \xi \geq \theta_r(x, y)\}.$$

With threshold rule θ in place, the probability that x matches with some y during time interval $[t, \tau]$ is denoted by

$$P_t^\theta(\tau)(y|x) = \text{Prob}[r_t^\theta(y|x) \leq \tau \text{ and } r_t^\theta(y|x) \leq r_t^\theta(y'|x) \forall y'],$$

and agent type x 's expected discounted match payoff is

$$W_t^\theta(x) \equiv \mathbb{E} \left[\int_t^T e^{-\rho(\tau-t)} \sum_{y \in Y} \pi_\tau^\theta(y|x) P_t^\theta(d\tau)(y|x) + e^{-\rho(T-t)} \left(1 - \sum_{y \in Y} P_t^\theta(T)(y|x)\right) h_T \middle| \mathcal{F}_t \right].$$

(h_T is the μ_T -dependent terminal payoff, common to all agents, which accrues if unmatched at time T . In the simplest interpretation, set this to zero.)

If the optimal individual stopping rule is to accept a match whenever $S_t(x, y) + \xi \geq 0$, optimal stopping exhibits a striking symmetry: agent type x will only want to match with agent type y , if agent type y wants to match with agent type x ; two agents would never disagree. If so, the individually optimal matching threshold and the equilibrium matching rule coincide. That this is indeed the case is established by the following lemma.

Lemma 1 (Intratemporal efficiency). $W_t^{\theta^*}(x) = \sup_{\theta} W_t^\theta(x)$ where $\theta_t^*(x, y) = -S_t(x, y)$.

The result derives from the efficiency of the Nash bargaining sharing rule: if there exist gains from matching, the transfer rule is such that both parties stand to benefit. We henceforth omit all superscripts as equilibrium matching thresholds will throughout be governed by θ_t^* . The time t value of search is given by

$$V_t(x) \equiv \sup_{\theta} W_t^\theta(x). \quad (4)$$

As we go forward, a final remark is in order.

Remark 2. *The defining quantities of the model, notably μ_t and V_t , are stochastic processes and depend on the particular realization of the sample path of $(B_\tau)_{0 \leq \tau \leq t}$. Stipulated inequalities must hold for all realizations $\omega \in \Omega$, unless we suppose by contradiction; by this we mean that there exists $\omega \in \Omega$ such that the reverse inequality holds at time t for the event $\cap_{F \in \mathcal{F}_t; \omega \in F} F$.*

3 Equilibrium existence and uniqueness

3.1 Definition of equilibrium

An equilibrium is a characterization of matching patterns, i.e., specifies who matches with whom. Matching patterns derive from individual choices. Lemma 1 asserts that optimal choices adhere to the non-negative ex-post surplus rule: two agents match if and only if ex-post surplus is non-negative. It follows that matching decisions are a function of agents' value of search, which in turn depends on the evolution of the search pool.

Given an optimal matching threshold θ^* , an equilibrium is taken to be double (μ, V) where μ is given by (2) and V is given by (4).

Noticeably, the evolution of the search pool and the agents' value of search interact with one another: agents' evaluate their value of search under expectations which must conform to the population dynamics. Population dynamics exhibit a higher hazard rate and increased match creation if the value of search is low. This creates a feedback loop between population dynamics which move forward from the origin of time and values of search which are constructed backward from a final time T . The forward-backward structure is a feature of virtually all dynamic general equilibrium models under rational expectations. Adopting terminology from physics, some authors refer to such a system as a mean field game.

In what follows we seek to represent (μ, V) as a system of forward-backward stochastic differential equations (FBSDEs). This allows to connect our problem of equilibrium existence and uniqueness to an established literature in mathematics which studies the well-posedness of systems of FBSDEs.

3.2 System of FBSDEs

The population dynamics were given by a forward stochastic differential equation, looking forward from some initial condition μ_0 . For analytical purposes we find it more convenient to work with a monotone transformation of $\mu_t(x)$ instead. Thus define $\gamma_t(x) = \ln \mu_t(x)$. Ito's lemma establishes that

$$\gamma_t(x) = \gamma_0(x) + \int_0^t \left[- \sum_{y \in Y} \beta_\tau e^{\gamma_\tau(y)} m_\tau(x, y) + \eta_\tau(x) - \frac{(\sigma_\tau(x))^2}{2} \right] d\tau + \int_0^t \sigma_\tau(x) dB_\tau,$$

where $\gamma_0(x) = \ln \mu_0(x)$. The motivation for this transformation is straightforward: $d\gamma_t(x)$ can be thought of as the flow percentage change of agents in the search pool whereas $d\mu_t(x)$ captures the absolute flow change in the mass of agent types x . In that we focus on the former (joint with the assumption that $\sigma_\tau(x)$ is element in some positive bounded interval), we rule out that (i) the mass of agents searching becomes negative due to large, negative realizations of random entry, and (ii) accommodate a search pool which may grow unboundedly large as long as the expected growth rate remains bounded.

We now establish that the continuation value of search can be expressed as a backward stochastic differential equation. This is called a backward stochastic differential equation, be-

cause the solution is constructed backwards from the terminal time T .

Lemma 2. *The value of search is, given μ_t , the solution to the following backward stochastic differential equation:*

$$V_t(x) = h_T + \int_t^T \left[\sum_{y \in Y} [\pi_\tau(y|x) - V_\tau(x)] \beta_\tau e^{\gamma_\tau(y)} m_\tau(x, y) - \rho V_\tau(x) \right] d\tau - \int_t^T Z_\tau(x) \cdot dB_\tau.$$

The lemma follows from an application of the martingale representation theorem which can be likened to the fundamental theorem of calculus in stochastic analysis. A transformation of $V_t(x)$ can be shown to be a martingale. Then $Z_t(x)$ can be thought of as its derivative.

To summarize, this establishes that an equilibrium is a solution to the following system of forward-backward stochastic differential equations.

Corollary 1. *$(\mu, V) \equiv (\exp \gamma, V)$ is an equilibrium if and only if (γ, V, Z) solves the following system of forward-backward stochastic differential equations*

$$\begin{cases} \gamma_t(x) = \gamma_0(x) + \int_0^t \left[- \sum_{y \in Y} \beta_\tau e^{\gamma_\tau(y)} m_\tau(x, y) + \eta_\tau(x) - \frac{(\sigma_\tau(x))^2}{2} \right] d\tau + \int_0^t \sigma_\tau(x) dB_\tau \\ V_t(x) = h_T + \int_t^T \left[\sum_{y \in Y} [\pi_\tau(y|x) - V_\tau(x)] \beta_\tau e^{\gamma_\tau(y)} m_\tau(x, y) - \rho V_\tau(x) \right] d\tau - \int_t^T Z_\tau(x) \cdot dB_\tau. \end{cases} \quad (\text{E})$$

where $\gamma_0 = \exp \mu_0$ and h_T are given.

3.3 Regularity conditions

In most applications economic fundamentals $\beta, \eta, \Xi, \alpha, \sigma, h$ will be time- and state-invariant parameters. A theory which treats those as parameters is meaningful, no intuition is lost focusing on this case, and numerical examples will consider them as such. For example, consider

$$(\eta_t, \alpha_t, \sigma_t, h_T) \equiv (0, \frac{1}{2}, \epsilon, 0).$$

Well-posedness and therefore the existence of a unique equilibrium will obtain with these parameter values as long as the meeting rate is bounded (as is readily implied by linear search or truncated quadratic search for instance), and Ξ_t captures a standard bounded distribution (such as a truncated normal with support exceeding deterministic output).

However, as theorists we strive for generality. As our theoretical results on sorting and equilibrium existence and uniqueness do not require economic fundamentals to be time- and state-invariant, we posit a more general relationship between economic fundamentals and both time t and state variables. Because it affords a cleaner representation, we consider here the log of the population, i.e., γ_t not μ_t . The relevant time t Markovian states are given by $\gamma_t \in \mathbb{R}^N$ and $V_t \in \mathbb{R}^N$. We will subsequently and to the extent that it is necessary impose regularity conditions

that govern this relationship.

$$\begin{aligned}
\beta &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}_+ & \text{where } \beta_t &\equiv \beta(t, \gamma_t, V_t) \\
\eta &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}_+^N & \text{where } \eta_t &\equiv \eta(t, \gamma_t, V_t) \\
\alpha^X, \alpha^Y &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow [0, 1] & \text{where } \alpha_t^X &\equiv \alpha^X(t, \gamma_t, V_t), \alpha_t^Y \equiv \alpha^Y(t, \gamma_t, V_t) \\
\Xi &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \{\text{cdf's on } \mathbb{R}\} & \text{where } \Xi_t &\equiv \Xi(t, \gamma_t, V_t) \\
\sigma &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}_+^N & \text{where } \sigma_t &\equiv \sigma(t, \gamma_t, V_t) \\
h &: \mathbb{R}^N \rightarrow \{\mathcal{F}_T\text{-measurable rv on } \mathbb{R}\} & \text{where } h_T &\equiv h(\gamma_T)
\end{aligned}$$

Here we write $\beta_t \equiv \beta(t, \gamma_t, V_t)$ to make explicit that whenever we write β_t we subsume the dependence of the meeting rate on the mass of agents searching at time t as well as their time t option value of search. The same remark applies for the functional relationships governing all other fundamentals.

We make the following assumption with regard to $\beta, \eta, \alpha, \Xi, \sigma, h$. For brevity of notation, denote $\Theta_t = (t, \gamma_t, V_t)$.

Assumption 1 (regularity condition). • β and η are non-negative, continuous in γ_t and Lipschitz continuous with constants L^β and L^η respectively in V_t ;

- there exists a constant $K^\beta > 0$ such that $\beta_t(\Theta_t) \leq \frac{K^\beta}{\sum_{z \in X \cup Y} \exp \gamma_t(z)}$;
- there exists a constant $K^\eta > 0$ which bounds η from above;
- Both α^X and α^Y are in $[0, 1]$ with $\alpha^X + \alpha^Y \leq 1$, Lipschitz continuous in γ_t with constant L^α and continuous in V_t ;
- Ξ_t admits a density Ξ'_t with bounded support $[\underline{\xi}, \bar{\xi}]$;
- Ξ is Lipschitz continuous in both (γ_t, V_t) and ξ with constant L^ξ , i.e. $(\gamma_t, V_t) \mapsto \Xi(t, \gamma_t, V_t)(\xi)$ is Lipschitz continuous for all $t \in [0, T]$ and $\xi \in [\underline{\xi}, \bar{\xi}]$, and $\xi \mapsto \Xi(t, \gamma_t, V_t)(\xi)$ is Lipschitz continuous for all $t \in [0, T]$ and $(\gamma_t, V_t) \in \mathbb{R}^N \times \mathbb{R}^N$;
- Ξ' is Lipschitz continuous in (γ_t, V_t) with constant L^ξ , i.e. $(\gamma_t, V_t) \mapsto \Xi'(t, \gamma_t, V_t)(\xi)$ is Lipschitz continuous for all $t \in [0, T]$;
- σ is continuous in (t, γ_t, V_t) and bounded from below and above by strictly positive constants $\underline{\sigma}$ and $\bar{\sigma}$;
- There exists a non-negative constant L^σ such that $\|\sigma(t, \gamma_t, V_t) - \sigma(t, \gamma'_t, V'_t)\| \leq L^\sigma (\|\gamma_t - \gamma'_t\|^2 + \|V_t - V'_t\|^2)$ uniformly for all $(t, \gamma_t, V_t) \in [0, T] \times \mathbb{R}^N \times \mathbb{R}^N$;
- h is bounded and Lipschitz continuous with constant L^h .

For analytical convenience, we also require that the support of Ξ_t admits large positive and negative values with sufficiently flat tails, such that $m_t(x, y)$ is always bounded away from zero and one.

3.4 Existence and uniqueness of equilibrium

With assumption 1 in place, we are ready to establish the existence of a unique equilibrium.

Theorem 1. *If regularity condition 1 is in place, there exists a unique equilibrium.*

Our existence result derives from an application of a deep theorem on the well-posedness of systems of FBSDEs due to ?. Refer to the appendix and related literature for further discussion.

4 Single-peaked preferences and single-crossing

In the next section we will introduce our definition of probabilistic positive assortative matching. In the present section we will concern ourselves with what can be seen as an intermediate step to our analysis: We introduce the notion of preferences over matched partners absent knowledge of the idiosyncratic pair-specific production shock. As we shall see those preferences are interchangeably represented by expected payoffs, the surplus function, and (perhaps closest to being identifiable) the probabilistic matching function. We will discuss two well-known ordinal properties of said preferences. In our view, those properties are of independent (and empirical interest). They constitute equilibrium sorting (and should thus be tested), but are weaker than thereafter introduced notion of probabilistic assortative matching.

4.1 The preference relation

Our definition of probabilistic positive assortative matching, as defined in the next section, can be equivalently recast in terms of ordinal properties of preferences. We begin our analysis by studying said preferences. We write $y_2 \succeq_t^x y_1$ when agent type x prefers to meet y_2 over y_1 at time t . A utility representation of said preference relation is readily given by the expected payoff conditional on meetings: $y \mapsto V_t(x) + m_t(x, y)(\pi_t(y|x) - V_t(x))$. When one is certain to meet a type y at time t , one always retains one's option value of search $V_t(x)$. Moreover, with probability $m_t(x, y)$ the productivity shock ξ will be large enough for agent types x and y to form a match, in which case agent type x receives the expected match payoff $\pi_t(y|x)$.

We posit sufficient conditions on the degree of complementarity of the ex-ante output, f , under which the preference relation \succeq_t^x satisfies two, arguably classical, ordinal properties: (i) *single-crossing* and (ii) *single-peaked preferences*. Preferences exhibit single-crossing if for two possible types $y < y'$, there exists a threshold type \hat{x} such that all types x smaller than \hat{x} prefer to meet y , and all other types x prefer to meet y' . This conveys the idea that higher types x prefer to meet higher types y . Preferences are single-peaked, if each agent type x has an ideal partner, denoted $y_t(x)$, such that whenever agent type x compares two partners that are both to the right or to the left of $y_t(x)$, she strictly prefers whichever partner is closest to $y_t(x)$.

Since both (i) and (ii) are ordinal properties of the preference relation \succeq_t^x , they hold true for any utility representation, three of which are of interest:

Remark 3. *The preference relation \succeq_t^x is represented by any of the following:*

- *the expected payoff conditional on meetings, $V_t(x) + m_t(x, y)(\pi_t(y|x) - V_t(x))$,*

- the probabilistic matching function, $m_t(x, y)$,
- the ex-ante surplus, $S_t(x, y)$,

for one is a monotonic transformation of the other.¹³

Note that this remark relies on the assumption that idiosyncratic shocks are independently and identically distributed across agents. In what follows, we prove single-crossing and single-peaked preferences for the ex-ante surplus, $y \mapsto S_t(x, y)$. This ensures that the expected payoff conditional on meetings as well as the probabilistic matching function also satisfy (i) and (ii). Refer to figure 2 for an illustration of these properties for two realizations of random ex-ante surplus $S_t(x_1, \cdot)$ and $S_t(x_2, \cdot)$.




Figure 2: Single-peaked preferences and single-crossing

4.2 Single-crossing

We begin our analysis with single-crossing. This property stipulates that higher types x exhibit a preference for meeting with higher types y , which is suggestive of preferences for positive assortative matching.

Definition 1. $S_t(x, y)$ satisfies single-crossing if for any two $x_1 < x_2$ and $y_1 < y_2$, both

$$\begin{aligned} S_t(x_2, y_1) - S_t(x_1, y_1) \geq 0 &\Rightarrow S_t(x_2, y_2) - S_t(x_1, y_2) > 0, \\ S_t(x_1, y_2) - S_t(x_1, y_1) \geq 0 &\Rightarrow S_t(x_2, y_2) - S_t(x_2, y_1) > 0. \end{aligned}$$

In a similar vein, $S_t(x, y)$ satisfies weak single-crossing if

$$\begin{aligned} S_t(x_2, y_1) - S_t(x_1, y_1) > 0 &\Rightarrow S_t(x_2, y_2) - S_t(x_1, y_2) \geq 0, \\ S_t(x_1, y_2) - S_t(x_1, y_1) > 0 &\Rightarrow S_t(x_2, y_2) - S_t(x_2, y_1) \geq 0. \end{aligned}$$

In particular, it obtains when the ex-ante match output f is supermodular.

¹³This is because we can write the expected payoff conditional on meetings as an increasing function of the ex-ante surplus, $S_t(x, y) \mapsto V_t(x) + \int_{-S_t(x, y)}^{+\infty} \alpha_t^X(S_t(x, y) + \xi_t) \Xi_t(d\xi)$; likewise for the probabilistic matching function, $S_t(x, y) \mapsto 1 - \Xi_t(-S_t(x, y))$.

Definition 2 (supermodular). *Ex-ante match output f is supermodular if for any two $x_2 > x_1$ and $y_2 > y_1$ we have*

$$f(x_2, y_2) + f(x_1, y_1) > f(x_1, y_2) + f(x_2, y_1).$$

Supermodularity is a well-known property in the theory of matching. In a frictionless matching market the type with whom one matches is also the type with whom one generates the greatest match payoff. Under supermodularity, Becker (1973) shows that said type is non-decreasing in one's own type. The next proposition generalizes this result to an environment with search frictions.¹⁴

Proposition 1 (Single-crossing). *Supermodularity and single-crossing relate as follows:*

1. $f(x, y)$ is supermodular;
1. \Leftrightarrow 2. $S_t(x, y)$ is supermodular;
2. \Rightarrow 3. $S_t(x, y)$ satisfies single-crossing,
3. \Rightarrow 4. there exists $\hat{y}_t(x) \in \arg \max_{y \in Y} S_t(x, y)$ which is non-decreasing in x .

Proof. Recall that $S_t(x, y) = f(x, y) - V_t(x) - V_t(y)$ from which 1. \Leftrightarrow 2. readily follows. Then 2. \Rightarrow 3. is immediate. To see 3. \Rightarrow 4., denote

$$S_t^\omega(x, y) \equiv \mathbb{E}[S_t(x, y) \mid \cap_{F \in \mathcal{F}_t; \omega \in F} F]$$

a particular realization of the ex-ante surplus. Suppose by contradiction that no such $\hat{y}_t(x)$ can be chosen. Then, in particular, there exist $x_2 > x_1$ and $\omega \in \Omega$ such that $y_1 \in \arg \max_y S_t^\omega(x_1, y)$, $y_2 \in \arg \max_y S_t^\omega(x_2, y)$, yet $y_1 > y_2$. Then $S_t(x_1, y_1) \geq S_t(x_1, y_2)$. Hence single-crossing implies that

$$S_t^\omega(x_1, y_1) \geq S_t^\omega(x_1, y_2) \quad \Rightarrow \quad S_t^\omega(x_2, y_1) > S_t^\omega(x_2, y_2).$$

But then y_2 cannot be the ex-ante strictly preferred partner's type of agent type x_2 , which is absurd. \square

This result demonstrates that supermodularity is a first-order reason for positive assortative matching, even in the presence of search frictions, aggregate uncertainty, and ex-post heterogeneity.

4.3 Single-peaked preferences

We now turn our attention to single-peaked preferences. Economically speaking, this captures the idea that each agent type x has an ex-ante preferred type such that agents more closely resembling her preferred type $y_t(x)$ are preferred over those agent types that do so to a lesser

¹⁴A similar result has been reported in the working paper by Shimer and Smith (2000) at the steady state.

degree. To introduce this notion formally, we make use of the finite differences operator: define $\Delta_y h(x, y) \equiv h(x, y_+) - h(x, y)$ for arbitrary function h . Here y_+ denotes the smallest type $y \in Y$ strictly greater than y . (Similarly, denote y_- the greatest type $y \in Y$ strictly smaller than y .)

Definition 3 (Single-peaked). $y \mapsto S_t(x, y)$ is single-peaked if there exists $\hat{y}_t(x)$ such that $\Delta_y S_t(x, y)$ is strictly positive for all $y < \hat{y}_t(x)$ and strictly negative for all $y > \hat{y}_t(x)$.

In the same vein, we say that a function is weakly single-peaked if the preceding holds for weakly positive and weakly negative differences. Finally, throughout refer to $y_t(x) = \arg \max_y S_t(x, y)$ and $x_t(y) = \arg \max_x S_t(x, y)$ as the correspondences giving the preferred partner's types. If preferences are (weakly) single-peaked, then those correspondences are convex-valued.

Single-peaked preferences are a much more demanding property to satisfy than single-crossing. As already observed by Shimer and Smith (2000), a given type x may express a desire to match with high and low types y , but not so with intermediate types, even when ex-ante match output is supermodular.¹⁵ As a keystone of their analysis, Shimer and Smith (2000) report sufficient conditions on f under which single-peaked preferences obtain nonetheless at the steady state. In this subsection, we prove that the sufficiency of those conditions extends to the richer environment studied in this paper.

To begin with, we introduce two conditions on ex-ante output f which ensure that the ex-ante surplus is single-peaked. To simplify the interpretation of the conditions, we express them in terms of differences, $\Delta_y f$.¹⁶

Definition 4. $\Delta_y f(x, y)$ is log supermodular if for any $x_1 < x_2$ and $y_1 < y_2$ we have

$$\frac{\Delta_y f(x_2, y_2)}{\Delta_y f(x_2, y_1)} \geq \frac{\Delta_y f(x_1, y_2)}{\Delta_y f(x_1, y_1)}.$$

If both $\Delta_y f(x, y)$ and $\Delta_x f(x, y)$ are log supermodular, we write that $\Delta f(x, y)$ is log supermodular.

Definition 5. A function $\Delta_y f(x, y)$ is log supermodular in differences if for any $x_1 < x_2 < x_3$ and $y_1 < y_2$ we have

$$\frac{\Delta_y f(x_3, y_2) - \Delta_y f(x_2, y_2)}{\Delta_y f(x_2, y_2) - \Delta_y f(x_1, y_2)} \geq \frac{\Delta_y f(x_3, y_1) - \Delta_y f(x_2, y_1)}{\Delta_y f(x_2, y_1) - \Delta_y f(x_1, y_1)}.$$

Observe that if $\Delta_y f(x, y)$ is log supermodular in differences, then so is $\Delta_x f(x, y)$ (with the roles of x and y) reversed. We thus simply write that $\Delta f(x, y)$ is log supermodular in differences.

These conditions ensure that the degree of complementarity in output is sufficiently high. First note that log supermodularity is a condition on the levels of the difference in ex-ante output, whereas log supermodularity in differences is a condition on the curvature of the difference in ex-ante output. Specifically, it asserts that under log supermodularity in differences $y \mapsto$

¹⁵They refer to this property as convexity of matching sets, but their proof relies on establishing that ex-ante match output is quasi-concave, i.e., agents' preference relations are single-peaked.

¹⁶We say that these conditions hold strictly if the defining equation is strict across all types.

$\Delta_x f(x_2, y)$ is more convex than $y \mapsto \Delta_x f(x_1, y)$, which admits an interpretation in terms of risk preferences that we draw on subsequently. It follows that neither of the two conditions implies the other. Provided that f is differentiable with strictly positive derivatives, these conditions are equivalent to saying that $\partial_x f$ and $\partial_{xy}^2 f(x, y)$ are log supermodular. The most natural ex-ante output which satisfies these conditions is $f(x, y) = xy$ (as first explored by Lu and McAfee (1996)).

Theorem 2 (single-peaked preferences). *Suppose that ex-ante output f is supermodular and Δf is log supermodular and log supermodular in differences, with at least one holding strictly. Then $y \mapsto S_t(x, y)$ is single-peaked.*

The ansatz of the proof proceeds by contraposition: suppose preferences were not single-peaked. Then there exists a type x and three adjacent agent types $y_1 < y_2 < y_3$ such that both $y_1 \succsim_t^x y_2$ and $y_3 \succsim_t^x y_2$, which is equivalent to $S_t(x, y_1) > S_t(x, y_2)$ yet $S_t(x, y_3) > S_t(x, y_2)$, or

$$\Delta_y V_t(y_1) > \Delta_y f(x, y_1) \quad \text{yet} \quad \Delta_y V_t(y_3) < \Delta_y f(x, y_3). \quad (5)$$

At this stage, it is unclear how to proceed, for the difference in the value of search $\Delta_y V_t(y)$ is an obscure object. It is only implicitly characterized by a system of FBSDEs which usually does not admit a closed-form solution. We will next present two results, the mimicking argument and a comparison of lotteries, which lead to a contradiction of the alleged violation of single-peaked preferences, namely condition 5.

Step 1: the mimicking argument

The main ingenuity of our proof is to provide tight bounds on the difference in the value of search between two agents, $\Delta_y V_t(y)$, solely in terms of the ex-ante match-output. Those are formalized in the following lemma that we refer to as the *mimicking argument*. It is reminiscent of the mimicking argument developed by Bonneton and Sandmann (2019) for the non-transferable utility paradigm.

Lemma 3 (mimicking argument). *For every $x_1 \in X$ and time t there exist non-negative random weights on Y , denoted $Q_t(y|x_1) : \mathcal{F}_t \rightarrow \mathbb{R}_+$ with $\sum_{y \in Y} Q_t(y|x_1) \leq 1$ such that for every $x_2 \in X$ the following inequality obtains in equilibrium:¹⁷*

$$V_t(x_2) - V_t(x_1) \geq \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] Q_t(y|x_1).$$

In words, the lemma states that the difference in the value of search between any two agents can be bounded by a weighted sum over the difference in match-output. Crucially, holding agent type x_1 fixed, these weights hold true for any other type x_2 .

The proof of this result relies on the efficiency of Nash bargaining as asserted by lemma 1. This lemma states that the value of search dominates the discounted expected match payoff one

¹⁷In the current version of the proof we make the additional assumption that α^X and α^Y are constant. We will seek to generalize this restriction in the future.

would obtain by adopting any other acceptance threshold θ , i.e., $V_t(\cdot) \geq W_t^\theta(\cdot)$. In particular, this affords a lower bound on a given type x_2 's value of search, where x_2 "mimics" the equilibrium acceptance threshold of agent type x_1 . Formally, set $\theta_t(x_2, \cdot) = \theta_t^*(x_1, \cdot)$ for all $t \in [0, T]$. In doing so, agent type x_2 exactly replicates agent type x_1 's matching probabilities. Finally, subtracting $V_t(x_1)$ from $V_t(x_2) \geq W_t^\theta(x_2)$ gives

$$V_t(x_2) - V_t(x_1) \geq \mathbb{E} \left[\int_t^T e^{-\rho(\tau-t)} \sum_{y \in Y} [\pi_\tau^\theta(y|x_2) - \pi_\tau(y|x_1)] \underbrace{P_t(d\tau)(y|x_1)}_{x_1 \text{'s matching probability}} \middle| \mathcal{F}_t \right].$$

When two agent types adopt identical acceptance thresholds, they match conditional on identical realizations of production shocks ξ . Since those are identically distributed across agents, the difference in match payoffs admits further simplification:

$$\pi_\tau^\theta(y|x_2) - \pi_\tau(y|x_1) = (1 - \alpha_\tau^X)(V_\tau(x_2) - V_\tau(x_1)) + \alpha_\tau^X(f(x_2, y) - f(x_1, y)).$$

Substituting back into the preceding, we get

$$\begin{aligned} V_t(x_2) - V_t(x_1) &\geq \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] \mathbb{E} \left[\int_t^T e^{-\rho(\tau-t)} \alpha_\tau^X P_t(d\tau)(y|x_1) \middle| \mathcal{F}_t \right] \\ &\quad + \mathbb{E} \left[\int_t^T (V_\tau(x_2) - V_\tau(x_1)) e^{-\rho(\tau-t)} (1 - \alpha_\tau^X) \bar{P}_t(d\tau)(x_1) \middle| \mathcal{F}_t \right]. \end{aligned}$$

The first term is, as desired, a weighted sum over the difference in match output, the second term is (a weighted discounted future average of) the expected difference in value of search. Note that if the difference in value of search, $V_\tau(x_2) - V_\tau(x_1)$, did not fluctuate over time, rearranging the inequality readily gives the desired bound.

The remainder of the proof proceeds by iteration. First, when applying the above reasoning to (the weighted discounted future average of) $V_\tau(x_2) - V_\tau(x_1)$, we once more obtain two terms: the first term is, as desired, a weighted sum over the difference in match output, the second term is (a weighted discounted future average over averages of) the expected difference in value of search. After k iterations of this reasoning, we obtain $k + 1$ terms, the first k of which converge to a geometric series of the desired weighted sum over the difference in match output, and the $k + 1$ th term is shown to converge to zero as k grows large - which allows us to conclude the proof.

The *Matryoshka dolls* offer an intuitive representation of our iterative process. The Matryoshkas are a countable number of dolls, one wrapped into another. In our proof a doll corresponds to the difference in value of search (an object we do not understand), whereas its shell corresponds to a weighted sum over the difference in match output (an object we do understand). As we open the first doll, we keep the shell and find inside another doll, inside of which there is another doll. As iterations go on, we accumulate shells. In the meantime, dolls become exceedingly small, which eventually allows us to get rid of them altogether.

We find it worthwhile to mention that in Bonneton and Sandmann (2019) we had presented

a different mimicking argument. In both papers, the underlying idea is to let one agent mimic someone else in order to facilitate the comparison of the value of search across types. In the NTU paradigm, the result hinges on *payoff monotonicity*. Superior types, being more desirable, can exploit their superior match offerings and replicate match outcomes of any inferior type. In this paper, the result obtains for a different reason, namely the *efficiency of the Nash bargaining sharing rule*. The ensuing result is somewhat stronger: an agent can mimic not only inferior types, but also superior ones.¹⁸

Step 2: comparison of lotteries

With the mimicking argument at hand, we can reformulate the negation of singled-peaked preferences (5) as a choice over lotteries:

$$\sum_{z \in X} \Delta_y f(z, y_1) Q_t(z|y_2) > \Delta_y f(x, y_1) \quad \text{yet} \quad \sum_{z \in X} f(z, y_3) Q_t(z|y_2) < \Delta_y f(x, y_3). \quad (6)$$

To give this condition a figurative meaning, take $\Delta_y f(x, y)$ to be a (fictitious) utility of agent type y over partners x . Under this figurative meaning, condition (6) means that the greater agent type y_3 prefers the certain partner x over an uncertain partner z (given by weights $Q_t(\cdot|y_2)$), whereas the lower agent type y_1 prefers the uncertain partner z instead. To make use of the theory of risk and draw on a beautiful theorem by Pratt (1964), we seek to normalize the weights $Q_t(\cdot|y_2)$ such that they correspond to a lottery, i.e., sum to one. Pratt (1964) shows that *given arbitrary* $y_3 > y_2$ *the following statements are equivalent:*

1. Agent type y_3 is more risk-averse than agent type y_1 ; that is, y_3 does not accept a lottery that is rejected by y_1 ;
2. $x \mapsto \Delta_y f(x, y)$ is log supermodular in differences.

In the appendix we present detailed normalization arguments drawing on supermodularity of f and log supermodularity of Δf such that, indeed, (6) obtains for some lottery, i.e., $x \mapsto Q_t(\cdot|y_2)$ sums to one. Then for agent type y_1 not to prefer the lottery $Q_t(\cdot|y_2)$ that agent type y_2 rejected, it suffices to require that $\Delta_y f(x, y)$ is log supermodular in differences.

Shimer and Smith (2000) draw on the same connection with the theory of risk. They prove that identical conditions, supermodularity of f and log supermodularity and log supermodularity in differences of Δf , ensure single-peaked preferences at the steady state.

Whether those conditions are necessary is not resolved. It is an open question that shall be tackled in future research.

Although the ex-ante surplus, the matching probabilities, and the expected payoffs are single-peaked and satisfy single-crossing, the equilibrium transfers do not inherit these good properties; and the distribution over wages depends heavily on the distribution of pair-specific shocks (as shown in appendix C.3).

¹⁸In the NTU paradigm, the desired bound straightforwardly obtains from one agent mimicking someone else. This is due to the independence between payoffs and the value of search. Here we use an iterative argument to deal with such dependence.




Figure 3: \exists upper contour matching set: not a lattice \Leftrightarrow not all iso-lines \nearrow .

5 Probabilistic Assortative Matching

In this section we characterize aggregate matching patterns by drawing on the individual preferences studied in the preceding section. We then discuss when and why matching patterns exhibit PAM and NAM.

5.1 Definition of probabilistic assortative matching

Positive assortative matching is the intuitive notion that agents of similar characteristics tend to match with one another. However intuitive, the mathematical definition which corresponds to this notion is neither self-evident nor unique in a model that generates probabilistic matching patterns as we observe in the data.¹⁹

In this paper we are interested in agents' choices. Concurrently, we take positive assortative matching to be a property of preferences \succsim_t^x and \succsim_t^y . In our work the probabilistic matching function $y \mapsto m_t(x, y)$ is a utility representation of preferences \succsim_t^x . It is thus without loss of generality to cast our definition of PAM in terms of *ordinal* properties of the probabilistic matching function $m_t(x, y)$. This function can best be visualized by a time-moving heat map over pairs of types x and y where lighter colors correspond to greater match probabilities.²⁰

The predominant definition of PAM within the realm of random search is due to Shimer and Smith (2000). Our definition is a generalization of theirs to environments in which match outcomes are probabilistic, not binary. Motivated by their work, we represent agents' joint decisions $m_t(x, y)$ by upper contour matching sets, that is, all those pairs of types (x, y) which match with probability weakly greater than some given p :

$$U_t(p) = \{(x, y) : m_t(x, y) \geq p\}.$$

What Shimer and Smith (2000) require on $U_t(1)$ for positive assortative matching to obtain, we require on $U_t(p)$ for all p . In particular, we require as a necessary condition for PAM that upper

¹⁹In the frictionless environment proposed by Becker (1973) said notion admits a unique formal definition: every agent matches with a unique agent type which is increasing in his own type.

²⁰Another choice of domain of the definition is the flow rate of match creation, i.e. $m_t(x, y)\lambda_t(y)\mu_t(x)$, which we observe in the data. We choose not to do so because we believe it is important to disentangle choices, m_t , from the physical environment, μ_t and λ_t .

contour matching sets $U_t(p)$ be a lattice for all $p \in [0, 1]$.

Definition 6. *The upper contour matching set $U_t(p)$ is a lattice if for any types $x_1 < x_2$ and $y_1 < y_2$, the upper contour matching set $U_t(p)$ is such that*

$$(x_1, y_2) \in U_t(p) \text{ and } (x_2, y_1) \in U_t(p) \quad \Rightarrow \quad (x_1, y_1) \in U_t(p) \text{ and } (x_2, y_2) \in U_t(p).$$

Take four types $x_1 < x_2$ and $y_1 < y_2$, such that both non-assortative pairs, (x_1, y_2) and (x_2, y_1) , match with probability greater than p . Then the lattice property ensures that both assortative pairs, (x_1, y_1) and (x_2, y_2) , equally match with probability greater than p .

As a second necessary condition for PAM we require that individual upper contour matching sets be convex.²¹ Formally, the individual upper contour set, defined as $U_t(x; p) = \{y : m_t(x, y) \geq p\}$, is convex if for any three agent types $y_1 < y_2 < y_3$, $y_1 \in U_t(x; p)$ and $y_3 \in U_t(x; p)$ imply that $y_2 \in U_t(x; p)$. This conveys the idea that the greatest match probabilities of any fixed agent type x are concentrated around an interval of types y , the ex-ante ideal partners for type x . The more a type differs from the ideal partners y , the lower is the probability of matching.

To summarize, we generalize the definition of PAM provided by Shimer and Smith (2000) as follows:

Definition 7 (PAM). *There is probabilistic positive assortative matching (PAM) if*

- (i) *the upper contour matching set $U_t(p)$ is a lattice,*
- (ii) *individual upper contour matching sets $U_t(x; p)$ are convex for all agent types,*

for all $p \in [0, 1]$.

We wish to emphasize that this definition is about ordinal properties of the preference relations \succeq_t^x , not its particular representation: a given utility representation of \succeq_t^x satisfies the PAM-defining properties if and only if every other representation satisfies those properties. In particular, we could equivalently show that all upper contour sets of the ex-ante surplus, $S_t(x, y)$, are a lattice and satisfy individual convexity. We chose to state the definition using the probabilistic matching function, rather than the preference relation, as to make explicit the connection between Shimer and Smith's definition and ours.²² In the subsequent section, we show that PAM can equivalently be stated in terms of intuitive conditions on the preference relations \succeq_t^x .

²¹In Shimer and Smith (2000), the definition of PAM tacitly assumes convexity. Indeed, the convexity of $U_t(x; 1)$ is implied by the lattice property in their framework. Here is a sketch of the argument: in a stationary environment with binary matching probabilities, $U_t(x, 1)$ is necessarily non-empty (for otherwise the agent searches forever). More generally, when the individual upper contour matching set $U_t(x, p)$ is non-empty for every type x , then $U_t(p)$ being a lattice implies that $U_t(p)$ is convex (proof in proposition 4 in the appendix). In non-stationary and/or probabilistic environment, individual upper contour sets can be non-empty; the lattice property does not imply convexity of individual upper contour matching sets.

²²In contrast to our definition, Shimer and Smith do not cast their definition in terms of preferences. If one were to introduce an identical preference relation in their framework as we do, i.e. a ranking over the preferred partner's type that one would like to meet, their definition of PAM would not be an ordinal property of said preference relation. And their binary $m_t(x, y)$ as opposed to their $S_t(x, y)$ would not be a utility representation of said preference relation. This follows from the simple fact that their definition of PAM bears on properties of a single upper counter set, $\{y : S_t(x, y) \geq 0\}$. In order to characterize ordinal properties of a preference relation via upper counter sets, we know from consumer choice theory that imposed properties must be true for all such sets.

For illustrative purposes, it is convenient to visualize PAM by drawing on lower and upper level lines: $l_t(x; p) = \min\{y : m_t(x, y) \geq p\}$ and $u_t(x; p) = \max\{y : m_t(x, y) \geq p\}$, where $l_t(x, p)$ ($u_t(x; p)$) is the smallest (highest) type y with who agent type x matches with probability greater than p . As shown in proposition in the appendix, provided that individual counter sets are convex, there is PAM if and only if $x \mapsto l_t(x; p)$ and $x \mapsto u_t(x; p)$ are non-decreasing for all p . This is depicted in figure ???. One can check visually that the upper contour matching set from figure 3 (center) is a lattice as both upper and lower level lines are increasing. However, the upper contour matching set in figure 3 (right) is not. The lattice property is upset for precisely those agent types for which level lines are decreasing. In particular, the heat map depicts an instance in time where PAM is not satisfied.

5.2 Preferences for assortative matching

In section 4 we had studied single-peaked preferences and single-crossing. As it turns out, these two properties, albeit necessary, do not suffice to establish PAM. The reason is that both are comparative properties regarding preferences *within* a given population: single-crossing is concerned with how a lower type x_1 's preferences compare to a higher type x_2 's preferences; single-peaked preferences are concerned with preferences of a single agent type. The lattice property, in contrast, is a property that considers both populations jointly. It bears on how preferences *across* populations relate to one another.²³ In order to provide an equivalent preference representation of PAM, we need not only weakly single-peaked preferences and weak single-crossing, but also a condition that ties preferences across the two populations. This condition, which to the best of our knowledge is new to the matching literature, is *reciprocity* of matching preferences.

Definition 8 (reciprocity). *We say that \succeq_t^x and \succeq_t^y as represented by $y \mapsto m_t(x, y)$ and $x \mapsto m_t(x, y)$ exhibit reciprocity if for arbitrary types $(x, y) \in X \times Y$*

$$y < \underline{y}_t(x) \Rightarrow \underline{x}_t(y) \leq x \quad \text{and} \quad y > \bar{y}_t(x) \Rightarrow \bar{x}_t(y) \geq x,$$

where $\underline{y}_t(x) = \min \arg \max_y m_t(x, y)$ and $\bar{y}_t(x) = \max \arg \max_y m_t(x, y)$; likewise for $\underline{x}_t(y)$ and $\bar{x}_t(y)$.

Interchangeably, we say that a function is reciprocal if preferences represented by said function are reciprocal. This definition is necessarily tangled, for it must grapple with the subtleties of indifference between several agent types. In words, reciprocity means: if a given agent type y is inferior to an agent type x 's set of preferred types, then such type y must have preference to meet someone inferior to type x ; likewise, if a given agent type y is superior to an agent type

²³A different way of making this point is to say that single-peaked preferences and single-crossing are *ordinal* properties of an individual agent's utility. The lattice property, while being an ordinal property of joint utilities as discussed in the preceding subsection, can not be cast as an ordinal property of an individual agent's utility: take arbitrary utility representations $u_t^x(y)$ of preference relation \succeq_t^x , e.g. let $u_t^x(y) \equiv m_t(x, y)$, and take arbitrary monotone transformations T^x for each x . Then $T^x \circ u_t^x(y)$ is another representation of preference relation \succeq_t^x which preserves single-peaked preferences and single-crossing from the viewpoint of agent types x (but not necessarily from the viewpoint of agent types y): $y \mapsto T^x \circ u_t^x(y)$ is single-peaked and $T^{x_1} \circ u_t^{x_1}(y_2) \geq T^{x_1} \circ u_t^{x_1}(y_1)$ implies that $T^{x_2} \circ u_t^{x_2}(y_2) > T^{x_2} \circ u_t^{x_2}(y_1)$. Meanwhile, the lattice property of upper contour matching sets $U_t(p)$ induced by some $m_t(x, y)$ is not preserved under x -contingent monotone transformations of $y \mapsto m_t(x, y)$.

x 's set of preferred types, then such type y must have a preference to meet someone superior to type x .

Having introduced all the pieces, we can now make precise the relationships between introduced primitives on preferences and the constituent parts of the definition of PAM. We begin by stating the obvious:

Remark 4. \succeq_t^x is weakly single-peaked if and only if $U_t(x; p)$ is convex for all $p \in [0, 1]$.

The link between preferences and PAM is established by the following proposition:

Proposition 2. *Suppose that upper contour sets $U_t(p)$ are a lattice for all p . Then preferences \succeq_t^x and \succeq_t^y satisfy reciprocity and weak single-crossing. Suppose that preferences \succeq_t^x and \succeq_t^y satisfy reciprocity, weak single-crossing and are weakly single-peaked. Then upper contour sets $U_t(p)$ are a lattice for all p .*

The key insight from this subsection is then that PAM admits an equivalent representation, expressed in terms of primitives of preferences:

Corollary 2 (Preference representation result of PAM). *There is PAM if and only if preferences are weakly single-peaked, and satisfy reciprocity and weak single-crossing.*

Single-peaked preference is the main sufficient condition for PAM in Shimer and Smith (2000), as it implies that individual matching set $U_t(x, 1)$ are convex. But in their context, it is not necessary, as single-peaked preference is not implied by the convexity of individuals matching set. In Atakan (2006), individual matching sets are proven to be convex when the match output is supermodular, but supermodularity alone does not give single-peaked preference.

Lack of reciprocity

Reciprocity owes its name to the fact that, provided that preferences are also weakly single-peaked and satisfy weak single-crossing, preferences over the preferred partner's type are reciprocated: every agent type x 's set of preferred partner's type is at the very least adjacent to some agent type y for whom x is a preferred partner. To formalize this claim, define \mathcal{P}_t the set of pairs of types where at least one type is the other's preferred partner's type:

$$\mathcal{P}_t \equiv \{(x, y) : x \in x_t(y) \text{ or } y \in y_t(x)\}.$$

We establish that PAM has the following, implausibly strong, property:

Proposition 3. *Suppose there is PAM. Then,*

- \mathcal{P}_t is convex in the x and the y dimension, i.e., for $x_1 < x_2 < x_3$, if (x_1, y) and (x_3, y) are in \mathcal{P}_t , then so is (x_2, y) , likewise for types y ,
- \mathcal{P}_t is path-connected along the horizontal, vertical, or diagonal,²⁴

²⁴Formally, for any two $(\underline{x}, \underline{y})$ and (\bar{x}, \bar{y}) there exists a path $((\underline{x}, \underline{y}), (x_1, y_1), \dots, (x_K, y_K), (\bar{x}, \bar{y}))$ such that if (x', y') succeeds (x, y) , then $(x', y') \in \{(x_-, y_-), (x_-, y), (x_+, y), (x_+, y_+)\}$, where y_+ denotes the smallest type $y \in Y$ strictly greater than y . And, Here y_- denotes the smallest type $y \in Y$ strictly greater than y .

- (x_{\min}, y_{\min}) and (x_{\max}, y_{\max}) are in \mathcal{P}_t .

Visually, this property means that matching patterns, as depicted by a heat-map, can be likened to a mountain ridge. This ridge ranges from the pair composed of the lowest types, (x_{\min}, y_{\min}) , to the pair composed of the highest types, (x_{\max}, y_{\max}) . None of the three conditions is generically satisfied when there are single-peaked preferences that satisfy single-crossing, but not reciprocity.

We emphasize that the lack of reciprocal preferences, and thus PAM, arises generically for any $f(x, y)$ such that the preferences represented by $f(x, y)$, i.e., $y \mapsto f(x, y)$ and $x \mapsto f(x, y)$, do not satisfy reciprocity; that is, reciprocity of $f(x, y)$ is a necessary condition for PAM. This follows from the simple observation that when search frictions are sizeable and terminal values are zero, ex-ante surplus $S_t(x, y)$ is approximately equal to ex-ante match output $f(x, y)$. Thereby preferences over meetings are approximately given by those represented by $f(x, y)$.

Does reciprocity of $f(x, y)$, conversely, imply that preferences over meetings (as represented by $S_t(x, y)$ or $m_t(x, y)$) are reciprocal? In appendix D.6, we provide a counter-example which revolves around a symmetric population comprising three types; $y \mapsto f(x, y)$ is maximal for complementary types $y = x$ and $f(x, y)$ is supermodular (but does not satisfy log supermodularity in differences). We show that there exist economic fundamentals for which reciprocity does not obtain: low and high types exhibit self-preferences, whereas the intermediate type prefers to match with the lowest type. This shows that reciprocity of preferences over meetings is difficult to satisfy. We believe, however, that reciprocity of $f(x, y)$ and preceding complementarity conditions taken together are sufficient for reciprocity of preferences over meetings—and PAM.

In what follows we establish sufficiency for the last item, self-preferences of boundary types. It will play an important role in generalizing Shimer and Smith (2000)'s sorting result.

Lemma 4 (Boundary conditions). *Suppose that output f is supermodular. If $y \mapsto f(x_{\max}, y)$ is maximal for y_{\max} , then $\bar{y}_t(x_{\max}) = y_{\max}$; likewise, if $y \mapsto f(x_{\min}, y)$ is maximal for y_{\min} , then $\underline{y}_t(x_{\min}) = y_{\min}$.*

We stress that $y \mapsto f(x, y)$ non-increasing is a strong assumption as it rules out strict vertical differentiation of types. As a by-product of our analysis, we show this conditions, albeit sufficient is not necessary and can be weakened at the steady state. We prove that log supermodularity of f , as implied by their conditions under vertically differentiated types, renders this boundary condition obsolete (refer to appendix D.5 for the proof).

5.3 Main result

Our main result provides a comparative characterization of matching patterns when preferences are weakly single-peaked and satisfy single-crossing, but need not satisfy reciprocity. Those properties arise when f is supermodular and Δf is log supermodular and log supermodular in differences, with at least one holding strictly (as established by theorem 2 and proposition 1). To that end, we partition the set of pairs of types along the preferred partner's type correspondence, and show how match probabilities $m_t(x, y)$ compare across pairs of types. As a corollary to our comparative characterization of matching patterns, we can determine those pairs of types for

whom our theoretical definition of PAM obtains, and those for whom NAM (i.e., PAM upon reordering of types X) obtains.

To proceed, we partition the space of pairs $X \times Y$ into four regions, which partially overlap at the graphs of the preferred partner's types' correspondences $x_t(y)$ and $y_t(x)$, i.e., \mathcal{P}_t . We call the upper and lower enclosure of \mathcal{P}_t those sets of pairs of types that lie in between the two correspondences:

$$\bar{\mathcal{E}}_t \equiv \{(x, y) : y \leq \underline{y}_t(x) \wedge x \leq \underline{x}_t(y)\} \quad \text{and} \quad \underline{\mathcal{E}}_t \equiv \{(x, y) : y \geq \bar{y}_t(x) \wedge x \geq \bar{x}_t(y)\}.$$

Similarly we define its complements, the set above and below those correspondences the upper and lower outer set:

$$\underline{\mathcal{O}}_t \equiv \{(x, y) : y \leq \bar{y}_t(x) \wedge x \geq \underline{x}_t(y)\} \quad \text{and} \quad \bar{\mathcal{O}}_t \equiv \{(x, y) : y \geq \underline{y}_t(x) \wedge x \leq \bar{x}_t(y)\}.$$

Theorem 3 (characterization of matching patterns). *Suppose that preferences are weakly single-peaked. Then,*

- *match probabilities for pairs in $\underline{\mathcal{E}}_t$ rise for lower x and lower y , i.e., \swarrow ,*
- *match probabilities for pairs in $\bar{\mathcal{E}}_t$ rise for higher x and higher y , i.e., \nearrow ,*
- *match probabilities for pairs in $\bar{\mathcal{O}}_t$ rise for higher x and lower y , i.e., \searrow ,*
- *match probabilities for pairs in $\underline{\mathcal{O}}_t$ rise for lower x and higher y , i.e., \nwarrow .*

Suppose further that preferences satisfy single-crossing. Then both the upper and lower enclosure $\bar{\mathcal{E}}_t, \underline{\mathcal{E}}_t$ and the upper and lower outer set $\bar{\mathcal{O}}_t, \underline{\mathcal{O}}_t$ are a lattice.

Proof. The first statement is an immediate consequence of single-peaked preferences, of which $y \mapsto m_t(x, y)$ and $x \mapsto m_t(x, y)$ are representations. To see that the enclosures and outer sets are a lattice, observe that proposition 5 establishes an identical result: provided that $U_t(x; p)$ is convex, $U_t(p)$ is a lattice if and only if the p -level lines (which bound $U_t(p)$) are non-decreasing. Here, both enclosures and upper outer sets are non-empty by construction, convex due to single-peaked preferences and bounded by the lower and upper bounds of the preferred partner's type correspondences, $y \mapsto \underline{x}_t(y), \bar{x}_t(y)$ and $x \mapsto \underline{y}_t(x), \bar{y}_t(x)$. Proposition 1, item 4 establishes that those are non-decreasing, hence the result. \square

A characterization of matching patterns in terms of PAM and NAM readily follows from theorem 3. NAM is converse notion of PAM; there is NAM if by reversing the order of types for one side of the population, PAM obtains. Or, NAM arises if, given four types $x_1 < x_2$ and $y_1 < y_2$, the non-assortative pairs (x_1, y_2) and (x_2, y_1) have weakly greater match probability than at least one of the assortative pairs.

Corollary 3 (Equilibrium sorting). *Under theorem 3's assumptions, there is PAM in the outer sets $\bar{\mathcal{O}}_t$ and $\underline{\mathcal{O}}_t$, and NAM in the enclosure $\bar{\mathcal{E}}_t$ and $\underline{\mathcal{E}}_t$.*




Figure 4: Illustration of theorem 2: Single-crossing and single-peaked preferences

In the introduction we ascribed our results to a tension between two conflicting forces, type complementarity, i.e., $f(x, y)$ is supermodular, and vertical differentiation of types, i.e., $f(x, y)$ is increasing. We explored a channel through which search frictions give rise to unanimous preferences for superior types. If the value of search is low for everyone, then everyone will want to meet those agents with whom they produce the greatest ex-ante match output, not their complementary type. We wish to emphasize, however, that agents may likewise deviate and prefer inferior types over their complementary type—even when types are vertically differentiated: if the value of search is high for greater and low for inferior types, the loss in output of matching with inferior types may be outweighed by the lesser compensation those inferior types demand.²⁵

We caution that our main result should not be interpreted as a quantitative assessment. To us, PAM and NAM are ordinal properties. Their study gives rise to an ordinal characterization of sorting patterns that is robust to changes in the physical environment that do not concern $f(x, y)$. From an empirical perspective this is desirable, because it avoids mis-attributing changes in the frequency of matches of given characteristics x and y to changes in match complementarity $f(x, y)$.

5.4 Revisiting Shimer and Smith (2000)

The main contribution of Shimer and Smith (2000) is to identify sufficient conditions on $f(x, y)$ for which PAM obtains in the steady state and without pair-specific production shocks. Their set of conditions include the strong complementarity conditions required in theorem 2 and the

²⁵To obtain this, we can just consider the counter-example in appendix D.6, but change the payoff matrix. Now, $f(x, y) = xy$ and $q = 0.8$, $x_3 = 1$, $x_2 = 0.5$, $x_1 = 0.4$, we obtain the following: $V(x_2) = 0.4$, $V(x_1) = 0.07$, $V(x_3) = 0$, $S(x_2, x_2) = 0.11$, $S(x_1, x_2) = 0.13$, $S(x_3, x_2) = 0.03$. So the medium type prefers to match with the lowest type.

boundary conditions from lemma 4 (which prohibit vertically differentiated types).

Here we show that a version of Shimer and Smith's original result can be recovered in our more general environment under similar conditions as those imposed in their paper. We emphasize that those conditions are notably stronger than the assumptions made so far: we require their boundary conditions, and we also have to endow the type space with a distance and impose continuity of match output.

To proceed, let \hat{p}_t the min max-probability, the smallest of all agent's maximal match probabilities:

$$\hat{p}_t \equiv \min_x \max_y m_t(x, y)$$

Observe that, equivalently, $\hat{p}_t = \min_{(x,y) \in \mathcal{P}_t} m_t(x, y)$. Since $\mathcal{P}_t \subseteq \underline{\mathcal{E}}_t \cup \overline{\mathcal{E}}_t$, this probability is greater than the the smallest match probability in the enclosure of pairs (where PAM is upset):

$$p_t^e \equiv \min_{(x,y) \in \underline{\mathcal{E}}_t \cup \overline{\mathcal{E}}_t} m_t(x, y),$$

i.e., $\hat{p}_t \geq p_t^e$. p_t^e gives a threshold probability up to which upper counter matching sets do satisfy the PAM-defining probabilities. Indeed, under the assumptions of theorem 3, for all $p \leq p_t^e$, upper counter sets satisfy the PAM-defining properties. To reduce the significance of NAM, we seek to infer from large maximal match probabilities \hat{p}_t , that also p_t^e is large. If so, then (approximately) there is PAM. It turns out that this assertion requires two properties: first we need to endow the type space with a distance and require continuity of match output. Thus, denote d a distance on $X \times Y$. Secondly, we need to impose boundary conditions on preferences.

Assumption 2. *Assume that preferences satisfy the following:*

- *there are weakly single-peaked preferences that satisfy single-crossing;*
- $(x_{\min}, y_{\min}), (x_{\max}, y_{\max}) \in \mathcal{P}_t$

Moreover, we assume that output is uniformly continuous and Ξ_t is Lipschitz continuous:

- *the distance between horizontally or vertically adjacent types is less than $\frac{r}{N}$ where $\underline{N} \equiv \min\{N^X, N^Y\}$; $f(x, y)$ is Lipschitz-continuous with constant $\frac{q}{r}$;*
- Ξ_t *is Lipschitz-continuous with constant $\frac{L^\xi}{3q}$.*

As in Shimer and Smith, boundary conditions orient matching sets. Here this means that \mathcal{P}_t encloses the NAM-region $\underline{\mathcal{E}}_t \cup \overline{\mathcal{E}}_t$ both from above and below. As shown in lemma 4, self-preference at the boundary, i.e., $(x_{\min}, y_{\min}), (x_{\max}, y_{\max}) \in \mathcal{P}_t$, is satisfied when output $y \mapsto f(x_{\min}, y)$ is maximal for $y = y_{\min}$ and $y \mapsto f(x_{\max}, y)$ is maximal for $y = y_{\max}$.

Lemma 5. *Suppose that the preceding assumption is in place. Then*

$$p_t^e \geq \hat{p}_t - \frac{L^\xi}{\underline{N}}.$$

As an immediate consequence, the additional assumption ensures that \hat{p}_t gives a threshold probability up to which upper counter matching sets do satisfy the PAM-defining probabilities.

Theorem 4 (non-stationary Shimer and Smith with pair-specific production shocks). *Suppose the preceding assumption is in place. Then for all $p \leq \hat{p}_t - \frac{L\xi}{N}$, $U_t(p)$ and $U_t(x; p)$ satisfy the PAM-defining properties.*

This result encompasses the original result due to Shimer and Smith (2000). In the limit case when ξ is negligible and the economy is stationary, $\hat{p}_t \rightarrow 1$, for otherwise some agent type would never match, leading to a zero option value of search. As they consider a continuum of types, corresponding to $N \rightarrow \infty$, PAM can be recovered. More generally, our result establishes that with negligible pair-specific production shocks matching patterns exhibit PAM even in non-stationary environments (provided that search frictions are sufficiently large, so that no agent type can afford to reject all other agent types). In contrast, for diffuse pair-specific production shocks ξ , a large \hat{p}_t will be difficult to satisfy. Thus we do not view this result as a vindication of PAM, but rather as one exploring its limits.

If search frictions upset PAM, why does PAM obtain in Shimer and Smith (2000)? The key difference is that in their model, $m_t(x, y)$ is not a representation of preferences over meetings. Consequently, our characterization of PAM in terms of preferences does not apply to their environment. However, this does not invalidate the robust prediction of preferences for NAM, i.e., $\bar{\mathcal{E}}_t \cup \underline{\mathcal{E}}_t$ is non-empty. In other words, the preference for NAM is not an artefact of pair-specific shocks. It is already invisibly present in Shimer and Smith's model.

To clarify this point, it is instructive to draw on the deterministic match surplus. In both frameworks, for all pairs in the NAM region $\bar{\mathcal{E}}_t$ and $\underline{\mathcal{E}}_t$, the following property obtains: for all $x_2 > x_1$ and $y_2 > y_1$, and $k \in \mathbb{R}$,

$$S_t(x_1, y_1) \geq k \text{ and } S_t(x_2, y_2) \geq k \implies S_t(x_1, y_2) \geq k \text{ and } S_t(x_2, y_1) \geq k$$

In words, take four types $x_1 < x_2$ and $y_1 < y_2$, such that both assortative pairs, (x_1, y_1) and (x_2, y_2) , generate a deterministic surplus greater than a constant k . Then both non-assortative pairs, (x_1, y_2) and (x_2, y_1) , generate a deterministic surplus at least greater than k . What differs in both framework is how $S_t(x, y)$ translates into matching probabilities. In our model, due to pair-specific shocks, the probability of matching is a monotone transformation of the deterministic surplus. Thereby, we get,

$$m_t(x_1, y_1) \geq p \text{ and } m_t(x_2, y_2) \geq p \xrightarrow{\text{NAM}} m_t(x_1, y_2) \geq p \text{ and } m_t(x_2, y_1) \geq p,$$

which is precisely the theoretical definition of NAM. In Shimer and Smith (2000), all pairs in this subset match with certainty because the deterministic surplus is positive in this region. They need the aforementioned boundary conditions to guarantee that the surplus is indeed positive within $\bar{\mathcal{E}}_t$ and $\underline{\mathcal{E}}_t$. Therefore, matching probabilities are identical across pairs, so PAM, but also NAM obtains:

$$m_t(x_1, y_1) = 1 \text{ and } m_t(x_2, y_2) = 1 \xrightleftharpoons[\text{PAM}]{\text{NAM}} m_t(x_1, y_2) = 1 \text{ and } m_t(x_2, y_1) = 1,$$

As a result, predictions differ simply because in Shimer and Smith's model the probability of matching imperfectly captures the heterogeneity in surplus across pairs.

6 Conclusion

We have presented here two forays into the theory of random search and matching. First, we provided an existence and uniqueness result of equilibrium. This casts doubt on the robustness of multiple self-fulfilling (non-stationary) equilibrium paths frequently reported in the literature. Secondly, we formulated a new insight into assortative matching under Nash bargaining: search frictions impede assortative matching from occurring because frictions disproportionately erode the value of search and hence the bargaining power of more productive agents. In consequence, unproductive agents prioritize matching with productive agents. As a main result, we provided a comparative characterization of matching patterns that holds away from the steady state and regardless of the degree of search frictions.

Our objective was to provide a theory simple enough to allow for theoretical insights into sorting, yet rich enough to be plausibly at the origin of empirically observed matching patterns. In line with our objective, we proposed a theory in which match outcomes were probabilistic and the economy non-stationary. We hope that this theory will further invigorate the ongoing quest of identifying sorting in the labor market and consider with greater attention the role of pair-specific production shocks and non-stationary dynamics. We equally hope that our sorting result provides a useful lens through which empirical researchers can view matching patterns.

Our theoretical inquiry indicates several lines for future research. First, we conjecture that our definition of probabilistic assortative matching obtains under reasonable conditions over the match output in the frictionless model. If that is proven to be true, it could help us delineate more clearly the trade-offs involved when adopting the frictionless matching framework for empirical research. Secondly, our main sorting results makes a number of empirical predictions. As is well known, unobserved heterogeneity of workers' types renders the refutation of those predictions a difficult task. We view it as imperative however that the attempt be made to verify whether the ordinal properties asserted are empirically valid. Finally, we believe that our framework is a suitable benchmark for applied theory; it could be studied more extensively to investigate labor market interventions such as the minimum wage. Here sorting, in particular the reallocation of workers from less to more productive firms as well as the length of employment spells has an important role to play. Relatedly, our uniqueness result is a step forward in that direction as it potentially allows to study welfare.

Appendix

A Introductory proofs

The probability of meeting k agents with types $\{y_1, \dots, y_k\} \subset Y$ during time interval $(t_0, t_1]$ is given by

$$\frac{1}{k!} \prod_{\ell=1}^k \int_{t_0}^{t_1} \lambda_t(y_\ell) dt \exp \left\{ - \sum_{y \in Y} \int_{t_0}^{t_1} \lambda_t(y) dt \right\},$$

and follows from the definition of the inhomogenous Poisson point process.

A.1 Remark 1

Proof. Internal consistence of the model requires that, for all $x \in X$ and $y \in Y$,

$$\lambda_t(y)\mu_t(x) = \lambda_t(x)\mu_t(y),$$

where either side expresses the flow number of meetings between agent types x and y . Fix $x \in X$ for which $\mu_t(x) > 0$. Then for all $y \in Y$

$$\lambda_t(y) = \beta_t^Y \mu_t(y) \quad \text{where} \quad \beta_t^Y = \frac{\lambda_t(x)}{\mu_t(x)}.$$

Following identical arguments there exists β_t^X such that $\lambda_t(x) = \beta_t^X \mu_t(x)$. Substituting back into the balance condition where $\mu_t(x), \mu_t(y)$ are non-zero we note that $\beta_t \equiv \beta_t^X = \beta_t^Y$. \square

The proof implicitly assumes that there always exist x and y such that $\mu_t(x), \mu_t(y) > 0$. This will be a natural property of the model whenever the initial population of said types is non-zero.

As stated in the main text, the order of lemmata 1 and 2 is reversed for expositional purposes. We begin with lemma 2, which we use to establish lemma 1.

A.2 Proof of Lemma 2

Proof. We seek to transform the value of search into a martingale: according to Bayes' law for $\tau > t$, $P_0(\tau)(y|x) = P_0(t)(y|x) + \mathbb{P}(\text{don't match during } [0, t])P_t(\tau)(y|x)$, and therefore

$$P_t(\tau)(y|x) = \frac{P_0(\tau)(y|x) - P_0(t)(y|x)}{1 - \sum_{k=1}^N P_0(t)(k|x)}.$$

Then define

$$\begin{aligned}\tilde{V}_t(x) &\equiv e^{-\rho t} \left(1 - \sum_{k=1}^N P_0(t)(k|x)\right) V_t(x) \\ &= \mathbb{E} \left[\int_t^T \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) + e^{-\rho T} \left(1 - \sum_{k=1}^N P_0(T)(k|x)\right) h(\mu_T)(x) \middle| \mathcal{F}_t \right].\end{aligned}$$

The desired martingale is then given by

$$\begin{aligned}M_t(x) &\equiv \tilde{V}_t(x) + \int_0^t \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) \\ &= \mathbb{E} \left[\int_0^T \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) + e^{-\rho T} \left(1 - \sum_{k=1}^N P_0(T)(k|x)\right) h(\mu_T)(x) \middle| \mathcal{F}_t \right].\end{aligned}$$

And thus by the martingale representation theorem there exists a unique square integrable N -valued process $Z(x)$ such that

$$M_t(x) = M_T(x) - \int_t^T Z_\tau(x) \cdot dB_\tau.$$

Substituting back into the definition of $\tilde{V}_t(x)$, this gives

$$\begin{aligned}\tilde{V}_t(x) &= M_t(x) - \int_0^t \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) \\ &= M_T(x) - \int_t^T Z_\tau(x) \cdot dB_\tau - \int_0^t \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) \\ &= \tilde{V}_T(x) + \int_t^T \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) - \int_t^T Z_\tau(x) \cdot dB_\tau \\ &= e^{-\rho T} \left(1 - \sum_{k=1}^N P_0(T)(k|x)\right) h(\mu_T)(x) + \int_t^T \sum_{y \in Y} e^{-\rho \tau} \pi_\tau(y|x) P_0(d\tau)(y|x) - \int_t^T Z_\tau(x) \cdot dB_\tau,\end{aligned}$$

or, written in differential form,

$$d\tilde{V}_t(x) = - \sum_{y \in Y} e^{-\rho t} \pi_t(y|x) P_0(dt)(y|x) + Z_t(x) \cdot dB_t.$$

Now observe that due to Ito's lemma,

$$\begin{aligned} d\tilde{V}_t(x) &= d\left[e^{-\rho t}\left(1 - \sum_{k=1}^N P_0(t)(k|x)\right)V_t(x)\right] \\ &= e^{-\rho t}\left\{-\rho\left(1 - \sum_{k=1}^N P_0(t)(k|x)\right)V_t(x)dt - V_t(x)\sum_{y\in Y}P_0(dt)(y|x) + \left(1 - \sum_{k=1}^N P_0(t)(k|x)\right)dV_t(x)\right\}. \end{aligned}$$

By identifying the two we obtain

$$\begin{aligned} &-\sum_{y\in Y}e^{-\rho t}\pi_t(y|x)P_0(dt)(y|x) + Z_t(x) \cdot dB_t \\ &= e^{-\rho t}\left\{-\rho\left(1 - \sum_{k=1}^N P_0(t)(k|x)\right)V_t(x)dt - V_t(x)\sum_{y\in Y}P_0(dt)(y|x) + \left(1 - \sum_{k=1}^N P_0(t)(k|x)\right)dV_t(x)\right\}. \end{aligned}$$

Or, equivalently,

$$-\sum_{y\in Y}[\pi_t(y|x) - V_t(x)]\frac{P_0(dt)(y|x)}{1 - \sum_{k=1}^N P_0(t)(k|x)} + \rho V_t(x)dt + \tilde{Z}_t(x) \cdot dB_t = dV_t(x),$$

where

$$\tilde{Z}_t(x) = e^{\rho t}\frac{Z_t(x)}{1 - \sum_{k=1}^N P_0(t)(k|x)}.$$

Then note that $P_0(d\tau)(y|x)/(1 - \sum_{k=1}^N P_0(\tau)(k|x)) = P_t(d\tau)(y|x)$. Since the probability of meeting more than one agent in a given time interval of length ℓ is $o(\ell)$, $P_t(d\tau)(y|x)$ admits a neat representation: it is the hazard rate $\lambda_\tau(k|x)m_\tau(x, y)$ at which agent type x exits the search pool with an agent type y .

Substituting back into the former expression gives the BSDE

$$V_t(x) = h(\mu_T)(x) + \int_t^T \left[\sum_{y\in Y} [\pi_\tau(y|x) - V_\tau(x)]\lambda_\tau(y|x)m_\tau(x, y) - \rho V_\tau(x) \right] d\tau - \int_t^T \tilde{Z}_\tau(x) \cdot dB_\tau,$$

with endpoint constraint $V_T(x) = h(\mu_T)(x)$, as desired. \square

Since $V_t(x) = \mathbb{E}[V_t(x)|\mathcal{F}_t]$, taking expectations we deduce the following corollary of lemma 2:

Corollary 4.

$$V_t(x) = \mathbb{E}\left[h_T + \int_t^T \left[\sum_{y\in Y} [\pi_\tau(y|x) - V_\tau(x)]\lambda_\tau(y|x)m_\tau(x, y) - \rho V_\tau(x) \right] d\tau \middle| \mathcal{F}_t\right].$$

The corollary is an immediate implication of the lemma, since the expectation of stochastic integrals of square integrable processes is zero.

A.3 Proof of Lemma 1

Proof. Consider (for the purposes of this proof) expected ex-post surplus conditional on matching, $\hat{S}_\tau^\theta(x, y) \equiv \pi_\tau^\theta(y|x) - V_\tau(x)$. According to the preceding corollary,

$$\begin{aligned} & W_t^{\theta^*}(x) - W_t^\theta(x) \\ &= \mathbb{E} \left[\int_t^T \left\{ \sum_{y \in Y} [\hat{S}_\tau^{\theta^*}(x, y) m_\tau^{\theta^*}(x, y) - \hat{S}_\tau^\theta(x, y) m_\tau^\theta(x, y)] \lambda_\tau(y|x) - \rho [V_\tau(x) - W_\tau^\theta(x)] \right\} d\tau \middle| \mathcal{F}_t \right]. \end{aligned}$$

Note that by virtue of the choice of θ^* :

$$\hat{S}_\tau^{\theta^*}(x, y) m_\tau^{\theta^*}(x, y) - \hat{S}_\tau^\theta(x, y) m_\tau^\theta(x, y) = \int_{\theta_\tau^*(x, y)}^{\theta_\tau(x, y)} \frac{f(x, y) + \xi - V_\tau(x) - V_\tau(y)}{2} \Xi_\tau(x, y)(d\xi) \geq 0.$$

To simplify the notation, define for all $\tau \geq t$

$$\begin{aligned} u(\tau) &\equiv -\mathbb{E} \left[V_\tau(x) - W_\tau^\theta(x) \middle| \mathcal{F}_t \right] \\ \alpha(\tau) &\equiv -\mathbb{E} \left[\int_\tau^T \sum_{y=1}^N [S_r^{\theta^*}(x, y) m_r^{\theta^*}(x, y) - S_r^\theta(x, y) m_r^\theta(x, y)] \lambda_r(y|x) dr \middle| \mathcal{F}_t \right]. \end{aligned}$$

Then

$$u(t) = \rho \int_T^t u(\tau) d\tau + \alpha(t).$$

It remains to prove that $u(t) \leq 0$. This follows from Gröwall's lemma, since $\tau \mapsto \alpha(\tau)$ is non-positive and non-decreasing. Therefore

$$u(t) \leq \alpha(t) \exp \{ -\rho(T - t) \} \leq 0.$$

□

B Uniqueness and existence: proof of theorem 1

B.1 Delarue's theorem

We present here the system of FBSDEs studied by Delarue (2002), where we slightly adapt the problem to our purposes. Our choice of notation follows Delarue (2002). We then proceed and restate appropriate assumptions under which he proves the well-posedness of the system, namely existence and uniqueness of a solution.²⁶

²⁶Delarue (2002) considers a slightly richer problem where \tilde{X}_0 is random, square integrable and measurable with respect to some σ -algebra $\mathcal{G}_0 : \mathcal{F}_0 \subseteq \mathcal{G}_0$; the dimensions of the forward- and backward SDE need not match, coefficients of \tilde{f} and \tilde{g} may in addition depend on \tilde{Z} . This more general description of FBSDEs does not add any

Throughout, denote $\|\cdot\|$ the 1-norm, i.e., for any vector $v \in \mathbb{R}^N$ let $\|v\| = \frac{1}{N} \sum_{k=1}^N |v^k|$, where $|v^k|$ denotes the absolute value of the k th coordinate of v . Let

$$\begin{aligned}\tilde{f} &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N, \\ \tilde{g} &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N, \\ \tilde{\sigma} &: [0, T] \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N, \\ \tilde{h} &: \mathbb{R}^N \rightarrow \mathbb{R}^N\end{aligned}$$

measurable functions with respect to the Borelian σ -algebras.

For any \mathbb{R}^N -valued initial condition \tilde{X}_0 we are seeking $\mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^{N \times N}$ -valued and $\{\mathcal{F}_t\}$ -progressively measurable processes $(\tilde{X}_t, \tilde{Y}_t, \tilde{Z}_t)_{t \in [0, T]}$, solution of the of FBSDEs:

$$\left\{ \begin{array}{l} \forall t \in [0, T], \\ \tilde{X}_t = \tilde{X}_0 + \int_0^t \tilde{f}(s, \tilde{X}_s, \tilde{Y}_s) ds + \int_0^t \tilde{\sigma}(s, \tilde{X}_s, \tilde{Y}_s) dB_s \\ \tilde{Y}_t = \tilde{h}(\tilde{X}_T) + \int_t^T \tilde{g}(s, \tilde{X}_s, \tilde{Y}_s) ds - \int_t^T \tilde{Z}_s dB_s \\ \mathbb{E} \left[\int_0^T ((\tilde{X}_t)^2 + (\tilde{Y}_t)^2 + (\tilde{Z}_t)^2) dt \right] < \infty \end{array} \right. \quad (\text{E}')$$

We now present the (slightly adapted) regularity conditions from Delarue (2002). The first condition is Lipschitz continuity of \tilde{f} , \tilde{g} , \tilde{h} and $\tilde{\sigma}$ and corresponds to Delarue (2002)'s conditions (A1.1), (A1.2) and (A2.1). The second condition imposes a linear bound on those functions and corresponds to (A1.3) and (A2.2). The third condition imposes continuity of the drift and corresponds to (A1.4). Finally, the fourth condition imposes non-degeneracy of the variance of the diffusion process as well as continuity and corresponds to (A2.3) (non-degeneracy, after which Delarue (2002)'s article is named) and (A2.4).

Assumption A (Delarue (2002): assumptions (A1) and (A2)). *Functions \tilde{f} , \tilde{g} , \tilde{h} and $\tilde{\sigma}$ satisfy assumption A if there exist two non-negative constants \tilde{L} and $\tilde{\Lambda}$, as well as $\underline{\sigma} > 0$ such that:*

- $\forall t \in [0, T]$ and $(x, y), (x', y') \in \mathbb{R}^N \times \mathbb{R}^N$

$$\begin{aligned}\|\tilde{f}(t, x, y) - \tilde{f}(t, x, y')\| &\leq \tilde{L}\|y - y'\| \\ \|\tilde{g}(t, x, y) - \tilde{g}(t, x', y)\| &\leq \tilde{L}\|x - x'\| \\ \|\tilde{h}(x) - \tilde{h}(x')\| &\leq \tilde{L}\|x - x'\| \\ \|\tilde{\sigma}(t, x, y) - \tilde{\sigma}(t, x', y')\| &\leq \tilde{L}^2(\|x - x'\|^2 + \|y - y'\|^2)\end{aligned}$$

- $\forall t \in [0, T]$ and $(x, y), (x', y') \in \mathbb{R}^N \times \mathbb{R}^N$

$$\|\tilde{f}(t, x, y)\| \leq \tilde{\Lambda}(1 + \|y\|)$$

$$\|\tilde{g}(t, x, y)\| \leq \tilde{\Lambda}(1 + \|y\|)$$

$$\|\tilde{\sigma}(t, x, y)\| \leq \tilde{\Lambda}(1 + \|y\|)$$

$$\|\tilde{h}(x)\| \leq \tilde{\Lambda}$$

- $\forall t \in [0, T]$ and $(x, y) \in \mathbb{R}^N \times \mathbb{R}^N$ the functions $u \mapsto \tilde{f}(t, u, y)$ and $v \mapsto \tilde{g}(t, x, v)$ are continuous.

- $\forall t \in [0, T]$ and $(x, y) \in \mathbb{R}^N \times \mathbb{R}^N$ and $k \in \{1, \dots, N\}$

$$\tilde{\sigma}^k(t, x, y) \geq \underline{\sigma},$$

that is, the function $\tilde{\sigma}$ is uniformly bounded from below. Finally, $(t, x, y) \mapsto \tilde{\sigma}(t, x, y)$ is continuous.

We now proceed to state the first part of his main result:

Theorem A (Delarue (2002): theorem 2.6.1). *Problem (E') admits a unique solution under assumption A.*

At the heart of this result is the connection of FBSDEs with a quasi-linear parabolic system of PDEs. This had first been presented by Ma, Protter and Yong in their four step scheme. Delarue (2002) improves on this initial result in a number of ways, notably by relaxing the smoothness of parameters as seen in assumption A.

Drawing on Delarue's theorem A, we may then formulate the proof of theorem 1 which establishes the existence of a unique equilibrium under assumption 1.

B.2 Transfer of notation

Rewrite the system of FBSDEs (E) as

$$\begin{aligned} \gamma_t(x) &= \int_0^t \left[- \sum_{y \in Y} \beta_\tau e^{\gamma_\tau(y)} m_\tau(x, y) + \eta_\tau(x) - \frac{(\sigma_\tau(x))^2}{2} \right] d\tau + \int_0^t \sigma_\tau(x) dB_\tau \\ V_t(x) &= h_T + \int_t^T \alpha_\tau^X \sum_{y \in Y} \beta_\tau e^{\gamma_\tau(y)} m_\tau(x, y) [f(x, y) + \xi_\tau^e(x, y) - V_\tau(x) - V_\tau(y)] - \rho V_\tau(x) d\tau \\ &\quad - \int_t^T Z_\tau(x) \cdot dB_\tau, \end{aligned}$$

where

$$\xi_t^e(x, y) \equiv \frac{1}{m_t(x, y)} \int_{V_t(x)+V_t(y)-f(x, y)}^{\infty} \xi \Xi_t(d\xi)$$

$$m_t(x, y) \equiv 1 - \Xi_t[V_t(x) + V_t(y) - f(x, y)]$$

is the time t expected pair-specific production shock and meeting-contingent matching rate respectively between types x and y .

We then map our notation into Delarue's. Recall the type space is $X \cup Y = \{1, \dots, N^X, N^X + 1, \dots, N^X + N^Y \equiv N\}$. For any type $k \in X \cup Y$ denote $\chi(k)$ her own population: $\chi(k) = X$ if $k \in X$ and $\chi(k) = Y$ if $k \in Y$. Similarly, let $\chi^c(k) = X \cup Y \setminus \chi(k)$ the population of k 's potential partners.

Denote $\tilde{X}_t^k = \gamma_t(k)$ and $\tilde{Y}_t^k = V_t(k)$ for any $k \in \{1, \dots, N\}$. Finally, to abbreviate notation, denote $\Theta_t = (t, \tilde{X}_t, \tilde{Y}_t)$. Then identify

$$\tilde{f}^k(\Theta_t) \equiv - \sum_{\ell \in \chi(k)} \beta(\Theta_t) \exp \tilde{X}_t^\ell [1 - \Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))] + \eta(\Theta_t)(k)$$

$$\tilde{g}^k(\Theta_t) \equiv \alpha^{\chi(k)}(\Theta_t) \sum_{\ell \in \chi^c(k)} \beta(\Theta_t) \exp \tilde{X}_t^\ell [1 - \Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))] \cdot [f(k, \ell) + \xi_t^e(k, \ell) - \tilde{Y}_t^k - \tilde{Y}_t^\ell] - \rho \tilde{Y}_t^k \quad (\text{T})$$

$$\tilde{\sigma}^k(\Theta_t) \equiv \sigma(\Theta_t)(k)$$

$$\tilde{h}^k(\tilde{X}_T) \equiv h(\tilde{X}_T).$$

Under thus introduced notation the system E corresponds, as desired, to the system (E') studied by Delarue (2002).

B.3 Regularity conditions and proof

Lemma 6. *Suppose that the regularity condition 1 is in place. Then $\tilde{f}, \tilde{g}, \tilde{\sigma}, \tilde{h}$ as defined in (T) satisfy assumption A.*

Proof. We proceed in the same order as assumption A.

- Lipschitz continuity

Lipschitz continuity of \tilde{h} and $\tilde{\sigma}$ in (γ_t, V_t) is required by assumption 1. As to \tilde{f} , denote $\Theta = (t, \tilde{X}, \tilde{Y})$ and $\Theta'_t = (t, \tilde{X}'_t, \tilde{Y}'_t)$. Then

$$\begin{aligned}
& |\tilde{f}^k(\Theta_t) - \tilde{f}^k(\Theta'_t)| \\
&= \left| \sum_{\ell \in \chi(k)} \beta(\Theta_t) \exp \tilde{X}_t^\ell \{ \Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)) - \Xi(\Theta'_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)) \} + \eta(\Theta_t) - \eta(\Theta'_t) \right| \\
&\leq K^\beta \left\{ \sum_{\ell \in \chi(k)} |\Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)) - \Xi(\Theta'_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))| \right. \\
&\quad \left. + \sum_{\ell \in \chi(k)} |\Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)) - \Xi(\Theta'_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))| \right\} + |\eta(\Theta_t) - \eta(\Theta'_t)| \\
&\leq K^\beta \sum_{\ell \in \chi(k)} L^\xi |\tilde{Y}_t^k - \tilde{Y}_t^k + \tilde{Y}_t^\ell - \tilde{Y}_t^\ell| + K^\beta N L^\xi \|\tilde{Y}_t - \tilde{Y}'_t\| + L^\eta \|\tilde{Y}_t - \tilde{Y}'_t\| \\
&\leq K^\beta (\#\chi(k) L^\xi |\tilde{Y}_t^k - \tilde{Y}_t^k| + \sum_{\ell \in \chi(k)} L^\xi |\tilde{Y}_t^\ell - \tilde{Y}_t^\ell|) + (K^\beta N L^\xi + L^\eta) \|\tilde{Y}_t - \tilde{Y}'_t\| \\
&\leq K^\beta (N L^\xi \sum_{\ell=1}^N |\tilde{Y}_t^\ell - \tilde{Y}_t^\ell| + \sum_{\ell=1}^N L^\xi |\tilde{Y}_t^\ell - \tilde{Y}_t^\ell|) + (K^\beta N L^\xi + L^\eta) \|\tilde{Y}_t - \tilde{Y}'_t\| \\
&\leq (K^\beta (2N + 1) L^\xi + L^\eta) \|\tilde{Y}_t - \tilde{Y}'_t\|.
\end{aligned}$$

Therefore $\|\tilde{f}(\Theta_t) - \tilde{f}(\Theta'_t)\| \leq N(K^\beta(2N + 1)L^\xi + L^\eta) \|\tilde{Y}_t - \tilde{Y}'_t\|$ as desired.

As to \tilde{g} , denote $\Theta_t = (t, \tilde{X}_t, \tilde{Y}_t)$ and $\Theta'_t = (t, \tilde{X}'_t, \tilde{Y}_t)$. Likewise, denote $\xi_t^e(k, \ell) \equiv \xi^e(\Theta_t)(k, \ell)$ (the expected pair-specific production shock between types k and ℓ conditional on matching when the distribution function is $\Xi(\Theta_t)$ and the value of search is \tilde{V}_t). Then

$$\begin{aligned}
& \tilde{g}^k(\Theta_t) - \tilde{g}^k(\Theta'_t) \\
&= (\alpha^{\chi(k)}(\Theta_t) - \alpha^{\chi(k)}(\Theta'_t)) \sum_{\ell \in \chi^c(k)} \beta(\Theta_t) \exp \tilde{X}_t^\ell [1 - \Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))] \\
&\quad [f(k, \ell) + \xi_t^e(k, \ell) - \tilde{Y}_t^k - \tilde{Y}_t^\ell] \\
&\quad + \alpha^{\chi(k)}(\Theta'_t) \sum_{\ell \in \chi^c(k)} (\beta(\Theta_t) \exp \tilde{X}_t^\ell - \beta(\Theta'_t) \exp \tilde{X}'_t^\ell) [1 - \Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))] \\
&\quad [f(k, \ell) + \xi_t^e(k, \ell) - \tilde{Y}_t^k - \tilde{Y}_t^\ell] \\
&\quad + \alpha^{\chi(k)}(\Theta'_t) \sum_{\ell \in \chi^c(k)} \beta(\Theta'_t) \exp \tilde{X}'_t^\ell \\
&\quad \left\{ [\Xi(\Theta'_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)) - \Xi(\Theta_t)(\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell))] [f(k, \ell) - \tilde{Y}_t^k - \tilde{Y}_t^\ell] \right. \\
&\quad \left. + \int_{\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)}^{\tilde{\xi}} \xi \Xi(\Theta_t)(d\xi) - \int_{\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)}^{\tilde{\xi}} \xi \Xi(\Theta'_t)(d\xi) \right\}.
\end{aligned}$$

Now observe that $\tilde{X}_t \mapsto \beta(t, \tilde{X}_t, \tilde{Y}_t) \exp \tilde{X}_t^k$ is continuous and bounded, whence Lipschitz continuous. Denote associated Lipschitz constant L^β . Further observe that, due to Lipschitz continuity

of the density Ξ' ,

$$\begin{aligned} & \left| \int_{\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)}^{\bar{\xi}} \xi \Xi(\Theta_t)(d\xi) - \int_{\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)}^{\bar{\xi}} \xi \Xi(\Theta'_t)(d\xi) \right| \\ & \leq \int_{\tilde{Y}_t^k + \tilde{Y}_t^\ell - f(k, \ell)}^{\bar{\xi}} \xi \left| \Xi'(\Theta_t)(\xi) - \Xi'(\Theta'_t)(\xi) \right| d\xi \leq (\bar{\xi} - \underline{\xi}) L^\xi \|\tilde{X}_t - \tilde{X}'_t\|. \end{aligned}$$

Finally, for ease of notation, denote $\Pi = \max_{k, \ell} f(k, \ell) + \bar{\xi}$. It follows that

$$\begin{aligned} & \left| \tilde{g}^k(\Theta_t) - \tilde{g}^k(\Theta'_t) \right| \\ & \leq L^\alpha \|\tilde{X}_t - \tilde{X}'_t\| NK^\beta \Pi \\ & \quad + NL^\beta \|\tilde{X}_t - \tilde{X}'_t\| \Pi \\ & \quad + NK^\beta L^\xi \|\tilde{X}_t - \tilde{X}'_t\| \Pi \\ & \quad + NK^\beta (\bar{\xi} - \underline{\xi}) L^\xi \|\tilde{X}_t - \tilde{X}'_t\| \\ & = N \left(K^\beta (L^\alpha \Pi + L^\xi \Pi + (\bar{\xi} - \underline{\xi}) L^\xi) + L^\beta \Pi \right) \|\tilde{X}_t - \tilde{X}'_t\|, \end{aligned}$$

so that

$$\|\tilde{g}(\Theta_t) - \tilde{g}(\Theta'_t)\| \leq (N)^2 \left(K^\beta (L^\alpha \Pi + L^\xi \Pi + (\bar{\xi} - \underline{\xi}) L^\xi) + L^\beta \Pi \right) \|\tilde{X}_t - \tilde{X}'_t\|.$$

- Boundedness

\tilde{h} and $\tilde{\sigma}$ are bounded by assumption. Furthermore

$$\left| \tilde{f}^k(\Theta_t) \right| \leq NK^\beta + K^\eta \quad \text{and} \quad \left| \tilde{g}^k(\Theta_t) \right| \leq NK^\beta \Pi.$$

Multiplying by the number of types N gives the desired bounds of \tilde{f} and \tilde{g} .

- Continuity

$\tilde{X}_t \mapsto \tilde{f}(t, \tilde{X}_t, \tilde{Y}_t)$, being the sum and product of continuous functions (in \tilde{X}_t) is continuous, and so is $\tilde{Y}_t \mapsto \tilde{g}(t, \tilde{X}_t, \tilde{Y}_t)$.

- Non-degeneracy

Non-degeneracy of σ follows from the last item of assumption A. \square

The proof of theorem 1 follows immediately:

Proof. It suffices to prove that there exists a unique solution to the system of FBSDEs (E). Having verified in the preceding lemma that assumption A holds under stated regularity conditions 1, this follows from theorem A due to Delarue (2002). \square

C Single-peaked preferences

C.1 Proof of Lemma 3

Proof. Fix arbitrary types x_1, x_2 in X (a symmetric proof applies for types in Y) and time t . Write $\bar{P}_t(\tau)(x) \equiv \sum_{y \in Y} P_t(\tau)(y|x)$. Let θ such that $m_t^\theta(x_2, y) = m_\tau(x_1, y)$ for all $\tau \in [0, T]$, types $y \in Y$ and across states $\omega \in \Omega$ (so that x_2 matches for identical realizations of ξ as x_1). Denote for $k = 1, 2, \dots$

$$M_t^k(y|x_1) = \mathbb{E} \left[\int_{\tau_0=t}^T \int_{\tau_1}^T \cdots \int_{\tau_{k-1}}^T e^{-\rho(\tau_k-t)} \alpha_{\tau_k}^X P_{\tau_{k-1}}(d\tau_k)(y|x_1) \prod_{l=k-1}^1 (1 - \alpha_{\tau_l}^X) \bar{P}_{\tau_{l-1}}(d\tau_l)(x_1) \middle| \mathcal{F}_t \right]$$

$$R_t^k(x_1, x_2) = \mathbb{E} \left[\int_{\tau_0=t}^T \int_{\tau_1}^T \cdots \int_{\tau_{k-1}}^T (V_{\tau_k}(x_2) - V_{\tau_k}(x_1)) e^{-\rho(\tau_k-t)} \prod_{l=k}^1 (1 - \alpha_{\tau_l}^X) \bar{P}_{\tau_{l-1}}(d\tau_l)(x_1) \middle| \mathcal{F}_t \right].$$

We then prove by induction that

$$V_t(x_2) - V_t(x_1) \geq \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] \sum_{l=1}^k M_t^l(y|x_1) + R_t^k(x_1, x_2).$$

Indeed, the optimality condition $V_\tau(x_2) \geq W_\tau^\theta(x_2)$ gives

$$\begin{aligned} V_t(x_2) - V_t(x_1) &\geq W_t^\theta(x_2) - V_t(x_1) \\ &= \mathbb{E} \left[\int_t^T \sum_{y \in Y} [\pi_\tau^\theta(y|x_2) - \pi_\tau(y|x_1)] e^{-\rho(\tau-t)} P_t(d\tau)(y|x_1) \middle| \mathcal{F}_t \right] \\ &= \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] \underbrace{\mathbb{E} \left[\int_t^T e^{-\rho(\tau-t)} \alpha_\tau^X P_t(d\tau)(y|x_1) \middle| \mathcal{F}_t \right]}_{M_t^1(y|x_1)} \\ &\quad + \underbrace{\mathbb{E} \left[\int_t^T (V_\tau(x_2) - V_\tau(x_1)) e^{-\rho(\tau-t)} (1 - \alpha_\tau^X) \bar{P}_t(d\tau)(x_1) \middle| \mathcal{F}_t \right]}_{R_t^1(x_1, x_2)}. \end{aligned}$$

Now suppose by induction hypothesis that

$$V_t(x_2) - V_t(x_1) \geq \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] \sum_{l=1}^{k-1} M_t^l(y|x_1) + R_t^{k-1}(x_1, x_2).$$

We show that

$$R_t^{k-1}(x_1, x_2) \geq \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] M_t^k(y|x_1) + R_t^k(x_1, x_2)$$

from which the claim follows. Indeed, once more applying the optimality condition gives

$$\begin{aligned}
R_t^{k-1}(x_1, x_2) &\geq \mathbb{E} \left[\int_{\tau_0=t}^T \int_{\tau_1}^T \cdots \int_{\tau_{k-2}}^T (W_{\tau_{k-1}}^\theta(x_2) - V_{\tau_{k-1}}(x_1)) e^{-\rho(\tau_{k-1}-t)} \prod_{l=k-1}^1 (1 - \alpha_{\tau_l}^X) \bar{P}_{\tau_{l-1}}(d\tau_l)(x_1) \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[\int_{\tau_0=t}^T \int_{\tau_1}^T \cdots \int_{\tau_{k-2}}^T \left[\int_{\tau_{k-1}}^T \sum_{y \in Y} [\pi_{\tau_k}^\theta(y|x_2) - \pi_{\tau_k}(y|x_1)] e^{-\rho(\tau_k - \tau_{k-1})} P_{\tau_{k-1}}(d\tau_k)(y|x_1) \right] \right. \\
&\quad \left. e^{-\rho(\tau_{k-1}-t)} \prod_{l=k-1}^1 (1 - \alpha_{\tau_l}^X) \bar{P}_{\tau_{l-1}}(d\tau_l)(x_1) \middle| \mathcal{F}_t \right] \\
&= \sum_{y \in Y} [f(x_2, y) - f(x_1, y)] \\
&\quad \underbrace{\mathbb{E} \left[\int_{\tau_0=t}^T \int_{\tau_1}^T \cdots \int_{\tau_{k-1}}^T e^{-\rho(\tau_k-t)} \alpha_{\tau_k}^X P_{\tau_{k-1}}(d\tau_k)(y|x_1) \prod_{l=k-1}^1 (1 - \alpha_{\tau_l}^X) \bar{P}_{\tau_{l-1}}(d\tau_l)(x_1) \middle| \mathcal{F}_t \right]}_{M_t^k(y|x_1)} \\
&\quad + \underbrace{\mathbb{E} \left[\int_{\tau_0=t}^T \int_{\tau_1}^T \cdots \int_{\tau_{k-1}}^T (V_{\tau_k}(x_2) - V_{\tau_k}(x_1)) e^{-\rho(\tau_k-t)} \prod_{l=k}^1 (1 - \alpha_{\tau_l}^X) \bar{P}_{\tau_{l-1}}(d\tau_l)(x_1) \middle| \mathcal{F}_t \right]}_{R_t^k(x_1, x_2)}.
\end{aligned}$$

Then define

$$Q_t(y|x_1) = \lim_{k \rightarrow \infty} \sum_{l=1}^k M_t^l(y|x_1).$$

It remains to show that

$$\sum_{y \in Y} Q_t(y|x_1) \leq 1 \quad \text{and} \quad \lim_{k \rightarrow \infty} R_t^k(x_1, x_2) = 0 \quad \text{almost surely.}$$

We here prove this result when α_t^X is constant equal to α^X . A similarly straightforward result can be shown when the meeting rate appropriately bounded. The proof of the general result is deferred to a future version of this work. First note that

$$R_t^k(x_1, x_2) \leq (1 - \alpha^X)^k$$

Further observe that

$$\sum_{y \in Y} \int_t^T e^{-\rho(\tau-t)} \alpha^X P_t(d\tau)(y|x_1) \leq \alpha^X,$$

for the LHS describes the discounted expected probability times the bargaining weight. It follows that

$$\sum_{y \in Y} M_t^k(y|x_1) \leq \alpha^X (1 - \alpha_t^X)^{k-1},$$

whence

$$\sum_{y \in Y} Q_t(y|x_1) = \sum_{k \geq 1} \sum_{y \in Y} M_t^k(y|x_1) \leq \sum_{k \geq 1} \alpha^X (1 - \alpha^X)^{k-1} = 1.$$

□

C.2 Proof of theorem 2

We here prove convexity of matching sets. This proof relies on two intermediate observations. First, we need to establish quasi-concavity of output f . Clearly, in thin markets the value of search may be very small. Then for ex-ante surplus $S_t(x, y)$ to be quasi-concave, so must be output f :

Lemma 7. *If f is supermodular and log supermodular in differences with respect to y , then $f(x, y)$ is quasi-concave with respect to y for any $x > x_{min}$.*

Proof. Suppose that there exist $x_1 > x_0$ and $y_3 > y_2 > y_1$ such that $f(x_1, y_1) > f(x_1, y_2)$ and $f(x_1, y_3) > f(x_1, y_2)$. Supermodularity ensures that $f(x_1, y_3) - f(x_1, y_2) > f(x_0, y_3) - f(x_0, y_2)$. Multiplying on both side by $\frac{1}{f(x_1, y_2) - f(x_1, y_1)} < 0$, we obtain:

$$\frac{f(x_0, y_3) - f(x_0, y_2)}{f(x_1, y_2) - f(x_1, y_1)} > \frac{f(x_1, y_3) - f(x_1, y_2)}{f(x_1, y_2) - f(x_1, y_1)}.$$

We can construct an lower bound of the right hand side using log supermodularity in differences:

$$\frac{f(x_0, y_3) - f(x_0, y_2)}{f(x_1, y_2) - f(x_1, y_1)} > \frac{f(x_1, y_3) - f(x_1, y_2)}{f(x_1, y_2) - f(x_1, y_1)} \geq \frac{f(x_0, y_3) - f(x_0, y_2)}{f(x_0, y_2) - f(x_0, y_1)}.$$

Case 1: if $f(x_0, y_3) - f(x_0, y_2) \geq 0$ then the above inequalities imply that:

$$\frac{1}{f(x_1, y_2) - f(x_1, y_1)} > \frac{1}{f(x_0, y_2) - f(x_0, y_1)}.$$

If $f(x_0, y_3) - f(x_0, y_2) \geq 0$ then log supermodularity in differences imply that $f(x_0, y_2) - f(x_0, y_1) \leq 0$, and so the above inequality contradicts supermodularity: $f(x_1, y_2) - f(x_1, y_1) > f(x_0, y_2) - f(x_0, y_1)$.

Case 2: if $f(x_0, y_3) - f(x_0, y_2) \leq 0$ then,

$$\frac{1}{f(x_1, y_2) - f(x_1, y_1)} < \frac{1}{f(x_0, y_2) - f(x_0, y_1)}.$$

If $f(x_0, y_3) - f(x_0, y_2) \leq 0$ then log supermodularity in differences imply that $f(x_0, y_2) - f(x_0, y_1) \geq 0$ and so the above inequality contradicts supermodularity: $f(x_1, y_2) - f(x_1, y_1) > f(x_0, y_2) - f(x_0, y_1)$.

□

Secondly, log supermodularity in differences plays an important role in the proof, because it relates to preferences over risk-preferences. More specifically, Pratt (1964) shows that *the following statements are equivalent for any two agent types $x_1 < x_2$ with respective utility functions $y \mapsto h(x_i, y)$, $x_i \in \{x_1, x_2\}$:*

1. $h(x, y)$ is log supermodular in differences with respect to y ;
2. Agent type x_1 is more risk-averse than agent type x_2 ; that is, x_1 does not accept a lottery that is rejected by x_2 .

We then proceed to the proof of the convexity result:

Proof. We provide a proof by contraposition. Suppose there exists a type $x_1 \in X$ whose surplus function $y \mapsto S_t(x_1, y)$ at a given time t and state ω was not strictly quasi-concave. Then, there exist types $y_3 > y_2 > y_1$ such that

$$S_t^\omega(x_1, y_3) \geq S_t^\omega(x_1, y_2) \quad \text{and} \quad S_t^\omega(x_1, y_1) \geq S_t^\omega(x_1, y_2)$$

which is equivalent to

$$V_t^\omega(y_3) - V_t^\omega(y_2) \leq f(x_1, y_3) - f(x_1, y_2) \quad \text{and} \quad V_t^\omega(y_2) - V_t^\omega(y_1) \geq f(x_1, y_2) - f(x_1, y_1).$$

Due to lemma 3, there exist ω -contingent non-negative weights $Q_t^\omega(\cdot|y_2) : X \rightarrow \mathbb{R}_+$ with $Q_t^\omega(y_2) \equiv \sum_{x \in X} Q_t^\omega(x|y_2) \leq 1$ satisfying

$$\begin{aligned} \sum_{x \in X} [f(x, y_2) - f(x, y_1)] Q_t^\omega(x|y_2) &\geq f(x_1, y_2) - f(x_1, y_1) \\ \sum_{x \in X} [f(x, y_3) - f(x, y_2)] Q_t^\omega(x|y_2) &\leq f(x_1, y_3) - f(x_1, y_2). \end{aligned}$$

Now distinguish between two cases. Suppose (i) that $f(x_1, y_2) - f(x_1, y_1) \geq 0$. Then let x_2 the largest type in $[0, 1]$ such that

$$\frac{1}{Q_t^\omega(y_2)} (f(x_1, y_2) - f(x_1, y_1)) \geq f(x_2, y_2) - f(x_2, y_1).$$

Due to supermodularity, the fact that $Q_t^\omega(y_2) \in (0, 1)$, and $f(x_1, y_2) - f(x_1, y_1) \geq 0$ we have that $x_2 \geq x_1$. Moreover, unless $x_2 = 1$, due to the intermediate value theorem the preceding holds as an equality. Now, if $x_2 = 1$, supermodularity alone implies that

$$\sum_{x \in X} [f(x, y_3) - f(x, y_2)] \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} \leq f(1, y_3) - f(1, y_2).$$

Otherwise, log supermodularity in differences of f with respect to y gives

$$\begin{aligned} \frac{1}{Q_t^\omega(y_2)}(f(x_1, y_3) - f(x_1, y_2)) &= \frac{f(x_1, y_3) - f(x_1, y_2)}{f(x_2, y_3) - f(x_2, y_2)} \frac{1}{Q_t^\omega(y_2)}(f(x_2, y_3) - f(x_2, y_2)) \\ &\leq \frac{f(x_1, y_2) - f(x_1, y_1)}{f(x_2, y_2) - f(x_2, y_1)} \frac{1}{Q_t^\omega(y_2)}(f(x_2, y_3) - f(x_2, y_2)) = f(x_2, y_3) - f(x_2, y_2). \end{aligned}$$

And so, there exists $\hat{x} = x_2$, such that

$$\begin{aligned} \sum_{x \in X} [f(x, y_2) - f(x, y_1)] \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\geq f(\hat{x}, y_2) - f(\hat{x}, y_1) \\ \sum_{x \in X} [f(x, y_3) - f(x, y_2)] \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\leq f(\hat{x}, y_3) - f(\hat{x}, y_2). \end{aligned}$$

Now suppose (ii) that $f(x_1, y_2) - f(x_1, y_1) < 0$. Since f is quasi-concave, also $f(x_1, y_3) - f(x_1, y_2) < 0$. We apply a symmetric reasoning as in (i): let x_0 the smallest type in $[0, 1]$ such that

$$\frac{1}{Q_t^\omega(y_2)}(f(x_1, y_3) - f(x_1, y_2)) \leq f(x_0, y_3) - f(x_0, y_2).$$

Due to supermodularity, the fact that $Q_t^\omega(y_2) \in (0, 1)$, and $f(x_1, y_3) - f(x_1, y_2) < 0$ we have that $x_0 \leq x_1$. Moreover, unless $x_0 = 0$, due to the intermediate value theorem the preceding holds as an equality. Now, if $x_0 = 0$, supermodularity alone implies that

$$\sum_{x \in X} [f(x, y_2) - f(x, y_1)] \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} \geq f(0, y_2) - f(0, y_1).$$

Otherwise, log supermodularity in differences of f with respect to y and the fact that $f(x_0, y_2) - f(x_0, y_1) < 0$ give

$$\begin{aligned} \frac{1}{Q_t^\omega(y_2)}(f(x_1, y_2) - f(x_1, y_1)) &= \frac{f(x_1, y_2) - f(x_1, y_1)}{f(x_0, y_2) - f(x_0, y_1)} \frac{1}{Q_t^\omega(y_2)}(f(x_0, y_2) - f(x_0, y_1)) \\ &\geq \frac{f(x_1, y_3) - f(x_1, y_2)}{f(x_0, y_3) - f(x_0, y_2)} \frac{1}{Q_t^\omega(y_2)}(f(x_0, y_2) - f(x_0, y_1)) = f(x_0, y_2) - f(x_0, y_1). \end{aligned}$$

And so, as in (i), there exists $\hat{x} = x_0$, such that

$$\begin{aligned} \sum_{x \in X} [f(x, y_2) - f(x, y_1)] \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\geq f(\hat{x}, y_2) - f(\hat{x}, y_1) \\ \sum_{x \in X} [f(x, y_3) - f(x, y_2)] \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\leq f(\hat{x}, y_3) - f(\hat{x}, y_2). \end{aligned}$$

This pair of inequalities is of course equivalent to summation over differences:

$$\begin{aligned} \sum_{x \in X} \sum_{\substack{y \in Y \\ y_1 \leq y < y_2}} \Delta_y f(x, y) \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\geq \sum_{\substack{y \in Y \\ y_1 \leq y < y_2}} \Delta_y f(\hat{x}, y) \\ \sum_{x \in X} \sum_{\substack{y \in Y \\ y_2 \leq y < y_3}} \Delta_y f(x, y) \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\leq \sum_{\substack{y \in Y \\ y_2 \leq y < y_3}} \Delta_y f(\hat{x}, y). \end{aligned}$$

In particular, this means there exist $y' : y_1 \leq y' < y_2$ and $y'' : y_2 \leq y'' < y_3$ such that

$$\begin{aligned} \sum_{x \in X} \Delta_y f(x, y') \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\geq \Delta_y f(\hat{x}, y') \\ \sum_{x \in X} \Delta_y f(x, y'') \frac{Q_t^\omega(x|y_2)}{Q_t^\omega(y_2)} &\leq \Delta_y f(\hat{x}, y''). \end{aligned}$$

This runs counter to the characterization of log supermodularity in differences of $\Delta_y f$ in terms of risk preferences; interpret $x \mapsto \Delta_y f(\cdot, y)$ as agent type y 's utility function. Then agent type y'' prefers the safe outcome \hat{x} over a lottery, whereas agent type y' , of lower rank than y'' , prefers the lottery when facing the identical decision problem, in spite of Pratt's theorem. \square

C.3 Observation: non single-peaked wages

We here make a remark that is interesting in its own right: provided that preferences are single-peaked, we show here that expected equilibrium payoffs, denoted $\pi_t^e(y|x)$, are not necessarily single-peaked. First, expected payoffs can be expressed as follows:

$$\pi_t^e(y|x) \equiv V_t(x) + \alpha_t^X S_t(x, y) + \frac{\alpha_t^X}{1 - \Xi_t(-S_t(x, y))} \int_{-S_t(x, y)}^{\bar{\xi}} \xi \Xi_t(d\xi).$$

We are interested in the shape of this function with respect to y . As $V_t(x)$ does not depend on y , it suffices to consider the second and third term. If those terms were a monotone transformation of S , π_t^e would inherit the properties of S and be single-peaked. The following reveals that absent further assumptions regarding Ξ this need not be the case. Consider

$$S \mapsto \alpha_t^X S + \frac{\alpha_t^X}{1 - \Xi_t(-S)} \int_{-S}^{\bar{\xi}} \xi \Xi_t(d\xi).$$

Taking the derivative with respect to S (using the Leibniz integral rule) we obtain

$$1 - \frac{\Xi_t'(-S)}{(1 - \Xi_t(-S))^2} \int_{-S}^{\bar{\xi}} \xi \Xi_t(d\xi) + \frac{\Xi_t'(-S)S}{1 - \Xi_t(-S)},$$

which boils down to

$$FOC : \quad 1 - \frac{\Xi'_t(-S)}{1 - \Xi_t(-S)} \mathbb{E}^\xi[\xi - S \mid \xi + S].$$

Observe that this expression is not necessarily non-negative. It follows that, unlike meeting-contingent match probabilities and surplus, match payoffs conditional on matching need not be single-peaked.

D Assortative matching: proofs

D.1 Illustrative results

Non-emptiness + lattice \implies convexity

Proposition 4. *Suppose that $U_t(y; p)$ is non-empty for all $y \in Y$. Then if $U_t(p)$ is a lattice, then $U_t(x; p)$ is convex for all $x \in X$.*

The proof is an immediate adaptation from Shimer and Smith (2000).

Proof. Suppose that there exist x and $y_3 > y_2 > y_1$ such that y_3 and y_1 are in $U_t(x; p)$, but y_2 is not. Or, equivalently, $(x, y_1), (x, y_3) \in U_t(p)$, but $(x, y_2) \notin U_t(p)$. Since $U_t(y_2; p)$ is non empty, there exists a type $x' \in U_t(y_2; p)$. If $x' > x$, the fact that $(y_3, x), (y_2, x') \in U_t(p)$ and the lattice property imply that $(x, y_2) \in U_t(p)$. Using the same reasoning, we find a contradiction for the case $x' < x$. \square

Proof of proposition 5

Proposition 5. *Suppose that $U_t(x; p)$ is convex for all agent types. Then $U_t(p)$ is a lattice if and only if $x \mapsto l_t(x; p)$ and $x \mapsto u_t(x; p)$ are non-decreasing.*

Proof. Maintain throughout that $U_t(x; p)$ is convex. (\Leftarrow). Suppose that $x \mapsto l_t(x; p)$ as well as $x \mapsto u_t(x; p)$ are non-decreasing. Consider $x_1 < x_2$ and $y_1 < y_2$ such that (x_1, y_2) and (x_2, y_1) in $U_t(p)$. Then $(x_1, y_2) \in U_t(p)$ implies that $y_2 \leq u_t(x_1; p)$. And due to $x \mapsto u_t(x; p)$ being non-decreasing, $y_2 \leq u_t(x_2; p)$. Likewise, $(x_2, y_1) \in U_t(p)$ implies that $y_1 \geq l_t(x_2; p)$, and due to $x \mapsto l_t(x; p)$ being non-decreasing, $y_2 \geq l_t(x_2; p)$. As a result, $l_t(x_2; p) \leq y_2 \leq u_t(x_2; p)$. Since $U_t(x_2; p)$ is convex, this implies that y_2 in $U_t(x_2; p)$, or that $(x_2, y_2) \in U_t(p)$. Using an identical reasoning, it follows that $(x_1, y_1) \in U_t(p)$.

(\Rightarrow). Let $U_t(p)$ a lattice. Consider arbitrary $x_1 < x_2$. We show that $u_t(x_1; p) \leq u_t(x_2; p)$. If not, for $(x_1, u_t(x_1, p)) \in U_t(p)$, $(x_2, u_t(x_2, p)) \in U_t(p)$, the lattice property implies that $(x_2, u_t(x_1, p)) \in U_t(p)$, or that $u_t(x_1, p) \in \{y : m_t(x_2, y) \geq p\}$. It follows that $u_t(x_2; p) \equiv \max\{y : m_t(x_2, y) \geq p\} \geq u_t(x_1, p)$ as was to be shown. An identical reasoning establishes the monotonicity of $l_t(x; p)$. \square

D.2 Proof of the preference representation

We here provide three results that make a connection between preferences and PAM, and jointly establish an equivalent representation of PAM, namely proposition 2.

Proposition 6. *Suppose that preferences \succsim_t^x and \succsim_t^y satisfy weak single-crossing, and are weakly single-peaked and reciprocal. Then upper contour matching sets $U_t(p)$ are a lattice for all p .*

Proof. Pick any four types $x_1 < x_2$ and $y_1 < y_2$. We show that

$$\min\{m_t(x_1, y_1), m_t(x_2, y_2)\} \geq \min\{m_t(x_2, y_1), m_t(x_1, y_2)\},$$

thus proving the lattice property. Distinguish between two cases:

case (i) $m_t(x_1, y_2) > m_t(x_1, y_1)$. Weak single-crossing implies that $m_t(x_2, y_2) \geq m_t(x_2, y_1)$. Then for the preceding min-min condition to hold it suffices that $m_t(x_1, y_1) \geq m_t(x_2, y_1)$. Suppose by contradiction that this held with strictly reversed sign. We are thus left with two inequalities. The first inequality, $m_t(x_1, y_2) > m_t(x_1, y_1)$, taken together with weakly single-peaked preferences, implies that the set of agent type x_1 preferred partners' types strictly majorizes agent type y_1 , therefore $y_1 < \underline{y}_t(x_1)$. The second inequality, $m_t(x_1, y_1) < m_t(x_2, y_1)$, taken together with weakly single-peaked preferences, implies that the set of agent type y_1 preferred partners' types strictly majorizes x_1 , therefore $x_1 < \underline{x}_t(y_1)$. As both conditions cannot obtain jointly under reciprocity, this poses the desired contradiction in case (i).

Case (ii) $m_t(x_1, y_2) \leq m_t(x_1, y_1)$. Then for the preceding min-min condition to hold it suffices that $m_t(x_2, y_2) \geq \min\{m_t(x_1, y_2); m_t(x_2, y_1)\}$. Suppose by contradiction that this held with strictly reversed sign: then $m_t(x_2, y_2) < m_t(x_1, y_2)$ and $m_t(x_2, y_2) < m_t(x_2, y_1)$. The first inequality, taken together with weakly single-peaked preferences, implies that the set of agent type y_2 preferred partners' types is inferior to agent type x_2 , therefore $x_2 > \bar{x}_t(y_1)$. The second inequality, taken together with weakly single-peaked preferences, implies that the set of agent type x_2 preferred partners' types is inferior to agent type y_2 , therefore $y_2 > \bar{y}_t(x_2)$. As both conditions cannot obtain jointly under reciprocity, this poses the desired contradiction in case (ii). \square

Proposition 7. *Suppose that upper contour matching sets $U_t(p)$ are a lattice for all p . Then the preference relation \succsim_t^x satisfies weak single-crossing.*

Proof. Suppose by contradiction that $m_t(x, y)$ did not satisfy weak single-crossing, i.e., suppose there exist state ω , time t and four types $x_2 > x_1$ yet $y_2 > y_1$ such that $m_t^\omega(x_1, y_2) > m_t^\omega(x_1, y_1)$ and $m_t^\omega(x_2, y_2) < m_t^\omega(x_2, y_1)$. $U_t(p)$ being a lattice for all p requires that $\min\{m_t^\omega(x_1, y_1), m_t^\omega(x_2, y_2)\} \geq \min\{m_t^\omega(x_1, y_2), m_t^\omega(x_2, y_1)\}$. For convenience, denote $a \equiv m_t^\omega(x_1, y_1)$, $b \equiv m_t^\omega(x_2, y_2)$, $c \equiv m_t^\omega(x_1, y_2)$, $d \equiv m_t^\omega(x_2, y_1)$. Then the problem is to contradict the following set of inequalities:

$$c > a, d > b \quad \text{and} \quad \min\{a, b\} \geq \min\{c, d\}.$$

We distinguish between five cases (throughout the outer terms of the inequalities will pose a contradiction to one of the preceding three inequalities): (i) if $a = b$, $c = d$, then $\min\{a, b\} = a < c = \min\{c, d\}$. Absurd. (ii) if $a \geq b$, $c \geq d$, then $a \geq b = \min\{a, b\} \geq \min\{c, d\} = d$.

Absurd. (iii) if $a \geq b$, $c \leq d$, then $a \geq b = \min\{a, b\} \geq \min\{c, d\} = c$. Absurd. (iv) if $a < b$, $c \geq d$, then $b > a = \min\{a, b\} \geq \min\{c, d\} = d$. Absurd. (v) if $a \leq b$, $c \leq d$, then $a = \min\{a, b\} \geq \min\{c, d\} \geq c$. Absurd. \square

Proposition 8. *Suppose that upper contour sets $U_t(p)$ are a lattice for all p . Then the preference relation \succeq_t^x satisfies reciprocity.*

Proof. There are two possible contrapositions to reciprocity: (i) there exist x_1, y_1 such that $y_1 < \underline{y}_t(x_1) \equiv y_2$, yet $x_1 < \underline{x}_t(y_1) \equiv x_2$. Or, this implies that $m_t(x_1, y_1) < m_t(x_1, y_2)$ and $m_t(x_1, y_1) < m_t(x_2, y_1)$. So $U_t(\min\{m_t(x_1, y_2); m_t(x_2, y_1)\})$ is not a lattice. (ii) there exist x_2, y_2 such that $y_2 > \bar{y}_t(x_2) \equiv y_1$, yet $x_2 > \bar{x}_t(y_2) \equiv x_1$. Or, this implies that, $m_t(x_2, y_2) < m_t(x_2, y_1)$ and $m_t(x_2, y_2) < m_t(x_1, y_2)$. So $U_t(\min\{m_t(x_1, y_2); m_t(x_2, y_1)\})$ is not a lattice. \square

D.3 Characterization of reciprocity via \mathcal{P}_t

We now provide a proof of proposition 3. We prove each claim in turn.

Claim 1: $\mathcal{P}_t(y)$ is convex.

Proof. Fix $y_1 < y_2 < y_3$ such that $(x, y_1), (x, y_3) \in \Gamma_t$. We show that then also $(x, y_2) \in \mathcal{P}_t$. Due to convexity of $Y_t(x)$ as implied by weakly single-peaked preferences, this is readily the case when both $y_1, y_3 \in Y_t(x)$. Three cases remain.

(i) $y_1 \in Y_t(x)$ and $x \in X_t(y_3)$. If $y_2 \notin Y_t(x)$, then both $m_t(x, y_1) > m_t(x, y_2)$ and $m_t(x, y_3) > m_t(x, y_2)$. Suppose to the contrary that $(x, y_2) \notin \mathcal{P}_t$. Then both, $y_2 \notin Y_t(x)$ so that $m_t(x, y_1) > m_t(x, y_2)$ and $m_t(x, y_2) > m_t(x, y_3)$, and $x \notin X_t(y_2)$ so that either $x < X_t(y_2)$ or $x > X_t(y_2)$. In the former case $\underline{x}_t(y_2) > x$ and $m_t(x, y_2) < m_t(\underline{x}_t(y_2), y_2)$. Denoting $p \equiv \min\{m_t(x, y_3), m_t(\underline{x}_t(y_2), y_2)\}$, we have that $(x, y_3), (\underline{x}_t(y_2), y_2)$ belong to $U_t(p)$, but not (x, y_2) . Then $U_t(p)$ is not a lattice. In the latter case $\bar{x}_t(y_2) < x$ and $m_t(x, y_2) < m_t(\bar{x}_t(y_2), y_2)$. Denoting $p \equiv \min\{m_t(x, y_1), m_t(\bar{x}_t(y_2), y_2)\}$, we have that $(x, y_1), (\bar{x}_t(y_2), y_2)$ belong to $U_t(p)$, but not (x, y_2) . Once more, $U_t(p)$ is not a lattice. Having established, provided preferences are weakly single-peaked and satisfy weak single-crossing and reciprocity, that $U_t(p)$ is a lattice for all p , this poses the desired contradiction.

(ii) $y_3 \in Y_t(x)$ and $x \in X_t(y_1)$. Symmetric arguments as in (i) apply.

(iii) $x \in X_t(y_1)$ and $x \in X_t(y_3)$. If $y_2 \in Y_t(x)$, convexity is satisfied. If $Y_t(x) < y_2$, then for $y'_1 = \bar{y}_t(x)$, $y'_1 \in Y_t(x)$ and $x \in X_t(y_3)$, so that (i) establishes that $(x, y_2) \in \mathcal{P}_t$. If $Y_t(x) > y_2$, then for $y'_3 = \underline{y}_t(x)$, $y'_3 \in Y_t(x)$ and $x \in X_t(y_1)$, so that (ii) establishes that $(x, y_2) \in \mathcal{P}_t$. \square

Claim 2: \mathcal{P}_t is connected (this means that between any two pairs of types in \mathcal{P}_t , one can stay within \mathcal{P}_t by moving along the horizontal, vertical axis or diagonal and connect the two pairs of types).

Proof. Consider two adjacent types $x_1 < x_2$. We show that $Y_t(x_1)$ and $Y_t(x_2)$ are connected. Since both $Y_t(x_1)$ and $Y_t(x_2)$ are convex due to weak single-crossing, it remains to consider two cases:

(i) Suppose that $\bar{y}_t(x_1) < \underline{y}_t(x_2)$. Then pick $y : \bar{y}_t(x_1) < y < \underline{y}_t(x_2)$. Then reciprocity implies that $\underline{x}_t(y) \leq x_2$ and $\bar{x}_t(y) \geq x_1$, so that $\{x_1, x_2\} \in X_t(y)$ for all such y . In particular, it follows one can connect $(x_1, \bar{y}_t(x_1))$ and $(x_2, \underline{y}_t(x_2))$ via a path in \mathcal{P}_t .

(ii) Suppose that $\underline{y}_t(x_1) > \bar{y}_t(x_2)$. Then there exists $y : \bar{y}_t(x_2) < y < \underline{y}_t(x_1)$. But this is impossible according to single-crossing: if type x_1 strictly prefers $\underline{y}_t(x_1)$ over $\bar{y}_t(x_2)$, then type x_2 must at weakly share this and weakly prefer $\underline{y}_t(x_1)$ over $\bar{y}_t(x_2)$. But then $\underline{y}_t(x_1) \in Y_t(x_1)$ which poses the desired contradiction. \square

Claim 3: \mathcal{P}_t contains both (x_{\max}, y_{\max}) and (x_{\min}, y_{\min})

Proof. We show that (x_{\max}, y_{\max}) is in \mathcal{P}_t . That (x_{\min}, y_{\min}) is in \mathcal{P}_t then follows along identical lines. Indeed, if $\bar{x}_t(y_{\max}) < x_{\max}$, reciprocity requires that $\bar{y}_t(x_{\max}) \geq y_{\max}$, so that $\bar{y}_t(x_{\max}) = y_{\max}$. \square

D.4 Proof of theorem 4, non-stationary Shimer and Smith

We first prove lemma 4.

Proof. We proof the claim for x_{\max} , identical arguments apply for x_{\min} . Suppose to the contrary. Then there exists time t , state ω and type $y_j \in Y$ such that $f(x_{\max}, y_{\max}) - V_t^\omega(x_{\max}) - V_t^\omega(y_{\max}) < f(x_{\max}, y_j) - V_t^\omega(x_{\max}) - V_t^\omega(y_j)$. Or, applying the mimicking argument 3, this implies that there exists a measure $\mathcal{Q}_t^\omega(x|y_{\max})$ whose sum is less than one such that

$$f(x_{\max}, y_{\max}) - f(x_{\max}, y_j) < V_t^\omega(y_{\max}) - V_t^\omega(y_j) \leq \sum_{x \in X} [f(x, y_{\max}) - f(x, y_j)] \mathcal{Q}_t^\omega(x|y_{\max}).$$

Then supermodularity implies the first, the boundary condition the second inequality in what follows:

$$\begin{aligned} \sum_{x \in X} [f(x, y_{\max}) - f(x, y_j)] \mathcal{Q}_t^\omega(x|y_{\max}) &\leq [f(x_{\max}, y_{\max}) - f(x_{\max}, y_j)] \sum_{x \in X} \mathcal{Q}_t^\omega(x|y_{\max}) \\ &\leq f(x_{\max}, y_{\max}) - f(x_{\max}, y_j). \end{aligned}$$

This poses the desired contradiction. \square

We next prove that $S_t(x, y)$ inherits Lipschitz-continuity from $f(x, y)$:

Lemma 8. *Suppose that output f is Lipschitz-continuous with constant L^f . Then, for any state and time, surplus S_t is Lipschitz-continuous with constant $3L^f$.*

Proof. Observe that

$$\begin{aligned} |S_t(x', y') - S_t(x'', y'')| &\leq |f(x', y') - f(x'', y'')| + |V_t(x') - V_t(x'')| + |V_t(y') - V_t(y'')| \\ &\leq |f(x', y') - f(x'', y'')| + \left| \sum_y [f(x', y) - f(x'', y)] \mathcal{Q}(y|x') \right| + \left| \sum_x [f(x, y') - f(x, y'')] \mathcal{Q}(x|y') \right| \end{aligned}$$

due to the mimicking argument 3. \square

We then provide a proof of the key lemma 5.

Proof. Pick arbitrary $(x_3, y_3) \in \bar{\mathcal{E}}_t$. Four cases are of interest. We show that in the first three cases $m_t(x_3, y_3) \geq \hat{p}_t$. First, if $(x_3, y_3) \in \mathcal{P}_t$, this is readily the case. Secondly, if $(x_3, y_3) \notin \mathcal{P}_t$, yet there exists y_0 such that $x_3 \in x_t(y_0)$, then monotonicity of the correspondence $y \mapsto x_t(y)$ due to single-crossing and $x_3 < \underline{x}_t(y_3)$ due to $(x_3, y_3) \in \bar{\mathcal{E}}_t$ imply that $y_0 < y_3$. Then weakly single-peaked preferences give that $m_t(x_3, \underline{y}_t(x_3)) \geq m_t(x_3, y_3) \geq m_t(x_3, y_0) \geq \hat{p}_t$ as desired. Thirdly, if $(x_3, y_3) \notin \mathcal{P}_t$, yet there exists x_0 such that $y_3 \in y_t(x_0)$, a symmetric reasoning applies.

Thus consider the fourth case: suppose that $(x_3, y_3) \notin \mathcal{P}_t$ and there do not exist x_0, y_0 such that $x_3 \in x_t(y_0)$ or $y_3 \in y_t(x_0)$. Consider, if well-defined, $y_1 \equiv \max\{y : \bar{x}_t(y_1) < x_3\}$ and $x_1 \equiv \max\{x : \bar{y}_t(x_1) < y_3\}$.

Since $(x_{\min}, y_{\min}) \in \mathcal{P}_t$, either $x_{\min} = \underline{x}_t(y_{\min})$ or $y_{\min} = \underline{y}_t(x_{\min})$. Since there do not exist x_0, y_0 with the desired properties, neither $y_3 \in y_t(x_{\min})$, nor $x_3 \in x_t(y_{\min})$. It follows that either $\{y : \bar{x}_t(y_1) < x_3\}$ or $\{x : \bar{y}_t(x_1) < y_3\}$ is non-empty (for y_{\min} or x_{\min} respectively belong to it). Thus x_1 or y_1 (but not necessarily both) is well-defined. Thus, without loss of generality, take x_1 to be well-defined

Then define $x_2 \equiv \min\{x : \underline{y}_t(x_2) > y_3\}$. Since $(x_3, y_3) \in \bar{\mathcal{E}}_t \setminus \mathcal{P}_t$, $y_3 < \underline{y}_t(x_3)$. It follows that $\{x : \underline{y}_t(x_2) > y_3\}$ is non-empty (for x_3 belongs to it), so that x_2 is well-defined. By construction, $x_1 < x_2 \leq x_3$ and x_1, x_2 are adjacent.

Next, notice that $x_3 < \underline{x}_t(y_3)$ in conjunction with weak single-peaked preferences implies that $m_t(x_3, y_3) \geq m_t(x_2, y_3) \geq m_t(x_2, y_1)$. Notice further that $m_t(x_1, y_1) \geq \hat{p}_t$. It thus suffices to prove that $m_t(x_2, y_1) \geq m_t(x_1, y_1) - \epsilon^N$.

This follows from the fact that $f(x, y)$ is Lipschitz-continuous. For then, according to the preceding lemma $S_t(x, y)$, is Lipschitz-continuous with constant $\frac{q}{r} \frac{3}{N}$. And therefore, since x_1 and x_2 are adjacent, $|S_t(x_2, y_1) - S_t(x_1, y_1)| < 3 \frac{q}{N}$. Therefore, $m_t(x_2, y_1) = 1 - \Xi_t(S_t(x_2, y_1)) \geq 1 - \Xi_t(S_t(x_1, y_1) + 3 \frac{q}{N}) = m_t(x_1, y_1) + \Xi_t(S_t(x_1, y_1)) - \Xi_t(S_t(x_1, y_1) + 3 \frac{q}{N})$. Since Ξ_t is Lipschitz-continuous with constant $\frac{L^\xi}{3q}$, it follows that $\Xi_t(S + 3 \frac{q}{N}) - \Xi_t(S) \leq \frac{L^\xi}{N}$ for all S . Then the preceding implies that $m_t(x_3, y_3) \geq m_t(x_2, y_1) \geq \hat{p}_t - \frac{L^\xi}{N}$, as desired. □

D.5 Self-acceptance at the boundary: steady state

We here provide a result which is interesting in its own right. Unlike the main text, we consider here an identical environment as in Shimer and Smith (2000). The economy is stationary and there are no pair-specific production shocks. We show that Shimer and Smith's boundary condition can be relaxed in the steady state: rather than imposing that output for lowest types is non-increasing, we add the complementarity condition that output is log supermodular.

Proposition 9. *Suppose that the economy is at the steady state, and populations are symmetric. If f is supermodular and log supermodular, then the lowest type self-accepts.*

To achieve this we construct an upper bound on the stationary value of search which does not depend on other agent types' value of search. We then prove that under given conditions this upper bound is smaller than the payoff the highest and lowest types can guarantee themselves by matching with an agent of equal type. In this proof, we denote $M(k|x)$ the expected discounted

probability that a type x with a type k . For ease of exposition, we denote $x_{min} = 1$ and $x_{max} = N$. In the steady state, the value of search of agent type x can be expressed as such:

$$V(x) = \sum_{k=1}^N \frac{1}{2} (f(x, k) + V(x) - V(k)) M(k|x).$$

Isolating $V(x)$ this gives

$$V(x) = \frac{1}{2 - \bar{M}_x} \sum_{k=1}^N (f(x, k) - V(k)) M(k|x),$$

where $\bar{M}_x = \sum_{k=1}^N M(k|x)$. Then by applying the mimicking argument we obtain the following upper bound on $V(x)$:

$$V(x) \leq \frac{1}{2 - \bar{M}_x} \left[\sum_{k=1}^N f(x, k) M(k|x) - \sum_{k=1}^N \sum_{h=1}^N \pi(k|h) M(h|x) M(k|x) \right].$$

Then set

$$V(x) \leq \bar{V}(x) \equiv \max_{\{M(k|x)\}_{k=1}^N} \frac{1}{2 - \bar{M}_x} \left[\sum_{k=1}^N f(x, k) M(k|x) - \sum_{k=1}^N \sum_{h=1}^N \frac{f(h, k)}{2} M(h|x) M(k|x) \right],$$

where the max is taken over discounted probabilities $\{M(k|x)\}_{k=1}^N$ which must be non-negative and sum to less or equal to one, i.e. $\sum_{k=1}^N M(k|x) \leq 1$. Given optimal discounted probabilities $\{M(k|x)\}_{k=1}^N$, shifting mass between two types m and n to which one has assigned positive discounted probability can not raise \bar{V} . Or, let $\hat{M}^\alpha(n|x) = \alpha M(m|x) + M(n|x)$, $\hat{M}^\alpha(m|x) = (1 - \alpha)M(m|x)$ and $\hat{M}^\alpha(k|x) = M(k|x)$ otherwise. And let

$$v(\alpha) = \sum_{k=1}^N f(x, k) \hat{M}^\alpha(k|x) - \sum_{k=1}^N \sum_{h=1}^N \frac{f(h, k)}{2} \hat{M}^\alpha(h|x) \hat{M}^\alpha(k|x).$$

Then due to the optimality of $\{M(k|x)\}_{k=1}^N$, it must be that $v(0) \geq v(\alpha)$ for all $\alpha \in [0, 1]$, and therefore $v'(0) \leq 0$. Since the roles of m and n are interchangeable, this implies that $v'(0) = 0$. Differentiating gives

$$0 = v'(0) = f(x, n) - f(x, m) - \sum_{k=1}^N [f(k, n) - f(k, m)] M(k|x).$$

Or, when $x = 1$ and $n < m$ this gives

$$f(1, n) - f(1, m) = \sum_{k=1}^N [f(k, n) - f(k, m)] M(k|x) < [f(1, n) - f(1, m)] \sum_{k=1}^N M(k|x)$$

where the latter inequality follows from supermodularity. This is impossible since $\sum_{k=1}^N M(k|x) \leq 1$. This establishes that for $x = 1$ the maximizing weights are zero for all but a single type k .

Likewise, when $x = N$ and $n > m$ this gives

$$f(N, n) - f(N, m) = \sum_{k=1}^N [f(k, n) - f(k, m)]M(k|x) < [f(N, n) - f(N, m)] \sum_{k=1}^N M(k|x),$$

where the latter inequality follows once more from supermodularity. This poses the desired contradiction. In conclusion we have shown that

$$\bar{V}(1) = \frac{\bar{M}_1}{2 - \bar{M}_1} [f(1, y) - \bar{M}_1 \frac{f(y, y)}{2}] \quad \text{and} \quad \bar{V}(N) = \frac{\bar{M}_N}{2 - \bar{M}_N} [f(1, y) - \bar{M}_N \frac{f(y, y)}{2}]$$

for some agent type y (not the same for $x = 1$ and $x = N$). In what remains we derive necessary and sufficient conditions that ensure both $\bar{V}(1) \leq \frac{f(1,1)}{2}$ and $\bar{V}(N) \leq \frac{f(N,N)}{2}$.

Notice that for both $x \in \{1, N\}$ the bound $\bar{M}_x \mapsto \bar{V}(x)$ is weakly convex if $f(x, y) \geq f(y, y)$ and strictly concave otherwise. If it is convex, the maximizing \bar{M}_x must lie on the boundary $\{0, 1\}$. If it is concave, the maximizing \bar{M}_x is characterized by, if interior, the first-order condition. Or,

$$\bar{M}_x = \begin{cases} 1 & \text{if } \frac{f(y,y)}{f(x,y)} \leq \frac{4}{3} \\ 2 - \frac{2\sqrt{f(y,y)(f(y,y)-f(x,y))}}{f(y,y)} & \text{if } \frac{f(y,y)}{f(x,y)} > \frac{4}{3}. \end{cases}$$

If $\bar{M}_x = 1$, then supermodularity implies that the $\bar{V}(x) \leq \frac{f(x,x)}{2}$, all but ensuring self-acceptance. (The same upper bound holds trivially for $\bar{M}_x = 0$.)

Thus focus on the case where $\frac{f(y,y)}{f(x,y)} > \frac{4}{3}$. Then

$$\bar{V}(x) = 2f(y, y) - f(x, y) - 2\sqrt{f(y, y)(f(y, y) - f(x, y))}$$

We then seek to contradict the claim that x 's option value of search may exceed $\frac{f(x,x)}{2}$, thereby ensuring that she will match with her own type.

To simplify the algebra, let $\gamma = f(y, y)/f(x, y)$, and $\theta = f(x, y)/f(x, x)$. Then $\bar{V}(x) = h(\gamma)y$ where h is defined as $h(\gamma) = 2\gamma - 1 - 2\sqrt{\gamma(\gamma - 1)}$ and $\gamma \geq \frac{4}{3}$. Then $\bar{V}(x) \leq \frac{f(x,x)}{2}$ if and only if $h(\gamma) \leq \frac{1}{2\theta}$, or equivalently

$$\theta \leq \frac{1}{2h(\gamma)} \quad \Leftrightarrow \quad \frac{f(x, y)}{f(x, x)} \leq \frac{1}{2h(\frac{f(y,y)}{f(x,y)})}.$$

Now observe that

$$\frac{f(y, y)}{f(x, y)} < \frac{1}{2h(\frac{f(y,y)}{f(x,y)})},$$

which means that log supermodularity is sufficient, but not necessary to achieve this.

D.6 Counter-example

Example 1 (Horizontal differentiation \Rightarrow reciprocity). Consider an economy with two symmetric populations, in the steady state, with three types $x_3 > x_2 > x_1$, and negligible pair-specific shocks. The supermodular match output, $f(x, y)$, exhibits horizontal differentiation, and is given by the following matrix:

f	x_1	x_2	x_3
x_3	0	9	10
x_2	9	10	9
x_1	10	9	0

Suppose there are extremely few agents of type x_1 and x_2 in the search pool, so the probability of meeting them is close to zero. Furthermore, suppose that the discounted expected probability of meeting x_3 is equal to $1/2$. In such case, the medium type x_2 prefers to meet x_1 over himself, upsetting reciprocity.

Proof. Denote q the discounted expected probability of meeting x_3 . We guess (then we will verify) that, upon meeting, x_3 matches with everyone but not x_1 ; and that x_2 prefers to match with x_1 than x_2 :

$$S(x_3, x_1) < 0 \text{ and } S(x_3, x_2) > 0 \text{ and } S(x_1, x_2) > S(x_2, x_1)$$

Accordingly, the value of search of each type is given by:

$$\begin{aligned} V(x_1) &= 0 \\ V(x_3) &= q \frac{f(x_3, x_3)}{2} \\ V(x_2) &= \frac{q}{2-q} (f(x_3, x_2) - \alpha \frac{f(x_3, x_3)}{2}) \end{aligned}$$

First, observe that $S(x_2, x_3) > 0$ is equivalent to $V(x_2) > 0$. So it's trivially satisfied when q and $f(x_3, x_3) - f(x_3, x_2)$ are not too high (it is satisfied in our case). Then, using the expression of the value of search, $S(x_3, x_1) < 0$ and $S(x_1, x_2) > S(x_2, x_2)$, can be expressed as such:

$$\begin{aligned} \frac{q}{2-q} (f(x_3, x_2) - q \frac{f(x_3, x_3)}{2}) &> f(x_2, x_2) - f(x_2, x_1) \\ f(x_3, x_1) &< q \frac{f(x_3, x_3)}{2} \end{aligned}$$

We can now plug the value of q and f , and check that these two inequalities are satisfied, and so our conjectured equilibrium is indeed an equilibrium. \square

E Discussion of Related Literature

E.1 Existence and uniqueness

Existence proofs feature prominently in the literature on matching with random search, with the exclusive focus being on the existence of steady state equilibria. Shimer and Smith (2000) prove existence of a steady state equilibrium under Nash bargaining and quadratic search. Smith (2006) derives an identical result when payoffs are not transferable. Lauer mann and Nöldeke (2015) and Lauer mann et al. (2020) generalize this literature to a broader class of meeting rates. While in Bonneton and Sandmann (2019) we had provided the first existence proof in the non-stationary NTU framework, we here present the first non-stationary *uniqueness* result in the theory of random search matching. Unlike in our earlier work we consider a stochastic environment.

Our uniqueness result casts doubt on the robustness of multiple self-fulfilling equilibrium paths frequently reported in the literature (as for instance in ?, ?, Burdett and Coles (1997), Eeckhout and Lindenlaub (2019)). As a common theme, multiple equilibrium paths arise because agents can perfectly foresee and coordinate their actions.²⁷ In our paper, random entry fosters aggregate uncertainty and thereby obstructs self-fulfilling equilibrium paths. Intuitively, this arises because starting from a given size and composition of the search pool, its evolution is not foreseeable. It follows (unlike in a deterministic framework) that the value of search is continuous in the size and composition of the search pool. This implies that if acceptance thresholds were unduly low along one particular path, they must be low across all neighboring paths, too. Further, if acceptance rules are dominated along some paths, this narrows the number of rationalizable acceptance threshold rules along neighboring paths. Proceeding iteratively, we can eliminate dominated acceptance threshold rules across neighboring paths, ultimately converging to a unique acceptance threshold rule (and thereby implied value of search). The idea that noise, however small, can break the perfect foresight property and thus restore equilibrium uniqueness was first presented in ?.

E.2 The case for pair-specific production shocks

In our view, pair-specific production shocks are essential to our understanding of matching data. We here present evidence that a random-search model without pair-specific production shocks is misidentified.

First, it is known that the unemployment rate at the beginning of an unemployment spell correlates negatively with the contemporary wage (refer to Bils (1985) and Beaudry and DiNardo (1991)). Hagedorn and Manovskii (2013) show empirically (alas without controlling

²⁷For instance, consider the quadratic search technology where the expected number of meetings rises in the mass of agents searching. Here, increased exit of types x decreases the expected number of meetings and thus the value of search for types y . This, in turn, fosters less selective acceptance rules (and greater exit) for agents types y , too. Such feedback loop may give rise to multiple equilibria, whereby agents coordinate on more or less selective match acceptance strategies.

for unobserved heterogeneity) that various variables summarizing past aggregate labor market conditions have explanatory power for current wages only because they are correlated with the sampling rate of production shocks. They lose any predictive power once pair-specific production shocks are accounted for. Quite intuitively, less search in times of high unemployment rates leads to less sampling of production shocks and thus less surplus creation and lower wages.

Secondly, there is evidence from a calibrated version of the stationary random search model due to Shimer and Smith (2000) with on-the-job search. Lopes de Melo (2016) shows that said calibration provides a good fit of key moments of matching data, but falls short on capturing the dispersion of firm fixed-effects in the standard Abowd et al. regression on log wages.²⁸ This dispersion is low in the calibrated model, because wages are not increasing in firm productivity; compared to firm y_1 , firm y_2 may pay a greater wage to x_2 , but a lower wage to x_1 . Reminiscent of here-considered pair-specific production shocks, De Melo (2015) suggests that stochastic worker types increase the dispersion in firm fixed effects in a calibrated model.

Thirdly, the random search matching model without pair-specific production shocks predicts that the time t flow of match creation between two types x, y , denoted $\zeta_t(x, y)$ satisfy a constant ratio property: whenever $\zeta_t(x, y)$ is non-zero for combination of pairs x_1, x_2 and y_1, y_2 , it must satisfy $\frac{\zeta_t(x_2, y_2)}{\zeta_t(x_2, y_1)} = \frac{\zeta_t(x_1, y_2)}{\zeta_t(x_1, y_1)}$.²⁹ This prediction is difficult to assess empirically, for characteristics observable to the econometrician very incompletely describe workers' heterogeneity. Individual types are for the most part unobservable. Hagedorn, Law, Manovskii (2016) show however that unobserved heterogeneity and matching decisions in Shimer and Smith's standard random search and matching model can be identified. Figure 6 of their work plots their estimates of $\zeta_t(x, y)$ for German data. Their empirical strategy identifies each individual with a unique type, workers' types, so that $\mu(x)$ is uniformly distributed and their figure 6 is proportional to both the flow rate of matching and an empirical estimate of $\mu(y)m(x, y)$. Visually, the ratio property does not obtain, rendering their estimates inconsistent with a stationary model without pair-specific production shocks. Reported estimates do satisfy single-peakedness of $x \mapsto \mu(y)m(x, y)$ however, as is implied by our theorem 2 and consistent with a model in which there are pair-specific production shocks.

²⁸Using longitudinal matched employer–employee data, Abowd et al. estimate firm and worker fixed effects of an otherwise standard wage regression (refer to de Melo (2015), table 1 for a synthesis of results across various studies). Initially, they argued that the estimate of fixed effects of the regression corresponds to workers' and firms' productivity, accounting for unobserved heterogeneity. According to this argument, the covariance between worker and firm fixed effects would thus identify sorting. This approach has since been discredited. The reason is that standard search and matching models predict that a given worker's expected wage is non-increasing in firm productivity (more specifically single-peaked under known complementarity conditions as we show in this paper). This is at odds with a key identifying assumption in Abowd et al. which requires that wages be monotone increasing in firm productivity (refer to Gautier Teuling (2006) and Eeckhout and Kircher (2011) for the initial finding, footnote 2 in Hagedorn, Law, Manovskii (2016) for further empirical corroboration, and de Melo (2015) for an in-depth discussion of implications for key moments).

²⁹The prediction that the flow rate of match creation neither depends on the functional form of the meeting rate other than that it satisfy anonymity, nor on stationarity of the economy. To see this, consider the flow rate of match creation of two types x, y in detail: $\zeta_t(x, y) \equiv \beta_t \mu_t(y) \mu_t(x) m_t(x, y)$, where β_t is a time-varying constant governing the speed of meetings, $\mu_t(y)$ is the mass of agent types searching, and $m_t(x, y)$ is the match indicator function, which is zero or one absent pair-specific production shocks.

Recursive Information Design

1 Introduction

In information design in games, a sender does not persuade a single receiver (as in Kamenica and Gentzkow (2011)), but a finite number of players. Thus higher-order beliefs—beliefs about other players’ beliefs—matter. If the sender is to manipulate higher-order beliefs, one must ask: What are higher-order beliefs? Which distributions over higher-order beliefs are consistent with the common prior assumption? And what is the link between higher-order beliefs and equilibrium play?

These are not new questions. The importance of higher-order beliefs has been recognized since Keynes (1936)’s famous ‘beauty contest’ metaphor, according to which traders in financial markets assess an asset’s value based on their belief about other traders’ assessments.¹ Yet, economics and game theory has been operational for the past four decades, not because of a mature understanding of higher-order beliefs, but because the field found ways to bypass them. The key insight is due to Harsanyi (1968): all relevant information about the state of nature, beliefs about other players’ beliefs about the state of nature, as well as beliefs about beliefs about other players’ beliefs can be encoded in what is known as a type. This is a theoretical argument with immense practical implications: information is no longer embedded in a vector of probabilities but can be deduced from the distribution over types. Owing to the success of Harsanyi’s program, economists settled for a construction of higher-order beliefs which works well for the purpose of its construction—showing that the universal type space (the space of coherent belief hierarchies) is compact under the product topology, as established by Mertens and Zamir (1985) and Brandenburger and Dekel (1993)—but is hopelessly untractable otherwise.

This paper proposes to unlearn the lesson that Harsanyi taught us. The notion of a type, crucial to the progress made in the literature, is dispensed with. Accordingly, I drop the canonical representation of information embedded through an information structure, and propose the notion of a belief structure instead: individual equilibrium strategies are solely determined by the higher-order beliefs held, not the distribution thereof.

This is a timely endeavor. There has been a recent surge in attention to study the trans-

¹On the theoretical side, higher-order beliefs have since led to striking predictions regarding uniqueness Carlsson and van Damme (1993) and robustness Kajii and Morris (1997) of equilibrium, thereby making the case that higher-order beliefs are oftentimes pivotal in shaping outcomes. Accordingly, Morris (2019), in his presidential address to the Econometric Society, reiterates the need for a research agenda encompassing both theoretical and empirical research that at the very least makes explicit implicit assumptions on higher-order beliefs.

mission of information in sender-receiver relationships where (i) the sender has the power to pre-commit to a transmission rule, sometimes referred to as a signal, and, (ii) the receiver forms beliefs according to Bayes rule. This framework, due to Brocas and Carrillo (2007), Rayo and Segal (2010) and the seminal paper by Kamenica and Gentzkow (2011), has now been extended to study the design of information transmission rules in games, as in Taneva (2019), Bergemann and Morris (2016b,a) and Mathevet et al. (2019), where there are many receivers. Here the difficulties from the single-receiver problem are compounded in that not only information regarding the state of nature but also higher-order beliefs concerning other players' beliefs matter.

How would a sender want to manipulate higher-order beliefs? In the single-receiver framework, Kamenica and Gentzkow (2011) show that the sender's optimization problem is best represented as a choice of a (Bayes plausible) distribution over first-order beliefs, not abstract messages. In contrast, the literature thus far suggests that the information design problem in games consists in selecting an information structure, i.e., a distribution over abstract messages or types, not explicit belief hierarchies. This leaves the analyses of the single- and many-receiver problem surprisingly disconnected.

Ignoring the epistemic insight afforded by higher-order beliefs has turned out to be a fruitful approach when belief hierarchies can be associated with action recommendations. This is sometimes referred to as a revelation principle (see Bergemann and Morris (2016a) (page 46)) and applies if the sender can not only choose the belief structure but also select the equilibrium. If so, understanding the content of implied belief hierarchies is superfluous. In contrast, I study information design where equilibrium selection is not dictated by the sender. As the computation of (Bayes) Nash equilibria is a hard problem (see Daskalakis et al. (2009)), I focus on the narrower class of supermodular games where equilibrium play is well-understood: individual best-responses rise in other players' actions; equilibria can be ranked; and the least and greatest Bayes Nash equilibrium also arise under weaker notions of rationality, i.e., iterated elimination of dominated strategies and rationalizability. Equipped with a natural equilibrium candidate, the selected equilibrium will be the greatest Bayes Nash equilibrium.

To illustrate, consider the well-known investment game, where two players' payoffs depend both on an unknown state of nature and exhibit complementarities in their investment decisions. (One interpretation is that uncertainty concerns the rule of law in a given country; greater investment leaves individual foreign investors less vulnerable to unlawful expropriations. Another interpretation may concern the efficiency of the country's bureaucracy which improves through learning by doing.)

		Player 2	
		invest	divest
Player 1	invest	s,s	s-1,0
	divest	0,s-1	0,0

If it is common knowledge that $s > 0$, this game exhibits a trivial Bayes Nash equilibrium where both players invest. To implement it, it suffices to send a single message recommending this course of action irrespective of the state of nature. In contrast, under greatest equilibrium se-

lection players play the Bayes Nash equilibrium with the least amount of investment (assuming that ‘divest’ corresponds to the high action).

As is well-known, the information structure which maximizes the probability of investment under greatest equilibrium selection takes the form of an infection argument. The infection argument considers a chain of messages alternating between players such that one message inducing ‘invest’ as a dominant action is contagious for another player to likewise play ‘invest’. As an application, I study the content of thus induced belief hierarchies when there are two states of nature, $\underline{s} < 1 < \bar{s}$. I show that the infection argument admits an important distinction based on the higher-order properties of thus induced beliefs: When $\underline{s} \leq 0$, players are primarily concerned with matching the state. Here, the infection argument features common knowledge of higher-order beliefs as in Carlsson and van Damme (1993) and Mathevet et al. (2019). When $0 < \underline{s} < 1$, players are primarily concerned with coordinating their behavior. Here, the infection argument admits an identical information partition as in the electronic mail game due to Rubinstein (1989) with infinitely many messages and uncertainty over higher-order beliefs at all orders.

Higher-order beliefs

My approach revolves around a novel explicit representation of higher-order beliefs that is bijective to the universal type space due to Mertens and Zamir (1985). A first-order belief is, as is standard, a probability measure over the possible states of nature. Higher-order beliefs are elements in a conditional probability space: a second-order belief assigns a probability, contingent on an assumed state of nature, to other players’ first-order beliefs. A $k + 1$ -order belief assigns a probability, contingent on an assumed state of nature and other players’ up-to- $k - 1$ order beliefs, to those players’ successive k -order beliefs. This construction is tractable, because it avoids Mertens and Zamir’s notion of coherence according to which lower and higher order beliefs can not disagree on events of lower order (simply so because, unlike in Mertens and Zamir, no belief hierarchy contains two higher-order beliefs which support other player’s lower-order beliefs of the same order). It is also natural. To illustrate, consider a patient (he) consulting a doctor (she). The patient’s first-order belief corresponds to the probability he assigns to the events of being healthy or ill. The patient’s second-order belief corresponds to the probability he assigns to the doctor believing that her patient is in good or bad health—conditional on the state of nature. This is a familiar object: suppose the doctor had conducted a test entailing a positive or negative health outlook. The patient’s second-order belief prior to learning what the doctor already knows is then given by a double, namely the type I error α and type II error β of said test: conditional on being healthy (ill), the patient assigns probability $1 - \alpha$ (β) to the doctor believing her patient to be in good health.

Bayes plausibility

In the first part of this work, I offer a novel characterization of the distributions over higher-order beliefs which are consistent with the common prior assumption (see theorem 2). The simplest such characterization is intuitive and easy to define: say that a belief structure is believable if for each player each induced belief hierarchy can be deduced via Bayes rule from the proposed

distribution over beliefs hierarchies. However intuitive, the full content of this definition is unfathomable: if a player is to believe a k -order belief, then he must not only be able to deduce it from the distribution over k -order beliefs, but said update must remain unchanged irrespective of the beliefs of order $\ell > k$ induced.

The need to replace Bayes rule with an equivalent, yet more practical characterization has long been recognized: when there is a single player or receiver as in the canonical work by Kamenica and Gentzkow (2011), Aumann and Maschler (1995) show that Bayes rule is equivalent to a martingale property, hereafter referred to as first-order Bayes plausibility: any player's expected first-order belief must be equal to the prior belief.

In games, knowledge of first-order beliefs does not characterize Bayes Nash equilibrium play. To extend the single-receiver analysis to games, one must characterize the space of believable distributions over entire belief hierarchies. In this work I propose a condition concerning a single player j 's unconditional marginal distributions that I coin (generalized) Bayes plausibility (following the original concept concerning first-order beliefs of an individual receiver). The theorem asserts that a belief structure is believable if and only if the unconditional marginal distributions satisfy Bayes plausibility.

The algebraic details of Bayes plausibility merit further description. How does one deduce the joint distribution over all players' beliefs from the marginal distribution over a single player's beliefs? Why are other players' lower-order beliefs irrelevant to the player's marginal distribution over higher-order beliefs? The key insights are twofold: First, a single player's higher-order beliefs embed information regarding the joint distribution over lower-order beliefs. Accordingly, I show that the joint distribution over (i.e. correlation of) first-order beliefs can be deduced from any player's marginal distribution over second-order beliefs; if the marginal distribution over that player's second-order beliefs satisfies a linear constraint—second-order Bayes plausibility—then the joint distribution over all players' first-order beliefs is believable. Proceeding recursively, I establish that the space of believable joint distributions over k -order beliefs is described by constraints on marginal distributions over $k + 1$ -order beliefs— $k + 1$ -order Bayes plausibility. Secondly, I show that the marginal distribution over an individual player's $k + 1$ -order beliefs does not vary with the state of nature nor other players' up-to- $k - 1$ lower-order beliefs (see corollary 1). For instance, any distribution over the patient's second-order beliefs—how likely is it that the doctor is optimistic with regard to and contingent on the patient's health?—must not depend on the patient's health. Such contingency would allow the patient to deduce additional information about the state of nature which is already embedded in his first-order belief. More generally, an induced $\ell + 1$ -order belief must not lead to a changed Bayesian update concerning other players' k -order beliefs where $k < \ell$.

Information Design in Supermodular Games

In the second part of this work, I study the information design problem in supermodular games under greatest equilibrium selection²: equilibria can be ranked from low to high, and nature

²I focus on supermodular games because equilibrium play is robust and well-understood. Applications include investment games, bank runs, currency attacks, or the coordination motive in research.

selects the greatest Bayes Nash equilibrium.

As is well-known (see Milgrom and Roberts (1990)), optimal individual strategies can be computed as the limit of greatest best-responses to greatest best-responses to ... greatest best-responses to the greatest action profile. Viewed from the traditional information structure perspective, the expectations under which best-responses are computed derive from the distributions over abstract types. Viewed from the belief structure perspective instead, beliefs, not distributions thereof, determine equilibrium play (see proposition 1). This allows me to identify up-to- k order beliefs with the k th best response to the greatest action profile —irrespective of the belief structure chosen by the sender. Finally, I derive a novel representation of the sender's problem (see proposition 4), analogous to the simplifying-the-problem proposition in Kamenica and Gentzkow (2011) in higher-order belief space.

This representation involves a modified utility of the sender which solely depends on the belief hierarchy held by a single player j . It does not depend on other players' actions nor the state of nature as expectations over those can be computed with the knowledge of player j 's sequential belief. The sender's expected utility then follows by taking the expectation over player j 's sequential belief as encoded by the unconditional marginal distribution. The economic insight here is that the sender's expected utility can be approximated by assuming he were facing a single level- k reasoning player and gradually letting the depth of reasoning converge to infinity. Conversely, the marginal gain to the sender of manipulating ever greater k -order beliefs vanishes gradually as k increases.

Related literature

The following briefly mentions some related literature. A description of higher order beliefs was first achieved by Mertens and Zamir (1985), later by Brandenburger and Dekel (1993). Also refer to the excellent book by Maschler et al. (2013), chapter 11. Further topological properties, not a subject of this study, were investigated by Mariotti et al. (2005).

The study of optimal information design has been pioneered by many authors. Here I only mention the seminal paper due to Kamenica and Gentzkow (2011), as well as Rayo and Segal (2010) and Brocas and Carrillo (2007), who study the single-agent case with (i) commitment and (ii) Bayesian updating. Kamenica and Gentzkow (2011)'s elegant simplifying-the-problem proposition builds on the notion of Bayes plausibility, first formulated by Aumann and Maschler (1995) (chapter 1, condition 5.6) in the study of repeated games where the focus is on average payoffs. Aumann and Maschler (1995) establish a concavification result of the informed player. See also Hörner and Skrzypacz (2016) for a beautiful application in repeated games. Later results in information design mirror this initial finding. Shmaya and Yariv (2016) formulate the martingale property of Bayes plausibility to provide testable implications of Bayesian updating in experiments

In recent work, Arieli et al. (2020), Ziegler (2020), Levy et al. (2021) study multi-receiver environments and ask which joint distributions over players' first-order beliefs are believable. Those authors point out that a joint distribution over players' first-order beliefs may satisfy first-order Bayes plausibility, but need not be feasible in that there exists no believable belief

structure inducing said beliefs.³ This seemingly negative result owes to its focus on distributions which do not condition on the state of a nature. In contrast, any conditional joint distribution—one distribution for each state of nature—whose unconditional marginals satisfy first-order Bayes plausibility is believable. Corollary 3 formalizes this point.

The formulation of the information design problem in games was proposed by Taneva (2019), Mathevet et al. (2019), Forges (1993), Bergemann and Morris (2016a), Arieli and Babichenko (2019). Of those, only Mathevet et al. (2019) do not focus on sender-preferred equilibrium selection. They do not yet offer a belief-based representation of the problem but operate under the implicit common prior assumption (see proposition 5 and the discussion thereafter).

The paper closest to mine is concurrent work by Morris et al. (2020). They ask which outcomes are implementable under greatest⁴ equilibrium selection when players have binary actions only. Their objective is different in that they seek to characterize outcomes without making higher-order beliefs explicit. They identify a novel condition, sequential obedience, which concerns distributions over sequences of (gradually decreasing) action recommendations. Here the sender-preferred action must be obediently followed conditional on the expectation that only players who received the recommendation before them choose the sender-preferred action. Section 4 discusses the higher-order belief content of sequential obedience at length and partially generalizes it to multi-action games.

Infection arguments have been proposed by Rubinstein (1989), Carlsson and van Damme (1993), Kajii and Morris (1997), Mathevet et al. (2019), Moriya and Yamashita (2020), Morris et al. (2020), Halac et al. (2021) and Li et al. (2019). I here provide a unified treatment of infection arguments. The key difference is the approach: typically, belief hierarchies are deduced from distributions over signals. Here, I posit belief hierarchies and deduce the unique belief structure which is believable.

³For example, two players may hold optimistic (probability 0.8) or pessimistic beliefs (0.2) that a binary state is high, not low; if both states are equally likely, then first-order Bayes plausibility stipulates that both players must hold the optimistic belief exactly half of the time. It is easy to conjure joint unconditional distributions over first-order beliefs which satisfy first-order Bayes plausibility but are infeasible. This is for instance the case if one player's optimistic first-order belief were to always be jointly induced with the other player's pessimistic first-order belief.

⁴N.B. smallest actually; their selected equilibrium is the least, not the greatest equilibrium. Note that those are interchangeable labels in a supermodular game, so I will maintain the interpretation that the greatest equilibrium is selected.

2 Information Structures

Consider players $\{1, \dots, N\}$ and a finite set of states of nature S (encoding payoff-relevant information). It is common to represent the distribution over information regarding the state of nature and knowledge about other players' knowledge by an information structure. An information structure consists of a (countable) message space, $M = \times_{j=1}^N M^j$, and a mapping $\pi : S \rightarrow \Delta(M)$. Messages carry no intrinsic meaning. Player j deduces, using Bayes rule, the probability that the state of nature is s and that other players have received message $m^{-j} \in \times_{i \neq j} M^i$ conditional on having received message $m^j \in M^j$.

Example 1 (patient-doctor). *Consider a patient, healthy H or ill L , consulting a doctor. The doctor conducts a test with binary outcome, $M^D = \{m'^D, m''^D\}$. The test comes with a type I error α and a type II error β . The information structure is given by*

$$\begin{array}{c|cc} L & m'^D & m''^D \\ \hline m^P & 1 - \beta & \beta \end{array} \qquad \begin{array}{c|cc} H & m'^D & m''^D \\ \hline m^P & \alpha & 1 - \alpha \end{array}$$

3 Belief Structures

In what follows, I replace the concept of an information structure with the concept of a belief structure. Messages will be given the intrinsic meaning of belief hierarchies. Instead of deducing beliefs from distributions over abstract messages, the choice of belief hierarchies determines the distribution thereof. As will become apparent in the next section on supermodular games, imposing this structure comes with the considerable advantage that the notion of Bayes Nash equilibrium is not a function of the chosen information structure, but solely the underlying first- and higher-order beliefs.

I begin by offering a formal description of higher-order beliefs.

Higher-order beliefs

The canonical representation of higher-order beliefs is due to Mertens and Zamir (1985). Its construction is not tractable for the purposes of information design. I thus derive a novel, sequential characterization of higher-order beliefs which, unlike Mertens and Zamirs' construction, draws on conditional probabilities. Here the tedious notion of coherence is vacuously satisfied.⁵

Definition 1 (sequential belief space). *Let $\mathcal{T}_1 = \Delta(S)$ the space of first-order beliefs, and*

$$\mathcal{T}_k = \times_{l=1}^k \mathcal{T}_l \quad \text{and} \quad \mathcal{T}_{k+1} = \{\tau_{k+1}^j : S \times (\mathcal{T}_{k-1})^{N-1} \longrightarrow \Delta((\mathcal{T}_k)^{N-1})\}$$

⁵Coherence, as defined in Mertens and Zamir, arises as an additional constraint because their construction of up-to- k order beliefs entails $k + 1 - \ell$ -copies of lower ℓ -order beliefs. Coherence asserts that all said copies must coincide. Coherence will not be an issue in here, because my construction avoids making duplicate copies of lower-order beliefs. The formal description of Mertens and Zamir's construction is deferred to the appendix.

the space of k -order beliefs. Then define $\mathcal{T}_\infty \equiv \times_{k=1}^\infty \mathcal{T}_k$ (an individual player's) sequential belief space.

Any player's belief hierarchy is an element in the sequential belief space \mathcal{T}_∞ . In words: a player j 's $k + 1$ -order belief $\tau_{k+1}^j \in \mathcal{T}_{k+1}$ gives—conditional on the state of nature $s \in S$ and other players holding up-to- $k - 1$ -order belief $\tau_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}$ —a probability distribution over other players' k -order beliefs in $(\mathcal{T}_k)^{N-1}$. I use the following notation: A profile of k -order beliefs writes as $\tau_k = (\tau_k^1, \dots, \tau_k^N)$. A player j 's up-to- k -order belief writes as $\tau_k^j = (\tau_1^j, \dots, \tau_k^j) \in \mathcal{T}_k$. (Note the difference between bold and normal font.) Similarly, let $\tau_\infty^j = (\tau_1^j, \tau_2^j, \dots) \in \mathcal{T}_\infty$. Next, there are two ways to write a profile of up-to- k -order beliefs $\tau_k = (\tau_k^1, \dots, \tau_k^N) = (\tau_1, \dots, \tau_k) \in (\mathcal{T}_k)^N$. Finally, a profile of up-to- k -order beliefs excluding player j writes as $\tau_k^{-j} = (\tau_k^1, \dots, \tau_k^{j-1}, \tau_k^{j+1}, \dots, \tau_k^N) = (\tau_1^{-j}, \dots, \tau_k^{-j}) \in (\mathcal{T}_k)^{N-1}$.

Example 2 (patient-doctor). Consider the preceding doctor-patient example. Under a common prior $p \in \Delta\{L, H\}$ Bayes rule gives the following doctor's first-order beliefs:

$$\tau_1^D(L) = \frac{p(L)(1 - \beta)}{p(H)\alpha + p(L)(1 - \beta)} \quad \text{and} \quad \tau_1^D(H) = \frac{p(H)(1 - \alpha)}{p(H)(1 - \alpha) + p(L)\beta}.$$

I seek to describe the patient's first- and second-order belief after the doctor has learnt the test outcome, and before she has communicated it to the patient: what does the patient believe the doctor to know? In fact, the patient's second-order belief is well-known: it corresponds to the type I and II error:

$$\tau_1^P = p \quad \text{and} \quad \tau_2^P(\tau_1^D|L) = 1 - \beta, \quad \tau_2^P(\tau_1^D|H) = 1 - \alpha.$$

Drawing on the sequential representation of higher-order beliefs rather than the canonical representation due to Mertens and Zamir (1985) is without loss of generality. This is established by the following theorem, the proof of which is relegated to the appendix.

Theorem 1 (equivalence). Canonical and sequential belief spaces are bijective.

Here, strictly speaking, the bijection is established for equivalence classes of sequential beliefs with countable support, which are said to be identical if they agree on their support.

Bayes plausibility

I now define the central notion of a recursive belief structure which determines the distribution over players' belief hierarchies.⁶

Definition 2 (recursive belief structure). A recursive belief structure is the infinite tuple $\pi_\infty = (\pi_1, \dots)$ where $\pi_1 : S \rightarrow (\mathcal{T}_1)^N$ and $\pi_k : S \times (\mathcal{T}_{k-1})^N \rightarrow \Delta((\mathcal{T}_k)^N)$ for all $k \geq 2$.

Naturally, a recursive belief structure is also an information structure. Indeed, the first- and higher-order beliefs spelled out by a recursive belief structure can be identified with abstract

⁶To avoid technical difficulties, I assume throughout that $\pi_1(\cdot|s)$ and $\pi_{k+1}(\cdot|s, \tau_k)$ have finite support.

messages. Then $\prod_{k=1}^{\infty} \pi_k(\tau_k|s, \tau_{k-1})$ is the probability of inducing the message profile associated with the profile of belief hierarchies $\tau_{\infty} \in (\mathcal{T}_{\infty})^N$. Marginals $\pi_{k+1}^j : \mathcal{S} \times (\mathcal{T}_k)^N \rightarrow \Delta(\mathcal{T}_{k+1})$ are given by

$$\pi_{k+1}^j(\tau_{k+1}^j|s, \tau_k) = \sum_{\tilde{\tau}_{k+1}^{-j} \in (\mathcal{T}_{k+1})^{N-1}} \pi_{k+1}(\tau_{k+1}^j, \tilde{\tau}_{k+1}^{-j}|s, \tau_k).$$

Whereas there is no constraint on the space of information structures consistent with the common prior assumption (simply so because abstract message carry no inherent meaning), the choice of recursive belief structure is not arbitrary. Players must be able to deduce the suggested beliefs via Bayes rule (as formally defined in the appendix). Whenever this is the case, say that a belief structure is believable.

Definition 3 (believable recursive belief structure). *A recursive belief structure is believable if its induced beliefs can be deduced via Bayes rule.*

The space of believable recursive belief structures circumscribes the space of first-and higher order beliefs that can be induced by an information structure. If one wishes to construct a believable recursive belief structure the definition is of little help; to check Bayes rule for each supported belief is computationally intensive and can only be done ex-post, once a belief structure has been proposed. One particular problem is that the distribution over beliefs $(\mathcal{T}_k)^N$ must depend on all up-to- $k-1$ -order beliefs and the state of nature. This allows in principle that distributions over higher-order beliefs change lower-order beliefs (when deduced via Bayes rule). Clearly, this is not a property of believable belief structures, but the associated properties that the belief structure must satisfy to be believable are non-trivial.

The objective here is to propose an equivalent, yet more efficient characterization of believable information structures. To do this, I introduce unconditional distributions: $\lambda_{\infty} = (\lambda_{\infty}^j)_{j \in \{1, \dots, N\}}$ where $\lambda_{\infty}^j = (\lambda_1^j, \dots)$ is an infinite tuple and

$$\lambda_1^j \in \Delta(\mathcal{T}_1) \quad \text{and} \quad \lambda_k^j : \mathcal{T}_{k-1} \rightarrow \Delta(\mathcal{T}_k).$$

The interpretation is that $\lambda_k^j(\tau_k^j|\tau_{k-1}^j)$ is the probability that the belief structure induces player j 's k -order belief τ_k^j conditional on having induced player j 's up-to- $k-1$ -order belief τ_{k-1}^j .

I then introduce a novel (generalized) notion of Bayes plausibility taking as primitives unconditional marginal distributions. As it turns out, this notion subsumes Bayes rule and is the appropriate generalization of Bayes plausibility regarding first-order beliefs, i.e., condition (1), due to Aumann and Maschler (1995) (refer to chapter 1, condition 5.6): any player's expected first-order belief must be equal to the prior belief.

Definition 4 (Bayes plausibility). *An unconditional marginal distribution over higher order*

beliefs λ_∞ is Bayes plausible if it satisfies for all $i, j \in \{1, \dots, N\}$ and $k \geq 1$

$$p(s) = \sum_{\tilde{\tau}_1^j \in \mathcal{T}_1} \tilde{\tau}_1^j(s) \lambda_1^j(\tilde{\tau}_1^j) \quad (1)$$

$$\begin{aligned} \text{and} \quad & \left[\sum_{\tilde{\tau}_{k+1}^j \in \mathcal{T}_{k+1}} \tilde{\tau}_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j}) \lambda_{k+1}^j(\tilde{\tau}_{k+1}^j|\tau_k^j) \right] \pi_k^j(\tau_k^j|s, \tau_{k-1}) \\ & = \left[\sum_{\tilde{\tau}_{k+1}^i \in \mathcal{T}_{k+1}} \tilde{\tau}_{k+1}^i(\tau_k^{-i}|s, \tau_{k-1}^{-i}) \lambda_{k+1}^i(\tilde{\tau}_{k+1}^i|\tau_k^i) \right] \pi_k^i(\tau_k^i|s, \tau_{k-1}), \end{aligned} \quad (2)$$

where $(\pi_k^j)_{j \in \{1, \dots, N\}, k \geq 1}$ follows from λ_∞ via the recursion

$$\begin{aligned} \pi_1^j(\tau_1^j|s) &= \tau_1^j(s) \lambda_1^j(\tau_1^j) / p(s) \\ \pi_k^j(\tau_k^j|s, \tau_{k-1}) \pi_{k-1}(\tau_{k-1}^j, \tau_{k-1}^{-j}|s, \tau_{k-2}) &= \tau_k^j(\tau_{k-1}^{-j}|s, \tau_{k-2}^{-j}) \lambda_k^j(\tau_k^j|\tau_{k-1}^j) \pi_{k-1}^j(\tau_{k-1}^j|s, \tau_{k-2}). \end{aligned} \quad (3)$$

As a first main result of this paper I then report that any believable belief structure must have this property: distributions over induced first- and higher-order beliefs can be deduced via Bayes rule, if and only if they satisfy the (more efficient) notion of Bayes plausibility.

Theorem 2 (characterization result). *Given a common prior $p \in \Delta(S)$, a recursive belief structure π_∞ and an unconditional distribution λ_∞ , the following two statements are equivalent:*

- π_∞ is believable;
- λ_∞ is Bayes plausible.

Here π_∞ and λ_∞ are linked via (3).

Bayes plausibility embeds two slightly subtle properties. First, a belief structure where (2) holds is such that the probability of inducing player j 's $k + 1$ -order belief τ_{k+1}^j —conditional on player j holding up-to- k order belief τ_k^j —varies with other player's k -order belief τ_k^{-j} , but does not depend on lower-order events, namely (s, τ_{k-1}^{-j}) . For instance, at order two this means that the distribution over player j 's second-order beliefs conditional on his first-order belief is identical across all states of nature.

Corollary 1. *Fix a Bayes plausible λ_∞ . Define π_∞ via (3). Then for all $j \in \{1, \dots, J\}$ and $k \geq 1$ the unconditional marginal is independent of the state of nature and other players' lower order beliefs removed by more than one order; i.e.,*

$$\sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k^j, \tilde{\tau}_k^{-j}) \frac{\pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1})}{\pi_k^j(\tau_k^j|s, \tau_{k-1})} = \lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j)$$

for all $(s, \tau_{k-1}^{-j}) \in S \times (\mathcal{T}_{k-1})^{N-1}$.

Secondly, observe that one need not in fact specify unconditional distributions for all players. By manipulating (3), constraints (1) and (2) can equally be expressed as constraints concerning a single player j 's unconditional distribution $\lambda^j = (\lambda_1^j, \lambda_2^j, \dots)$. Accordingly, it is well-defined to say that an individual player's (rather than a profile of) unconditional distribution λ_∞^j

is Bayes plausible. The reason is that higher-order beliefs specify distributions over other players' lower-order beliefs. Whence knowledge of a single player's higher-order belief distribution gives the joint lower-order belief distribution.

Corollary 2. *Fix a player $j \in \{1, \dots, N\}$. If $\lambda_\infty = (\lambda_\infty^i)_{i=1}^N$ is Bayes plausible, then λ_∞^j uniquely defines λ_∞^i for all $i \neq j$. In particular, (3) gives*

$$p(s) \prod_{\ell=1}^{k-1} \pi_\ell(\tau_\ell | s, \tau_{\ell-1}) \pi_k^j(\tau_k^j | s, \tau_{k-1}) = \tau_1^j(s) \lambda_1^j(\tau_1^j) \prod_{\ell=1}^k \tau_{\ell+1}^j(\tau_\ell^{-j} | s, \tau_{\ell-1}^{-j}) \lambda_{\ell+1}^j(\tau_{\ell+1}^j | \tau_\ell^j). \quad (4)$$

The proof of the theorem proceeds via a nested double induction argument. It shows that if the probability of inducing a given higher-order belief were to depend on the state of nature or other player's induced lower order beliefs removed by more than one order, then the event of inducing said higher-order belief must equally change this player's lower-order beliefs as deduced via Bayes rule. If so, finite order beliefs are not sufficient statistics to compute finite order beliefs, thereby contradicting the notion that induced beliefs must be believable.

Feasible joint posterior beliefs

Much effort in the literature has gone into characterizing the feasibility of distributions over joint first-order beliefs (see Arieli et al. (2020), Ziegler (2020), Levy et al. (2021)). In contrast, the scope of theorem 2 is much broader in that it characterizes the space of distributions over belief hierarchies (not only first-order beliefs) which are believable.

Regarding first-order beliefs, one can deduce a simple, yet practical corollary: any joint conditional distribution over posteriors $\pi_1 : S \rightarrow \Delta((\mathcal{T}_1)^N)$ is believable as long as player's marginal distribution over first-order beliefs satisfies Aumann and Maschler (1995)'s first-order Bayes plausibility, i.e., condition (1). No additional condition needs to be imposed.

Corollary 3 (feasibility of joint distributions over first-order beliefs). *A first-order belief structure $\pi_1 : S \rightarrow \Delta((\mathcal{T}_1)^N)$ is believable if and only if associated $\lambda_1^j \in \Delta(\mathcal{T}_1)$ defined through*

$$\lambda_1^j(\tau_1^j) = \sum_{s \in S} \pi_1^j(\tau_1^j | s) p(s) \quad \text{and} \quad \pi_1^j(\tau_1^j | s) = \sum_{\tau_1^{-j} \in (\mathcal{T}_1)^{N-1}} \pi_1(\tau_1^j, \tau_1^{-j} | s) \quad \text{for all } j \in \{1, \dots, N\}$$

is first-order Bayes plausible (1).

Proof of corollary 3. The implication is immediate. Thus, take as given $\pi_1 : S \rightarrow \Delta((\mathcal{T}_1)^N)$ with π_1^j its marginals satisfying first-order Bayes plausibility. One must identify distributions over second-order beliefs, namely λ_2^j for all $j \in \{1, \dots, N\}$, for which conditions (2) for $k = 1$ and (3) for $k = 2$ are satisfied. Many choices are admissible, the most straightforward is to consider degenerate distributions over second-order beliefs: set $\lambda_2^j(\tau_2^j | \tau_1^j) = 1$ if $\tau_2^j : \tau_2^j(\tau_1^{-j} | s) \pi_1^j(\tau_1^j | s) = \pi_1(\tau_1^j, \tau_1^{-j} | s)$ for all $s \in S$, zero otherwise. The conditions are verified immediately, whence theorem 2 asserts that π_1 is believable. \square

The corollary suggests that for the explicit construction of feasible joint distributions over first-order beliefs one begins with marginal distributions for each player. Then any joint distri-

bution of first-order beliefs —conditional on the state of nature —is feasible as long as the joint distribution adds up to the marginals.

4 Information Design in Supermodular Games

Economists study beliefs because they induce actions, not as an end in itself. What is of interest, then, is the mapping from belief structures into Bayes Nash equilibrium. Two obstacles arise. Bayes Nash equilibria tend not to be unique. And to characterize all Bayes Nash equilibria of any games of incomplete information is an insurmountable task: Daskalakis et al. (2009) show that Bayes Nash equilibria are hard to compute.

Thus constrained, the prior literature has (mostly) investigated the information design problem in the paradigm where the sender also selects the equilibrium (see Bergemann and Morris (2016b)). As is well-known, a revelation principle applies: it is without loss of generality to associate an information structure's messages with action recommendations. This insight renders the understanding of the content of induced belief hierarchies superfluous.

In what follows, my focus is different. I do not assume that equilibrium selection is dictated by the sender. Instead, I consider a smaller class of games, namely supermodular games, where equilibrium play is particularly well understood and empirically more plausible: individual best-responses rise in other players' actions; equilibria can be ranked; and the least and greatest Bayes Nash equilibrium also arise under weaker notions of rationality, i.e., iterated elimination of dominated strategies and rationalizability. Equipped with a natural equilibrium candidate, the selected equilibrium will be the greatest Bayes Nash equilibrium —irrespective of the sender's preferences.⁷

Set-up

I begin by defining a supermodular game. Each player $j \in \{1, \dots, N\}$ chooses an action a^j from a finite, ordered action set A^j . Denote $A = \times_{j \in \{1, \dots, N\}} A^j$ the action space of all players and $a = (a^j)_{j \in \{1, \dots, N\}}$ an action profile in A . For $a', a'' \in A$ write $a'' \geq a'$ if $a''^j \geq a'^j$ for all $j \in \{1, \dots, N\}$. Player j 's utility, provided the state of nature is $s \in S$ and players play $a \in A$, is denoted by $u^j(a; s)$.

Assumption 1 (supermodular game). *For all players $j \in \{1, \dots, N\}$ and states of nature $s \in S$ utilities satisfy increasing differences:*

$$u^j(a''^j, a''^{-j}; s) - u^j(a'^j, a''^{-j}; s) \geq u^j(a''^j, a'^{-j}; s) - u^j(a'^j, a'^{-j}; s).$$

The established description of games of incomplete information is that of a normal-form game and an information structure. In that information structures eclipse beliefs, so does the solution concept: the definition of Bayes Nash equilibrium takes as a primitive the distribution

⁷This selection rule equally arises under adversarial equilibrium selection, i.e., nature selects the equilibrium which gives the lowest expected utility to the sender, when the sender's utility is non-decreasing in individual actions.

over abstract types (as encoded by an information structure $\pi : S \rightarrow \Delta(M)$), not the explicit beliefs held. I do not follow this convention here. I define a (pure strategy) ⁸ Bayes Nash equilibrium irrespective of the underlying distribution over types (as encoded by a belief structure), and take belief hierarchies, not the distribution thereof, as primitives:

Definition 5 (Bayes Nash equilibrium). *A (pure strategy) Bayes Nash equilibrium is a strategy profile σ such that each player j 's strategy $\sigma^j : \mathcal{T}_\infty \rightarrow A^j$ constitutes a mutual best-response:*

$$\sigma^j(\tau_\infty^j) \in \arg \max_{a^j \in A^j} \sum_{\tau_\infty^{-j} \in (\mathcal{T}_\infty)^{N-1}} \sum_{s \in S} u^j(a^j, \sigma^{-j}(\tau_\infty^{-j}); s) \tau_1^j(s) \prod_{l=2}^{\infty} \tau_l^j(\tau_{l-1}^{-j} | s, \tau_{l-2}^{-j}).$$

Characterization of greatest equilibrium play

The Bayes Nash equilibrium of a supermodular game is typically not unique. However, equilibria can be ranked. This means that for any two equilibrium strategy profiles, σ' and σ'' , one of the two assigns greater actions to identical beliefs than the other: $\sigma''^j(\tau_\infty^j) \geq \sigma'^j(\tau_\infty^j)$ for all players j and sequential beliefs $\tau_\infty^j \in \mathcal{T}_\infty$. In light of equilibrium multiplicity, I select the greatest Bayes Nash equilibrium. Milgrom and Roberts (1990) show that the greatest Bayes Nash equilibrium follows from iterated elimination of dominated strategies.

Definition 6 (iterated elimination of dominated strategies). *Construct recursively:*

- $\sigma_1^j : \mathcal{T}_1 \rightarrow A^j$ is the greatest best-response of player j , when every other player $i \neq j$ plays the maximal action $\max A^i$:

$$\sigma_1^j(\tau_1^j) \text{ is the greatest element in } \arg \max_{a^j} \sum_{s \in S} u^j(a^j, \max A^{-j}; s) \tau_1^j(s).$$

- $\sigma_k^j : \mathcal{T}_k \rightarrow A^j$ is the greatest best-response of player j , when other players follow σ_{k-1}^{-j} :

$$\sigma_k^j(\tau_k^j) \text{ is the greatest element in } \arg \max_{a^j} \sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}} \sum_{s \in S} u^j(a^j, \sigma_{k-1}^{-j}(\tilde{\tau}_{k-1}^{-j}); s) \tau_1^j(s) \prod_{l=2}^k \tau_l^j(\tilde{\tau}_{l-1}^{-j} | s, \tilde{\tau}_{l-2}^{-j}).$$

The k th iteration of greatest best responses allows to assign to each action a^j the subset of up-to- k -order beliefs which induce it. I call thus defined $\mathfrak{T}_k^j : A^j \rightarrow \mathcal{T}_k$ the space of terminal (order- k) a^j beliefs: $\mathfrak{T}_k^j(a^j) \equiv (\sigma_k^j)^{-1}(a^j)$. An up-to- k -order belief $\tau_k^j \in \mathfrak{T}_k^j(a^j)$ can be thought of as a temporary action recommendation: play a^j or a lesser ranked, but never a greater ranked action. Terminal beliefs will play a crucial role in the application which is to come. But first, proposition 1 makes the formal claim as to how beliefs and equilibrium play relate to one another.

Proposition 1. *The following holds*

⁸The restriction to pure strategies is warranted in supermodular games. Vives (1990), and Van Zandt (2010), building on Topkis (1998), prove their existence. Echenique and Edlin (2004) moreover show they are the unique stable equilibria.

- (i) $\sigma_k^j(\tau_k^j) \geq \sigma_{k+1}^j(\tau_k^j, \tilde{\tau}_{k+1}^j)$ for all $\tilde{\tau}_{k+1}^j \in \mathcal{T}_{k+1}$;
- (ii) $\mathcal{T}_k^j(a^j)$ is convex;
- (iii) $\bar{\sigma} \equiv \lim_{k \rightarrow \infty} \sigma_k$ constitutes a Bayes Nash equilibrium;
- (iv) There exists no other Bayes Nash equilibrium $\hat{\sigma}$ such that $\hat{\sigma}^j(\tau_\infty^j) \geq \bar{\sigma}^j(\tau_\infty^j)$ for some $\tau_\infty^j \in \mathcal{T}_\infty$.

It follows that $\bar{\sigma} = (\bar{\sigma}^j)_{j=1}^N$ is the greatest Bayes Nash equilibrium. In contrast, strategy σ_k^j corresponds to player j 's equilibrium play if that player were level- k reasoning with the anchor that naive players play their greatest action $\max A^j$.

The implementation problem

Having characterized equilibrium play allows one to infer outcomes. An outcome is, of course, a distribution over action profiles and states of nature: $\nu \in \Delta(A \times S)$. What is the outcome induced by a believable belief structure π_∞ ? This is straightforward: following proposition 1, the probability of an action profile a is equal to the probability the belief structure assigns to those sequential beliefs τ_∞ for which $\bar{\sigma}(\tau_\infty) = a$:

$$\nu(a, s) = \sum_{\substack{\tau_\infty \in (\mathcal{T}_\infty)^N: \\ \bar{\sigma}(\tau_\infty) = a}} p(s) \pi_1(\tau_1 | s) \prod_{k=2}^{\infty} \pi_k(\tau_k | s, \tau_{k-1}).$$

In concurrent work, Morris et al. (2020) coin the notion of greatest equilibrium implementable: an outcome is greatest equilibrium implementable if there exist an information (not belief) structure whose greatest Bayes Nash equilibrium induces said outcome.⁹ Using the language of higher-order beliefs, this reads as follows:

Definition 7. *An outcome is greatest equilibrium implementable if there exists a believable belief structure π_∞ which induces it.*

The preceding section made clear, that the notion of a believable belief structure is unwieldy. Hence I proposed to work with Bayes plausible unconditional distributions instead. In keeping with the subsequent representation of the information design problem, I note the following:

Proposition 2. *The following are equivalent:*

- (i) *There exists a believable recursive belief structure π_∞ which induces $\nu \in \Delta(A \times S)$;*

⁹Every belief structure is an information structure. But not every information structure is a belief structure. Additional, uninformative messages may be sent that serve as a correlating device, thereby implementing different outcomes. Ely and Peski (2006)Ely and Peski (2006), and Liu (2009)Liu (2009) provide examples. The greatest equilibrium of a supermodular game is a notable exception. More specifically, any outcome implemented by an information structure can also be implemented by a belief structure. This is due to the result by Milgrom and Roberts (1990): the sets of pure strategy (Bayesian) Nash equilibria, correlated equilibria, and rationalizable strategies have identical bounds. Therefore, if one wishes to implement the greatest equilibrium, one need not draw on a correlation device to achieve that.

(ii) *There exists a player j Bayes plausible unconditional distribution λ_∞^j such that*

$$\nu(a, s) = \sum_{\substack{\tau_\infty \in (\mathcal{T}_\infty)^N: \\ \bar{\sigma}(\tau_\infty) = a}} \tau_1^j(s) \lambda_1^j(\tau_1^j) \prod_{k=1}^{\infty} \tau_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j}) \lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j).$$

This result is an immediate consequence of condition (4) and theorem 2 and does not require further proof. Equivalently one may re-state: an outcome $\nu \in \Delta(A \times S)$ is greatest equilibrium implementable if and only if there exists a player j Bayes plausible unconditional distribution λ_∞^j which induces it (where the meaning of ‘ λ_∞^j induces ν ’ follows from (ii)).

Detour: sequential obedience

Morris et al. (2020) seek to characterize the space of implementable outcomes without making higher-order beliefs explicit. They identify a necessary condition in binary-action games which concerns sequences of action recommendations, not beliefs. Building on proposition 1 and condition (4), I here extend their condition to multi-action games and describe its content in terms of higher-order beliefs.

Fix a profile of sequential beliefs $\tau_\infty = (\tau_\infty^j)_{j=1}^N \in (\mathcal{T}_\infty)^N$. Proposition 1 asserts that this profile corresponds to a sequence of profiles of temporary action recommendations, which, if revised, can only decrease:

$$\tau_\infty \mapsto ((\sigma_1^j(\tau_1^j))_{j=1}^N, (\sigma_2^j(\tau_2^j))_{j=1}^N, \dots).$$

If one wishes to extract the information embedded in temporary action recommendations, but not higher-order beliefs, then one can delete (i) the greatest action recommendation (which was the anchor players started with) as well as (ii) duplicate action recommendations from this sequence: (i) if $\sigma_k^j(\tau_k^j) = \max A^j$, then delete $\sigma_k^j(\tau_k^j)$ from the sequence; (ii) if $\sigma_k^j(\tau_k^j) = \sigma_{k+1}^j(\tau_{k+1}^j)$, then delete $\sigma_{k+1}^j(\tau_{k+1}^j)$. By doing so, one obtains a finite sequence of decreasing action recommendations. Denote \mathcal{A} the space of such sequences.¹⁰ Observe that a sequence of decreasing action recommendations \mathcal{A} may be induced by many distinct profiles of sequential beliefs. The example illustrates this point:

Example 3. *I here anticipate the horizontal infection argument (see figure 2). There are two players, $j \in \{1, 2\}$, and binary actions, $A^j = \{0^j, 1^j\}$ for both players j . Then*

$$\mathcal{A} = \{\emptyset, (0^1), (0^2), (0^1, 0^2), (0^2, 0^1)\}.$$

Any $\alpha \in \mathcal{A}$ determines whether players receive an action recommendation to play low, and if so in which order. Consider two profiles of sequential beliefs $(\tau_\infty^1, \tau_\infty^2) = ((\tau_1^1, \tau_2^1, \dots), (\tau_1^2, \tau_2^2, \dots))$ and $(\tau_\infty^1, \tau_\infty^2) = ((\tau_1^1, \tau_2^1, \dots), (\tau_1^2, \tau_2^2, \dots))$, both in $(\mathcal{T}_\infty)^2$. Given the specifics of the horizontal

¹⁰More specifically, any $\alpha \in \mathcal{A}$ contains finitely many elements from $\cup_{j=1}^N (A^j \setminus \max A^j)$. Write $\alpha = (a_1^{\rho(1)}, a_2^{\rho(2)}, \dots)$ (where $\rho(\ell) \in \{1, \dots, N\}$). Then α has the property that for all $\ell_1 < \ell_2$, if the same player receives a new temporary action recommendation, i.e., $\rho(\ell_1) = \rho(\ell_2)$, then the subsequent action recommendation is strictly inferior, i.e., $a_{\ell_1}^{\rho(\ell_1)} > a_{\ell_2}^{\rho(\ell_2)}$ (simply so, because $\sigma_k^j(\tau_k^j)$ is decreasing in k as asserted by proposition 1).

infection argument, both sequential beliefs map into the same sequence $\alpha = (0^1, 0^2)$. However, the order-belief at which recommendations are issued differs across the two:

$$1^1 = \sigma_1^1(\tau_1^1) \neq \sigma_3^1(\tau_3^1) = \sigma_1^1(\tau_1^1) = 0^1 \quad \text{and} \quad 1^1 = \sigma_2^2(\tau_2^2) \neq \sigma_4^2(\tau_4^2) = \sigma_2^2(\tau_2^2) = 0^1.$$

There are two ways in which a sequence of action recommendations $\alpha \in \mathcal{A}$ maps into an action profile $a \in A$. First, one can ask what is the smallest (and therefore last) action recommended to a given player j : this is given by

$$\underline{\alpha}^j = \begin{cases} \min\{a^j \in \alpha \cap A^j\} & \text{if non-empty} \\ \max A^j & \text{otherwise.} \end{cases}$$

Accordingly, denote $\underline{\alpha} \in A$ the profile of smallest action recommendations. Secondly, definition 6 makes clear that it is only other players' lower-order beliefs which can motivate a player's best-response at higher orders. The ordering of action recommendations α captures this logic. To ease the exposition, define $r_\alpha(a^j)$ the rank at which a given action recommendation in $\cup_{j=1}^N (A^j \setminus \max A^j)$, if ever, is made. Thus, for any $a^j \in \alpha$ define player $i \neq j$'s lowest (and therefore last) action recommendation which precedes the action recommendation a^j :

$$\underline{\alpha}^i(a^j) = \begin{cases} \min\{a^i \in \alpha \cap A^i : r_\alpha(a^i) < r_\alpha(a^j)\} & \text{if non-empty} \\ \max A^i & \text{otherwise.} \end{cases}$$

And denote $\underline{\alpha}^{-j}(a^j) \in A^{-j}$ the profile of smallest action recommendations preceding a^j .

In the context of a binary-action supermodular game, Morris et al. (2020) relate greatest equilibrium implementable outcomes with distributions over sequential action recommendations (and states of nature), $\mathcal{A} \times S$, not profiles of sequential beliefs (and states of nature), $(\mathcal{T}_\infty)^N \times S$. The advantage of this approach is that \mathcal{A} is a lower-dimensional object than $(\mathcal{T}_\infty)^N$. A distribution $\nu_{\mathcal{A}} \in \Delta(\mathcal{A} \times S)$ is referred to as an ordered outcome. An ordered outcome naturally induces an outcome via

$$\nu(a, s) = \sum_{\substack{\alpha \in \mathcal{A}: \\ \underline{\alpha} = a}} \nu_{\mathcal{A}}(\alpha, s).$$

Conversely, a belief structure π_∞ induces an ordered outcome $\nu_{\mathcal{A}}$ via

$$\nu_{\mathcal{A}}(\alpha, s) = \sum_{\substack{(k_\ell)_{\ell=1}^{|\alpha|} \in \mathbb{N}^{|\alpha|}: \\ k_1 \leq k_2 \leq \dots \leq k_{|\alpha|}}} \sum_{j=1}^N \sum_{\substack{\tau_\infty^j \in \mathcal{T}_\infty: \\ \forall a^j \in A^j \cap \alpha \\ k_{r_\alpha(a^j)} = \min\{k: \tau_k^j \in \mathcal{T}_k^j(a^j)\}}} \pi_1(\tau_1|s) \prod_{k=2}^{\infty} \pi_k(\tau_k|s, \tau_{k-1}). \quad (5)$$

Morris et al. (2020) coin the notion of sequential obedience which I here extend to multi-action games:

Definition 8. An ordered outcome $v_{\mathcal{A}} \in \Delta(\mathcal{A} \times S)$ satisfies sequential obedience if

$$a^j \in \arg \max_{a^j \in A^j} \sum_{s \in S} \sum_{\alpha \in \mathcal{A}} u^j(a^j, \underline{\alpha}^{-j}(a^j); s) v_{\mathcal{A}}(\alpha, s)$$

for all players $j \in \{1, \dots, N\}$ and actions $a^j \in A^j$ for which there exists $\alpha \in \mathcal{A}$ so that $a^j \in \alpha$ and $\max_{s \in S} v_{\mathcal{A}}(\alpha, s) > 0$

Sequential obedience says that whenever player j receives the action recommendation to play a^j , then playing a^j must be the greatest best-response given player j 's belief that other players play their smallest action recommendation that preceded a^j .

Proposition 3 (multi-action extension of theorem 1.1, Morris et al. (2020)). An outcome $v \in \Delta(A \times S)$ induced by an ordered outcome $v_{\mathcal{A}} \in \Delta(\mathcal{A} \times S)$ is greatest equilibrium implementable only if $v_{\mathcal{A}} \in \Delta(\mathcal{A} \times S)$ satisfies sequential obedience.

Proof. Suppose that $v_{\mathcal{A}} \in \Delta(\mathcal{A} \times S)$ induces an outcome which is greatest equilibrium implementable. Then there exists a believable belief structure π_{∞} which satisfies (5). It follows that

$$\begin{aligned} & \sum_{s \in S} \sum_{\substack{\alpha \in \mathcal{A}: \\ a^j \in \alpha}} u^j(a^j, \underline{\alpha}^{-j}(a^j); s) v_{\mathcal{A}}(\alpha, s) \\ &= \sum_{s \in S} \sum_{kj \in \mathbb{N}} \sum_{\substack{\tau_{\infty} \in (\mathcal{T}_{\infty})^N: \\ kj = \min\{k \in \mathbb{N} : \tau_k^j \in \mathcal{T}_k^j(a^j)\}}} u^j(a^j, \sigma_{kj-1}^{-j}(\tau_{kj-1}^{-j}); s) p(s) \pi_1(\tau_1 | s) \prod_{\ell=2}^{\infty} \pi_{\ell}(\tau_{\ell} | s, \tau_{\ell-1}) \\ &= \sum_{kj \in \mathbb{N}} \sum_{\substack{\tau_{kj}^j \in \mathcal{T}_{kj}^j: \\ kj = \min\{k \in \mathbb{N} : \tau_k^j \in \mathcal{T}_k^j(a^j)\}}} \sum_{s \in S} \sum_{\tau_{kj-1} \in (\mathcal{T}_{kj-1})^{N-1}} u^j(a^j, \sigma_{kj-1}^{-j}(\tau_{kj-1}^{-j}); s) p(s) \pi_1(\tau_1 | s) \prod_{\ell=2}^{kj-1} \pi_{\ell}(\tau_{\ell} | s, \tau_{\ell-1}) \pi_{kj}^j(\tau_{kj}^j | s, \tau_{kj-1}^{-j}) \\ &= \sum_{kj \in \mathbb{N}} \sum_{\substack{\tau_{kj}^j \in \mathcal{T}_{kj}^j: \\ kj = \min\{k \in \mathbb{N} : \tau_k^j \in \mathcal{T}_k^j(a^j)\}}} \left[\sum_{s \in S} \sum_{\tau_{kj-1} \in (\mathcal{T}_{kj-1})^{N-1}} u^j(a^j, \sigma_{kj-1}^{-j}(\tau_{kj-1}^{-j}); s) \tau_1^j(s) \prod_{\ell=2}^{kj} \tau_{\ell}^j(\tau_{\ell-1}^{-j} | s, \tau_{\ell-2}^{-j}) \right] \lambda_1^j(\tau_1^j) \prod_{\ell=2}^{kj} \lambda_{\ell}^j(\tau_{\ell}^j | \tau_{\ell-1}^j) \end{aligned}$$

where the last equality follows from condition (4). To conclude, note that by construction (see definition 6) the arg max of the term in the square brackets is uniformly a^j , so it must also be the arg max of the whole expression. \square

There are some open questions relating to sequential obedience: the example showed that it is with loss of information to consider sequences of action recommendations, not sequential beliefs. Accordingly, in the context of binary-action games, Morris et al. (2020) show that even if an ordered outcome satisfies sequential obedience (joint with the usual obedience and consistency constraint), this is not yet sufficient to ensure there exists an information structure which greatest equilibrium implements the outcome (induced by the ordered outcome). As a counterexample, it suffices to consider an incomplete information game, where the complete information game for all states of nature admits a Nash equilibrium where all players play their

highest action. If so, there may exist sequentially obedient ordered outcomes which support action recommendations lower than the highest action. Nevertheless, the unique greatest equilibrium implementable outcome is one where all players play their highest action.

How can one ensure that an obedient and consistent outcome where the ordered outcome which induces it satisfies sequential obedience is greatest equilibrium implementable? In the context of binary-action games, Morris et al. (2020) introduce an additional assumption, richness or the dominance state assumption, under which there exists a state of nature for which playing the low action is a dominant strategy. If so, an explicit construction of the information structure guarantees implementability of any (obedient and consistent) outcome induced by a sequentially obedient ordered outcome. But does there exist an extension of the dominance state or richness assumption to multi-action games? Whether sequential obedience is a useful criterion beyond binary-action games, depends on whether this question can be answered in the affirmative.

The information design problem

The literal interpretation of information design is not an implementation problem. Instead, the information (or belief) structure is chosen by an outside sender or information designer prior to learning the state of nature as to maximize expected utility. Thus denote $w(a; s)$ the sender's utility contingent on the state of nature and the action profile chosen. If the belief structure is π_∞ , equilibrium play (governed by the selection of the greatest equilibrium) is $\bar{\sigma}$, and the common prior is p , the sender's value of persuasion V is given by

$$V = \sum_{s \in S} \sum_{a \in A} w(a; s) \sum_{\substack{\tau_\infty \in (\mathcal{T}_\infty)^N: \\ \bar{\sigma}(\tau_\infty) = a}} p(s) \pi_1(\tau_1 | s) \prod_{k=2}^{\infty} \pi_k(\tau_k | s, \tau_{k-1}).$$

Note that if $a^j \mapsto w(a; s)$ is non-increasing for all j , then selection of the greatest equilibrium corresponds to adversarial equilibrium selection: nature adversely selects the Bayes Nash equilibrium which, for any fixed belief structure, gives the lowest expected utility to the sender.

Simplifying-the-problem

As before, to represent the sender's value of persuasion via a belief structure $\pi_\infty = (\pi_1, \dots)$ is unwieldy, for one must consider all players' belief hierarchies jointly. I now propose a representation of the sender's value in terms of an individual player j 's marginal distribution $\lambda_\infty^j = (\lambda_1^j, \dots)$. The key idea is that information regarding other players' beliefs can be deduced from a single player's sequential belief. Define

$$\hat{w}_k^j(\tau_k^j) = \sum_{\tilde{s} \in S} \sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}} w(\sigma_k^j(\tau_k^j), \sigma_{k-1}^{-j}(\tilde{\tau}_{k-1}^{-j}); s) \tau_1^j(\tilde{s}) \prod_{\ell=2}^k \tau_\ell^j(\tilde{\tau}_{\ell-1}^{-j} | \tilde{s}, \tilde{\tau}_{\ell-2}^{-j}).$$

This gives the following generalization of Kamenica and Gentzkow (2011)'s single-agent proposition.

Proposition 4 (simplifying-the-problem). *The following are equivalent:*

- (i) *There exists a believable recursive belief structure π_∞ with value V ;*
- (ii) *There exists a player j Bayes plausible unconditional distribution λ_∞^j such that*

$$V = \lim_{k \rightarrow \infty} \sum_{\tau_k^j \in \mathcal{T}_k} \hat{w}_k^j(\tau_k^j) \lambda_1^j(\tau_1^j) \prod_{\ell=2}^k \lambda_\ell^j(\tau_\ell^j | \tau_{\ell-1}^j).$$

As with proposition 2, this is an immediate consequence of condition (4) and theorem 2 and does not require further proof.

How does this result compare to the earlier representation result on information design in games due to Mathevet et al. (2019)? These authors leverage the fact that any information structure can be separated into a public (Blackwell) experiment which shifts all player's first-order beliefs, and subsequent information structures which are contingent on the outcome of the initial experiment.

Proposition 5 (representation theorem, Mathevet, Perego, Taneva (2020)). *The following are equivalent:*

- (i) *There exists an information structure with value V ;*
- (ii) *There exists a first-order Bayes plausible (1) distribution $\lambda_1^0 \in \Delta \mathcal{T}_1$, and, for all priors $\tau_1^0 \in \text{supp}(\lambda_1^0)$, information structures with value $V^*(\tau_1^0)$ such that*

$$V^*(p) = \sum_{\tau_1^0} V(\tau_1^0) \lambda_1^0(\tau_1^0).$$

The problem with that approach is that the value of persuasion associated with a particular information structure cannot be computed without further knowledge of the properties of induced belief hierarchies. In comparison, $\hat{w}_k^j(\tau_k^j)$ can be computed recursively; and Bayes plausibility of an up-to- k player j unconditional distribution amounts to a finite set of equations. In effect, proposition 4 proposes a well-defined objective function and a set of constraints for the information design problem in supermodular games under greatest equilibrium selection. Note however that in practice the optimization problem becomes exceedingly burdensome as k goes to infinity. Computational properties remain to be investigated.

5 An Application to the Infection Argument

Since the seminal articles due to Rubinstein (1989) and Carlsson and van Damme (1993), it is known that higher-order beliefs can be most potently leveraged in supermodular games through a so-called infection argument (see Kajii and Morris (1997)). Uncertainty about other players' beliefs and their optimal course of action transpires through a chain of individual beliefs held by alternating players: The first belief induces player 1 to choose the low action as a dominant strategy. A second belief held by player 2 assigns sufficiently high positive probability to player

1 holding the first belief, thereby inducing him to also play the low action. A third belief held by player 1 assigns sufficiently high positive probability to player 2 holding the second belief, thereby also inducing him to play the low action. This chain can possibly go on indefinitely and implement the smallest greatest Bayes Nash equilibrium consistent with the common prior.

In this section I ask: what is the content of the infection argument in higher-order belief space? I show that the infection argument admits an important distinction based on whether the chain of infected beliefs concerns uncertainty over the state of nature or higher-order beliefs.

My presentation is guided by the well-known investment game between two players choosing between two actions ('invest' or 'divest' corresponding to the low and high action) where there is a binary (\bar{s} 'good' or \underline{s} 'bad') state of nature. Investment is more worthwhile when the other player chooses to invest. When $\bar{s} > 1 > \underline{s} \geq 0$, the investment game depicts a situation where players must solve a pure coordination problem. When $\bar{s} > 1 > 0 > \underline{s}$, matching the state is the first-order concern instead.

		Player 2	
		invest	divest
Player 1	invest	s,s	s-1,0
	divest	0,s-1	0,0

Under greatest equilibrium selection, players will choose to divest whenever doing so can be sustained as a Bayes Nash equilibrium. As it turns out, the infection argument leads to higher-order belief uncertainty when coordination is players' first-order concern. In contrast, the infection argument does not obfuscate higher-order beliefs when matching the state is players' first-order concern.

Rooted tree depiction

To facilitate the discussion, I represent a belief structure as a rooted tree. The root's neighborhood corresponds to the possible states of nature, edges are labeled by the probabilities of those states. Vertices of depth k correspond to the profiles of $k - 1$ order beliefs, a path from the root to a vertex of length k corresponds to a profile of up-to- $k - 1$ order beliefs. Finally, given a vertex $\tau_k \in (\mathcal{T}_k)^N$, let $(s, \tau_k) \in \mathcal{S} \times (\mathcal{T}_k)^N$ the vertices on the path from the root to τ_k . Then the edge from vertex τ_k to vertex τ_{k+1} is labeled by the probability of inducing τ_{k+1} contingent on having induced s, τ_k already. This probability is given by $\pi_{k+1}(\tau_{k+1}|s, \tau_k)$.

Matching the state: the vertical infection argument

Suppose that $\bar{s} > 1 > 0 > \underline{s}$. Then matching the state is a first-order concern for players, coordination is of secondary importance. Then consider the following belief structure, coined the vertical infection argument.¹¹

¹¹The content of the vertical infection argument was first proposed by Carlsson and van Damme (1993). Mathévet et al. (2019), studying the special case of $\bar{s} = 2$ and $\underline{s} = -1$ and a uniform prior, assert that the sender-preferred information structure under adversarial equilibrium selection is of this form. This is but one solution to the problem,




Figure 1: The vertical infection argument

The vertical infection argument (see figure 1) relies on a belief structure which entails a ranking of player 1's ($[1] < [3] < \dots$) and player 2's ($[2] < [4] < \dots$) first-order beliefs. Player 1, when holding first-order belief $\tau_1^{[1]}$, knows that player 2 holds first-order belief $\tau_2^{[2]}$. There is no such certainty for all higher-ranked first-order beliefs. Player 2, when holding first-order belief $\tau_2^{[2k]}$ assigns positive probability to both player 1's first-order beliefs $\tau_1^{[2k-1]}$ and $\tau_1^{[2k+1]}$; player 1, when holding first-order belief $\tau_1^{[2k+1]}$ assign positive probability to both player 2's first-order beliefs $\tau_2^{[2k]}$ and $\tau_2^{[2k+2]}$. Second-order beliefs are common knowledge in that knowledge of the first-order belief affords knowledge of all higher-order beliefs. Formally, given Bayes plausible λ_1^1 and λ_2^2 which are yet to be characterized, let the second-order belief the unique belief $\tau_2^{[1]}$ for which $\lambda_2^2(\tau_2^{[1]}|\tau_1^{[1]}) = 1$ is Bayes plausible, similarly for player 2.

The following beliefs maximize the probability of joint investment:

- $\tau_1^{[1]}$ invests assuming that player 2 does not;

$$\tau_1^{[1]}(\bar{s})(\bar{s} - 1) + (1 - \tau_1^{[1]}(\bar{s}))(\underline{s} - 1) > 0 \quad \Leftrightarrow \quad \tau_1^{[1]}(\bar{s}) = \frac{1 - \underline{s}}{\bar{s} - \underline{s}} + \epsilon.$$

- $\tau_2^{[2k]}$ invests assuming that $\tau_1^{[2k-1]}$ does, but $\tau_1^{[2k+1]}$ does not;

$$\begin{aligned} & \tau_2^{[2k]}(\bar{s}) \left[(1 - \tau_2^{[2k]}(\tau_1^{[2k+1]}|\bar{s}))\bar{s} + \tau_2^{[2k]}(\tau_1^{[2k+1]}|\bar{s})(\bar{s} - 1) \right] \\ & + (1 - \tau_2^{[2k]}(\bar{s})) \left[(1 - \tau_2^{[2k]}(\tau_1^{[2k+1]}|\underline{s}))\underline{s} + \tau_2^{[2k]}(\tau_1^{[2k+1]}|\underline{s})(\underline{s} - 1) \right] > 0 \\ \Leftrightarrow \quad & \tau_2^{[2k]}(\bar{s}) = \frac{-\underline{s} + \tau_2^{[2k]}(\tau_1^{[2k+1]}|\underline{s})}{\bar{s} - \underline{s} - \tau_2^{[2k]}(\tau_1^{[2k+1]}|\bar{s}) + \tau_2^{[2k]}(\tau_1^{[2k+1]}|\underline{s})} + \epsilon. \end{aligned}$$

- $\tau_2^{[2k+1]}$ invests assuming that $\tau_1^{[2k-1]}$ does, but $\tau_1^{[2k+1]}$ does not;

$$\tau_2^{[2k+1]}(\bar{s}) = \frac{-\underline{s} + \tau_2^{[2k+1]}(\tau_1^{[2k+2]}|\underline{s})}{\bar{s} - \underline{s} - \tau_2^{[2k+1]}(\tau_1^{[2k+2]}|\bar{s}) + \tau_2^{[2k+1]}(\tau_1^{[2k+2]}|\underline{s})} + \epsilon.$$

- $\tau_1^{[0]}$ and $\tau_2^{[0]}$ never invest so that $\tau_1^{[0]}(\bar{s}) = \tau_2^{[0]}(\bar{s}) = 0$.

(Clearly, if $p(\bar{s})$ is large, then these inequalities will not be binding, and investment will be easier to sustain.) Optimality derives from the fact that supported beliefs are the ‘minimal’ terminal

others supporting more first-order beliefs will equally solve the sender's problem (see also the discussion in Morris et al. (2020)). Halac et al. (2021), studying the joint problem of manipulating information and bonus payments in a moral hazard team problem, show that while additional information received under the sender-preferred design under adversarial equilibrium selection are uninformative, the random interim assignments of bonus payments is qualitatively identical to the vertical email structure. Finally, Li et al. (2019) study a related game of regime change with a continuously distributed state, and show that with a continuum of players the sender-preferred information structure under adversarial equilibrium selection is qualitatively of the same form; the frequency $\lambda_1^1(\tau_1^1)$ with which a given first-order belief τ_1^1 is induced corresponds to the mass of agents holding said belief.

beliefs which induce investment (where ‘invest’ corresponds to the least action):

$$\begin{aligned} \tau_{2k-1}^{[2k-1]_1} = (\tau_{1, \dots, 2k-1}^{[2k-1]_1}) \in \mathfrak{T}_{2k-1}^1(\text{invest}) \quad \text{yet} \quad \tau_{2k-2}^{[2k-1]_1} \in \mathfrak{T}_{2k-2}^1(\text{divest}), \\ \text{and} \quad \tau_{2k}^{[2k]_2} = (\tau_{1, \dots, 2k}^{[2k]_2}) \in \mathfrak{T}_{2k}^2(\text{invest}) \quad \text{yet} \quad \tau_{2k-1}^{[2k]_2} \in \mathfrak{T}_{2k-1}^2(\text{divest}); \end{aligned}$$

here higher-order beliefs are the common-knowledge extensions of said beliefs (i.e., $\tau_\ell^{[k]_j}$ are degenerate for all k and $\ell \geq 3$).

What are the frequencies with which said beliefs can be induced? According to theorem 2, the belief structure is fully characterized by the two marginal distributions over first-order beliefs, $\lambda_1^1 \in \Delta\{\tau_1^{[0]_1}, \tau_1^{[1]_1}, \tau_1^{[3]_1}, \dots\}$ and $\lambda_1^2 \in \Delta\{\tau_1^{[0]_2}, \tau_1^{[2]_2}, \tau_1^{[4]_2}, \dots\}$. And Bayes plausibility is satisfied when (see condition (1))

$$p(\bar{s}) = \sum_{\ell \in \{0,1,3,\dots\}} \tau_1^{[\ell]_1}(\bar{s}) \lambda_1^1(\tau_1^{[\ell]_1}) \quad \text{and} \quad p(\bar{s}) = \sum_{\ell \in \{0,2,4,\dots\}} \tau_1^{[\ell]_2}(\bar{s}) \lambda_1^2(\tau_1^{[\ell]_2})$$

and for all $s \in \{\underline{s}, \bar{s}\}$ and $k \in \mathbb{N}$ (see condition (2)))

$$\begin{aligned} \tau_2^{[2k-1]_1}(\tau_1^{[2k]_2} | s) \tau_1^{[2k-1]_1}(s) \lambda_1^1(\tau_1^{[2k-1]_1}) &= (1 - \tau_2^{[2k]_2}(\tau_1^{[2k+1]_1} | s)) \tau_1^{[2k]_2}(s) \lambda_1^2(\tau_1^{[2k]_2}) \\ \tau_2^{[2k]_2}(\tau_1^{[2k+1]_1} | s) \tau_1^{[2k]_2}(s) \lambda_1^2(\tau_1^{[2k]_2}) &= (1 - \tau_2^{[2k+1]_1}(\tau_1^{[2k+2]_2} | s)) \tau_1^{[2k+1]_1}(s) \lambda_1^1(\tau_1^{[2k+1]_1}). \end{aligned}$$

Example 4 (A solution to the vertical infection argument). Consider the vertical infection argument; Mathevet et al. (2019) propose to use three ranks [1], [2], [3] so that $\tau_2^{[3]_1}(\tau_1^{[2]_2} | s) = 1$ for all s . Then the characterization of first-order beliefs (ignoring the $\epsilon > 0$) gives

$$\tau_1^{[1]_1}(\bar{s}) = \frac{1 - \underline{s}}{\bar{s} - \underline{s}}, \quad \tau_1^{[2]_2}(\bar{s}) = \frac{-\underline{s} + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})}{\bar{s} - \underline{s} - \tau_2^{[2]_2}(\tau_1^{[3]_1} | \bar{s}) + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})} \quad \text{and} \quad \tau_1^{[3]_1}(\bar{s}) = \frac{-\underline{s}}{\bar{s} - \underline{s}}.$$

One must then determine $\lambda_1^1 \in \Delta\{\tau_1^{[0]_1}, \tau_1^{[1]_1}, \tau_1^{[3]_1}\}$, $\lambda_1^2 \in \Delta\{\tau_1^{[0]_2}, \tau_1^{[2]_2}\}$ and $\tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})$ and $\tau_2^{[2]_2}(\tau_1^{[3]_1} | \bar{s})$. (Clearly, if $p(\bar{s})$ is large, one induces more optimistic first-order beliefs and still ensures investment with probability one.)

First-order believability gives

$$\lambda_1^1(\tau_1^{[1]_1}) \frac{1 - \underline{s}}{\bar{s} - \underline{s}} + \lambda_1^1(\tau_1^{[3]_1}) \frac{-\underline{s}}{\bar{s} - \underline{s}} = p(\bar{s}) = \lambda_1^2(\tau_1^{[2]_2}) \frac{-\underline{s} + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})}{\bar{s} - \underline{s} - \tau_2^{[2]_2}(\tau_1^{[3]_1} | \bar{s}) + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})}.$$

Second-order believability gives

$$\begin{aligned} \frac{1 - \underline{s}}{\bar{s} - \underline{s}} \lambda_1^1(\tau_1^{[1]_1}) &= (1 - \tau_2^{[2]_2}(\tau_1^{[3]_1} | \bar{s})) \frac{-\underline{s} + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})}{\bar{s} - \underline{s} - \tau_2^{[2]_2}(\tau_1^{[3]_1} | \bar{s}) + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})} \lambda_1^2(\tau_1^{[2]_2}) \\ \left(1 - \frac{1 - \underline{s}}{\bar{s} - \underline{s}}\right) \lambda_1^1(\tau_1^{[1]_1}) &= (1 - \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})) \left(1 - \frac{-\underline{s} + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})}{\bar{s} - \underline{s} - \tau_2^{[2]_2}(\tau_1^{[3]_1} | \bar{s}) + \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s})}\right) \lambda_1^2(\tau_1^{[2]_2}), \end{aligned}$$

where summing the two equations yields $\lambda_1^1(\tau_1^1) = \lambda_1^1(\tau_1^2) \left(1 - \frac{[2]_2([3]_1|\underline{s})\bar{s} - [2]_2([3]_1|\bar{s})\underline{s}}{\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s})}\right)$, and

$$\begin{aligned} \frac{[2]_2([3]_1|\bar{s}) \frac{-\underline{s} + [2]_2([3]_1|\underline{s})}{\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s})} \lambda_1^2(\tau_1^2) = \frac{-\underline{s}}{\bar{s} - \underline{s}} \lambda_1^1(\tau_1^1)} \\ [2]_2([3]_1|\underline{s}) \left(1 - \frac{-\underline{s} + [2]_2([3]_1|\underline{s})}{\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s})}\right) \lambda_1^2(\tau_1^2) = \left(1 - \frac{-\underline{s}}{\bar{s} - \underline{s}}\right) \lambda_1^1(\tau_1^1), \end{aligned}$$

where summing the latter two equations yields $\lambda_1^2(\tau_1^2) \frac{[2]_2([3]_1|\underline{s})\bar{s} - [2]_2([3]_1|\bar{s})\underline{s}}{\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s})} = \lambda_1^1(\tau_1^1)$. Manipulating the preceding, one demonstrates that Bayes plausibility alone characterizes the belief structure. Substituting $\lambda_1^1(\tau_1^1)$ and $\lambda_1^1(\tau_1^1)$ into first-order believability gives

$$[\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s}) - [2]_2([3]_1|\underline{s})\bar{s} + [2]_2([3]_1|\bar{s})\underline{s}] \frac{1 - \underline{s}}{\bar{s} - \underline{s}} + [2]_2([3]_1|\underline{s})\bar{s} - [2]_2([3]_1|\bar{s})\underline{s}] \frac{-\underline{s}}{\bar{s} - \underline{s}} = -\underline{s} + [2]_2([3]_1|\underline{s}).$$

This simplifies to $\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s}) = 2([2]_2([3]_1|\underline{s})\bar{s} - [2]_2([3]_1|\bar{s})\underline{s})$. Substituting back into second-order believability (namely $\lambda_1^2(\tau_1^2) \frac{[2]_2([3]_1|\underline{s})\bar{s} - [2]_2([3]_1|\bar{s})\underline{s}}{\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s})} = \lambda_1^1(\tau_1^1)$ and the second such equation) this implies that

$$\lambda_1^1(\tau_1^1) = \lambda_1^1(\tau_1^1) = \frac{1}{2} \lambda_1^2(\tau_1^2).$$

Then first-order Bayes plausibility gives

$$\frac{1}{2} \lambda_1^2(\tau_1^2) \frac{1 - \underline{s}}{\bar{s} - \underline{s}} + \frac{1}{2} \lambda_1^2(\tau_1^2) \frac{-\underline{s}}{\bar{s} - \underline{s}} = p(\bar{s}) \Leftrightarrow \frac{1}{2} \lambda_1^2(\tau_1^2) = p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - 2\underline{s}}.$$

The probability of joint divestment is given by


$$\lambda_1^1(\tau_1^1) = \lambda_1^2(\tau_1^2) = 1 - \lambda_1^2(\tau_1^2) = 1 - 2p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - 2\underline{s}}.$$

This describes the first-order belief structure. For completeness, also compute marginals $\pi_1^1(\tau_1^1|\bar{s}) = \frac{\lambda_1^1(\tau_1^1)\lambda_1^1(\tau_1^1|\bar{s})}{p(\bar{s})} = \frac{1 - \underline{s}}{1 - 2\underline{s}}$ and $\pi_1^1(\tau_1^1|\underline{s}) = \frac{\lambda_1^1(\tau_1^1)\lambda_1^1(\tau_1^1|\underline{s})}{p(\bar{s})} = \frac{-\underline{s}}{1 - 2\underline{s}}$ and, similarly for \underline{s} , $\pi_1^1(\tau_1^1|\underline{s}) = \frac{\lambda_1^1(\tau_1^1)\lambda_1^1(\tau_1^1|\underline{s})}{1 - p(\bar{s})} = \frac{p(\bar{s})}{1 - p(\bar{s})} \frac{\bar{s} - 1}{1 - 2\underline{s}}$, $\pi_1^1(\tau_1^1|\underline{s}) = \frac{\lambda_1^1(\tau_1^1)\lambda_1^1(\tau_1^1|\underline{s})}{1 - p(\bar{s})} = \frac{p(\bar{s})}{1 - p(\bar{s})} \frac{\bar{s}}{1 - 2\underline{s}}$ and $\pi_1^1(\tau_1^1|\underline{s}) = \frac{\lambda_1^1(\tau_1^1)}{1 - p(\bar{s})} = \frac{1 - 2p(\bar{s})}{1 - p(\bar{s})} \frac{\bar{s} - \underline{s}}{1 - 2\underline{s}}$. (Of course, $\sum_{\ell \in \{0,1,3\}} \lambda_1^1(\tau_1^1) = 1$. Thus $\lambda_1^1(\tau_1^1) = \lambda_1^1(\tau_1^1)$ implies $\lambda_1^1(\tau_1^1) \leq \frac{1}{2}$ or equivalently $\bar{s} - \underline{s} \geq p(\bar{s})$.)

Otherwise the probability of divestment is zero.) Lastly, deduce $\frac{[2]_2([3]_1|\underline{s})}{\tau_2(\tau_1^1|\underline{s})}$ and $\frac{[2]_2([3]_1|\bar{s})}{\tau_2(\tau_1^1|\bar{s})}$ from second-order Bayes plausibility:

$$\begin{aligned} \frac{1 - \underline{s}}{\bar{s} - \underline{s}} \frac{1}{2} = \left(1 - \frac{[2]_2([3]_1|\bar{s})}{\tau_2(\tau_1^1|\bar{s})}\right) \frac{-\underline{s} + [2]_2([3]_1|\underline{s})}{\bar{s} - \underline{s} - [2]_2([3]_1|\bar{s}) + [2]_2([3]_1|\underline{s})} \\ \bar{s} - \underline{s} - \frac{[2]_2([3]_1|\bar{s})}{\tau_2(\tau_1^1|\bar{s})} + \frac{[2]_2([3]_1|\underline{s})}{\tau_2(\tau_1^1|\underline{s})} = 2 \left(\frac{[2]_2([3]_1|\underline{s})}{\tau_2(\tau_1^1|\underline{s})} \bar{s} - \frac{[2]_2([3]_1|\bar{s})}{\tau_2(\tau_1^1|\bar{s})} \underline{s}\right) \Leftrightarrow \begin{aligned} \frac{[2]_2([3]_1|\underline{s})}{\tau_2(\tau_1^1|\underline{s})} &= \frac{\bar{s}}{2\bar{s} - 1} \\ \frac{[2]_2([3]_1|\bar{s})}{\tau_2(\tau_1^1|\bar{s})} &= \frac{\underline{s}}{2\underline{s} - 1}. \end{aligned} \end{aligned}$$

Thus compute $\tau_1^{[2]}(\bar{s}) = \frac{1-2s}{2(\bar{s}-s)}$. In conclusion, I have constructed the following belief structure:



Here $1 - \lambda_1^1(\tau_1^{[0]}) = 1 - (1 - p(\bar{s}))\pi_1^1(\tau_1^{[0]}|s)$ is the probability of joint investment. The left figure considers the belief structure when the probability of investment is less than 1, i.e., $p(\bar{s}) < \frac{1}{2} \frac{1-2s}{\bar{s}-s}$.

Coordination problem: the horizontal infection argument

In the preceding, players' first-order concern was to match the state. Equilibrium play was coordinated through first-order beliefs, whereas second-order beliefs were common knowledge: player 1 knew that, if player 2 held the first-order belief $\tau_1^{[2k]_2}$, then player 2's second-order belief was given by $\tau_2^{[2k]_2}$, symmetrically so for player 2.

Now consider a version of the investment game where coordination is players' first-order concern. Specifically, suppose that $\bar{s} > 1 > \underline{s} > 0$. Here, players need not both match the state and coordinate their actions. Provided that players coordinate successfully, investment is always advantageous—irrespective of the state. In effect, information about the state of nature is less relevant to players.

Then consider what I coin the horizontal infection argument depicted in figure 2.¹² Yet again, the optimistic state \bar{s} serves to prevent miscoordination by motivating one of the two players to invest even when the other player does not. Observe however that unlike in the vertical infection argument, the chain of infection is infinite. Denote

$$\tau_{2k-1}^{[2k+1]_1} = (\tau_{2k-2}^{[2k+1]_1}, \tau_{2k-1}^{[2k+1]_1}) = (\tau_1^{[3]_1}, \tau_2^{[5]_1}, \tau_3^{[5]_1}, \dots, \tau_{2k-2}^{[2k+1]_1}, \tau_{2k-1}^{[2k+1]_1})$$

and

$$\tau_{2k-2}^{[2k]_2} = (\tau_{2k-3}^{[2k]_2}, \tau_{2k-2}^{[2k]_2}) = (\tau_1^{[4]_2}, \tau_2^{[6]_2}, \tau_3^{[6]_2}, \dots, \tau_{2k-3}^{[2k]_2}, \tau_{2k-2}^{[2k]_2})$$

the belief hierarchies. Two things are noteworthy. First, with the exception of $\tau_1^{[1]_1}$ and $\tau_1^{[2]_2}$, first-

¹²The content of the horizontal infection argument was first introduced by Rubinstein (1989). Kajii and Morris (1997) employed it successfully to study robustness of equilibrium to incomplete information. Moriya and Yamashita (2020) find it to be optimal to encourage effort in a moral hazard problem involving teams.




Figure 2: The horizontal infection argument

order beliefs are degenerate and assign probability one to the low state \underline{s} . Secondly, second- and higher-order beliefs are not common knowledge for all finite k . Only as $k \rightarrow \infty$ does it become common knowledge that the state of nature is \underline{s} .

The following beliefs maximize the probability of joint investment:

- $\tau_1^{[1]_1}$ invests assuming that player 2 does not;

$$\tau_1^{[1]_1}(\bar{s})(\bar{s} - 1) + (1 - \tau_1^{[1]_1}(\bar{s}))(\underline{s} - 1) > 0 \quad \Leftrightarrow \quad \tau_1^{[1]_1}(\bar{s}) = \frac{1 - \underline{s}}{\bar{s} - \underline{s}} + \epsilon. \quad (6)$$

- $\tau_2^{[2]_2}$ invests assuming that $\tau_1^{[1]_1}$ does but $\tau_1^{[3]_1}$ does not;

$$\begin{aligned} & \tau_2^{[2]_2}(\bar{s})\bar{s} + (1 - \tau_2^{[2]_2}(\bar{s})) \left[(1 - \tau_2^{[2]_2}(\tau_1^{[3]_1}|\underline{s}))\underline{s} + \tau_2^{[2]_2}(\tau_1^{[3]_1}|\underline{s})(\underline{s} - 1) \right] > 0 \\ \Leftrightarrow & (1 - \tau_2^{[2]_2}(\bar{s}))\tau_2^{[2]_2}(\tau_1^{[3]_1}|\underline{s}) = \underline{s} + \tau_2^{[2]_2}(\bar{s})(\bar{s} - \underline{s}) + \epsilon. \end{aligned} \quad (7)$$

- $\tau_{2k+1}^{[2k+2]_2}$ invests assuming that $\tau_{2k}^{[2k+1]_1}$ does but $\tau_{2k}^{[2k+3]_1}$ does not;

$$\begin{aligned} & (1 - \tau_{2k+1}^{[2k+2]_2}(\tau_{2k}^{[2k+3]_1}|\underline{s}, \tau_{2k-1}^{[2k+1]_1}))\underline{s} + \tau_{2k+1}^{[2k+2]_2}(\tau_{2k}^{[2k+3]_1}|\underline{s}, \tau_{2k-1}^{[2k+1]_1})(\underline{s} - 1) > 0 \\ \Leftrightarrow & \tau_{2k+1}^{[2k+2]_2}(\tau_{2k}^{[2k+3]_1}|\underline{s}, \tau_{2k-1}^{[2k+1]_1}) = \underline{s} + \epsilon. \end{aligned} \quad (8)$$

- $\tau_{2k}^{[2k+1]_1}$ invests assuming that $\tau_{2k-1}^{[2k]_2}$ does but $\tau_{2k-1}^{[2k+2]_2}$ does not;

$$\begin{aligned} & (1 - \tau_{2k}^{[2k+1]_1}(\tau_{2k-1}^{[2k+2]_2} | \underline{s}, \tau_{2k-2}^{[2k]_2})) \underline{s} + \tau_{2k}^{[2k+1]_1}(\tau_{2k-1}^{[2k+2]_2} | \underline{s}, \tau_{2k-2}^{[2k]_2})(\underline{s} - 1) > 0 \\ \Leftrightarrow & \tau_{2k}^{[2k+1]_1}(\tau_{2k-1}^{[2k+2]_2} | \underline{s}, \tau_{2k-2}^{[2k]_2}) = \underline{s} + \epsilon. \end{aligned} \quad (9)$$

Divestment occurs for beliefs $\tau_{\infty}^{[\infty]_1} \equiv \lim_{k \rightarrow \infty} \tau_{2k-1}^{[2k+1]_1}$ and $\tau_{\infty}^{[\infty]_2} \equiv \lim_{k \rightarrow \infty} \tau_{2k-2}^{[2k]_2}$.

Optimality derives once more from the fact that supported beliefs are the ‘minimal’ terminal beliefs which induce investment (where ‘invest’ corresponds to the smallest action):

$$\begin{aligned} & \tau_{2k-1}^{[2k-1]_1} \in \mathfrak{T}_{2k-1}^1(\text{invest}) \quad \text{yet} \quad \tau_{2k-2}^{[2k-1]_1} \in \mathfrak{T}_{2k-2}^1(\text{divest}), \\ \text{and} & \quad \tau_{2k}^{[2k]_2} \in \mathfrak{T}_{2k}^2(\text{invest}) \quad \text{yet} \quad \tau_{2k-1}^{[2k]_2} \in \mathfrak{T}_{2k-1}^2(\text{divest}). \end{aligned}$$

What are the frequencies with which said beliefs can be induced? According to proposition 2 those are characterized by Bayes plausibility, i.e., definition 4. An application of said notion regarding first-order beliefs implies that

$$p(\bar{s}) = \lambda_1^1(\tau_1^{[1]_1}) \tau_1^{[1]_1}(\bar{s}) \quad \text{and} \quad p(\bar{s}) = \lambda_1^2(\tau_1^{[2]_2}) \tau_1^{[2]_2}(\bar{s}). \quad (10)$$

Regarding second-order beliefs one obtains

$$\tau_2^{[3]_1}(\tau_1^{[2]_2} | \underline{s}) \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1}) \pi_1^1(\tau_1^{[3]_1} | \underline{s}) = \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s}) \pi_1^2(\tau_1^{[2]_2} | \underline{s}) \quad (11)$$

$$\text{and} \quad [\tau_2^{[3]_1}(\tau_1^{[4]_2} | \underline{s}) \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1}) + 1 - \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1})] \pi_1^1(\tau_1^{[3]_1} | \underline{s}) = \pi_1^2(\tau_1^{[4]_2} | \underline{s}). \quad (12)$$

Noting that

$$\pi_1^1(\tau_1^{[3]_1} | \underline{s}) = \frac{\lambda_1^1(\tau_1^{[3]_1})}{p(\underline{s})} \quad \text{and} \quad \pi_1^2(\tau_1^{[2]_2} | \underline{s}) = \lambda_1^2(\tau_1^{[2]_2}) \frac{\tau_1^{[2]_2}(\underline{s})}{p(\underline{s})}, \quad \pi_1^2(\tau_1^{[4]_2} | \underline{s}) = \frac{\lambda_1^2(\tau_1^{[4]_2})}{p(\underline{s})} \quad (13)$$

this can be further simplified. Our first objective is to characterize $\tau_1^{[2]_2}$.

Claim 1. $\tau_1^{[2]_2}(\bar{s}) = \tau_1^{[1]_1}(\bar{s}) \frac{1-\bar{s}}{2-\bar{s}} = \frac{(1-\bar{s})(1-\underline{s})}{(\bar{s}-\underline{s})(2-\bar{s})}$ and $\tau_1^{[2]_2}(\underline{s}) = \frac{\bar{s}(2-\underline{s})-1}{(\bar{s}-\underline{s})(2-\underline{s})}$.

Proof. Consider the following system of equations:

$$\begin{aligned} & (1 - \tau_1^{[2]_2}(\bar{s})) \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s}) = \underline{s} + \tau_1^{[2]_2}(\bar{s})(\bar{s} - \underline{s}) \\ & (1 - \underline{s}) \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1}) \lambda_1^1(\tau_1^{[3]_1}) = (1 - \tau_1^{[2]_2}(\bar{s})) \tau_2^{[2]_2}(\tau_1^{[3]_1} | \underline{s}) \lambda_1^2(\tau_1^{[2]_2}) \\ & \underline{s} \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1}) \lambda_1^1(\tau_1^{[3]_1}) + (1 - \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1})) \lambda_1^1(\tau_1^{[3]_1}) = 1 - \lambda_1^2(\tau_1^{[2]_2}) \\ & \tau_1^{[2]_2}(\bar{s}) = \frac{1 - \bar{s}}{\bar{s} - \underline{s}} \frac{1 - \lambda_1^1(\tau_1^{[3]_1})}{\lambda_1^2(\tau_1^{[2]_2})}. \end{aligned}$$

The first equation is (7) (letting ϵ go to zero), the second and third correspond to second-order Bayes plausibility (11) (drawing on (13) and (9)), the fourth follows from equating conditions

in (10) (and drawing on (6)). Summing the second and the third equations, and substituting the first gives

$$\lambda_1^1(\tau_1^{[3]1}) = 1 + \left[\tau_1^{[2]2}(\bar{s}) (\bar{s} - \underline{s}) - (1 - \underline{s}) \right] \lambda_1^2(\tau_1^{[2]2}).$$

Substituting the fourth equation gives

$$(1 - \lambda_1^1(\tau_1^{[3]1})) = (1 - \underline{s})(\lambda_1^2(\tau_1^{[2]2}) + \lambda_1^1(\tau_1^{[3]1}) - 1).$$

Finally, invoking first-order Bayes plausibility (10) gives

$$\frac{p(\bar{s})}{\tau_1^{[1]1}} = (1 - \underline{s}) \left(\frac{p(\bar{s})}{\tau_1^{[2]2}} - \frac{p(\bar{s})}{\tau_1^{[1]1}} \right).$$

□

Having characterized $\tau_1^{[2]2}$, the belief characterization of $\tau_2^{[2]2}$ (7) (once more ignoring the ϵ) gives

$$\tau_2^{[2]2}(\tau_1^{[3]1} | \underline{s}) = \frac{\bar{s} - \underline{s}}{\bar{s}(2 - \underline{s}) - 1}.$$

Then let's take stock and compute the distribution over first-order belief: (10) gives

$$\begin{aligned} \lambda_1^1(\tau_1^{[1]1}) &= p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - \underline{s}}, & \lambda_1^1(\tau_1^{[3]1}) &= 1 - p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - \underline{s}}, \\ \text{and } \lambda_1^2(\tau_1^{[2]2}) &= p(\bar{s}) \frac{(\bar{s} - \underline{s})(2 - \underline{s})}{(1 - \underline{s})(1 - \underline{s})}, & \lambda_1^2(\tau_1^{[4]2}) &= 1 - p(\bar{s}) \frac{(\bar{s} - \underline{s})(2 - \underline{s})}{(1 - \underline{s})(1 - \underline{s})}. \end{aligned}$$

Next, compute conditional marginals through (13):

$$\begin{aligned} \pi_1^1(\tau_1^{[1]1} | \underline{s}) &= \frac{p(\bar{s})}{1 - p(\bar{s})} \frac{\bar{s} - 1}{1 - \underline{s}}, & \pi_1^1(\tau_1^{[3]1} | \underline{s}) &= \frac{1 - p(\bar{s})}{1 - p(\bar{s})} \frac{\bar{s} - \underline{s}}{1 - \underline{s}}, \\ \text{and } \pi_1^2(\tau_1^{[2]2} | \underline{s}) &= \frac{p(\bar{s})}{1 - p(\bar{s})} \frac{\bar{s}(2 - \underline{s}) - 1}{(1 - \underline{s})(1 - \underline{s})}, & \pi_1^2(\tau_1^{[4]2} | \underline{s}) &= \frac{1}{1 - p(\bar{s})} - \frac{p(\bar{s})}{1 - p(\bar{s})} \frac{(2 - \underline{s})(\bar{s} - \underline{s})}{(1 - \underline{s})(1 - \underline{s})}. \end{aligned}$$

Finally, deduce from (11) that

$$\lambda_2^1(\tau_2^{[3]1} | \tau_1^{[3]1}) = \frac{\tau_2^{[2]2}(\tau_1^{[3]1} | \underline{s}) \pi_1^2(\tau_1^{[2]2} | \underline{s})}{\tau_2^{[3]1}(\tau_1^{[2]2} | \underline{s}) \pi_1^1(\tau_1^{[3]1} | \underline{s})} = \frac{\bar{s} - \underline{s}}{\underline{s}(1 - \underline{s})(1 - \underline{s})} \frac{p(\bar{s})}{1 - p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - \underline{s}}}.$$

A remark is in order concerning the parameter values for which λ_2^1 is well-defined. First, if $p(\bar{s}) > \tau_1^{[1]1}(\bar{s}) = \frac{1 - \underline{s}}{\bar{s} - \underline{s}}$, then $\lambda_1^1(\tau_1^{[3]1}) = 0$; absent belief manipulation joint investment is the unique Bayes Nash equilibrium. Conversely, if $p(\bar{s}) \leq \tau_1^{[1]1}(\bar{s})$ then $\lambda_2^1(\tau_2^{[3]1} | \tau_1^{[3]1}) \geq 0$. It is possible that this value exceeds one, in which case investment can be induced with probability one, and higher-order belief manipulation becomes redundant. For now pursue the analysis assuming that $p(\bar{s})$

is sufficiently small.

Regarding third- and higher-order beliefs, an application of Bayes plausibility stipulates that

$$\begin{aligned} \pi_{2k}^1(\tau_{2k}^{[2k+3]_1} | \underline{\mathcal{S}}, \tau_{2k-1}^{[2k+1]_1}, \tau_{2k-1}^{[2k+2]_2}) \\ = \tau_{2k+1}^{[2k+2]_2}(\tau_{2k}^{[2k+3]_1} | \underline{\mathcal{S}}, \tau_{2k-1}^{[2k+1]_1}) \lambda_{2k+1}^2(\tau_{2k+1}^{[2k+2]_2} | \tau_{2k}^{[2k+2]_2}) + 1 - \lambda_{2k+1}^2(\tau_{2k+1}^{[2k+2]_2} | \tau_{2k}^{[2k+2]_2}) \end{aligned} \quad (14)$$

where

$$\pi_{2k}^1(\tau_{2k}^{[2k+3]_1} | \underline{\mathcal{S}}, \tau_{2k-1}^{[2k+1]_1}, \tau_{2k-1}^{[2k+2]_2}) = \frac{\lambda_{2k}^1(\tau_{2k}^{[2k+3]_1} | \tau_{2k-1}^{[2k+1]_1})}{\pi_{2k-1}^2(\tau_{2k-1}^{[2k+2]_2} | \underline{\mathcal{S}}, \tau_{2k-2}^{[2k+1]_1}, \tau_{2k-2}^{[2k]_2})}; \quad (15)$$

and

$$\begin{aligned} \tau_{2k+2}^{[2k+3]_1}(\tau_{2k+1}^{[2k+4]_2} | \underline{\mathcal{S}}, \tau_{2k}^{[2k+2]_2}) \lambda_{2k+2}^1(\tau_{2k+2}^{[2k+3]_1} | \tau_{2k+1}^{[2k+3]_1}) + 1 - \lambda_{2k+2}^1(\tau_{2k+2}^{[2k+3]_1} | \tau_{2k+1}^{[2k+3]_1}) \\ = \pi_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \underline{\mathcal{S}}, \tau_{2k}^{[2k+3]_1}, \tau_{2k}^{[2k+2]_2}) \end{aligned} \quad (16)$$

where

$$\pi_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \underline{\mathcal{S}}, \tau_{2k}^{[2k+3]_1}, \tau_{2k}^{[2k+2]_2}) = \frac{\lambda_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \tau_{2k}^{[2k+2]_2})}{\pi_{2k}^1(\tau_{2k}^{[2k+3]_1} | \underline{\mathcal{S}}, \tau_{2k-1}^{[2k+1]_1}, \tau_{2k-1}^{[2k+2]_2})}. \quad (17)$$

Higher-order beliefs had been characterized by (8) and (9). It remains to determine the Bayes plausible distribution over higher-order beliefs. Iterating over (15) and (17) gives

$$\pi_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \underline{\mathcal{S}}, \tau_{2k}^{[2k+3]_1}, \tau_{2k}^{[2k+2]_2}) = \prod_{\ell=1}^k \frac{\lambda_{2\ell+1}^2(\tau_{2\ell+1}^{[2\ell+4]_2} | \tau_{2\ell}^{[2\ell+2]_2})}{\lambda_{2\ell}^1(\tau_{2\ell}^{[2\ell+3]_1} | \tau_{2\ell-1}^{[2\ell+1]_1})} \frac{\lambda_1^2(\tau_1^{[4]_2})}{1 - p(\bar{s})} \quad (18)$$

$$\pi_{2k}^1(\tau_{2k}^{[2k+3]_1} | \underline{\mathcal{S}}, \tau_{2k-1}^{[2k+1]_1}, \tau_{2k-1}^{[2k+2]_2}) = \prod_{\ell=1}^k \frac{\lambda_{2\ell}^1(\tau_{2\ell}^{[2\ell+3]_1} | \tau_{2\ell-1}^{[2\ell+1]_1})}{\lambda_{2\ell-1}^2(\tau_{2\ell-1}^{[2\ell+2]_2} | \tau_{2\ell-2}^{[2\ell]_2})} (1 - p(\bar{s})) \quad (19)$$

The unconditional distribution λ_∞ can then be derived as a solution of a recurrence equation. Observe that

$$\lambda_1^2(\tau_1^{[4]_2}) \prod_{k=1}^{\infty} \lambda_{2k}^1(\tau_{2k}^{[2k+3]_1} | \tau_{2k-1}^{[2k+1]_1}) \lambda_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \tau_{2k}^{[2k+2]_2})$$

gives the ex-ante probability of divestment.

Claim 2. *The unconditional distribution λ_∞ is the solution to the following recurrence equation.*

$$\lambda_{2k+2}^1(\tau_{2k+2}^{[2k+5]_1} | \tau_{2k+1}^{[2k+3]_1}) = \frac{1 - 2\underline{s}(1 - \underline{s})}{(1 - \underline{s})^2} - \frac{\underline{s}^2}{(1 - \underline{s})^2} \frac{1}{\lambda_{2k}^1(\tau_{2k}^{[2k+3]_1} | \tau_{2k-1}^{[2k+1]_1})}$$

$$\lambda_{2k+3}^2(\tau_{2k+3}^{[2k+6]_2} | \tau_{2k+2}^{[2k+4]_2}) = \frac{1 - 2\underline{s}(1 - \underline{s})}{(1 - \underline{s})^2} - \frac{\underline{s}^2}{(1 - \underline{s})^2} \frac{1}{\lambda_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \tau_{2k}^{[2k+2]_2})}$$

(provided those value are non-negative, if the solution to the recurrence equation admits negative values the induced probability of divestment is zero); starting values are given by

$$\lambda_2^1(\tau_2^{[5]_1} | \tau_1^{[3]_1}) = \frac{\bar{s} - \underline{s}}{\underline{s}(1 - \underline{s})(1 - \underline{s})} \frac{p(\bar{s})}{1 - p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - \underline{s}}}$$

$$\lambda_3^2(\tau_3^{[6]_2} | \tau_2^{[4]_2}) = \frac{1}{1 - \underline{s}} \frac{\lambda_2^1(\tau_2^{[5]_1} | \tau_1^{[3]_1})(1 - \underline{s}(1 - \underline{s})) - \underline{s}^2}{\underline{s} + \lambda_2^1(\tau_2^{[5]_1} | \tau_1^{[3]_1})(1 - \underline{s})}$$

Proof. To simplify the notation, denote $x_2 \equiv \lambda_2^1(\tau_2^{[5]_1} | \tau_1^{[3]_1}) = 1 - \lambda_2^1(\tau_2^{[3]_1} | \tau_1^{[3]_1})$ and

$$x_{2k} \equiv \lambda_{2k}^1(\tau_{2k}^{[2k+3]_1} | \tau_{2k-1}^{[2k+1]_1}), \quad x_{2k+1} \equiv \lambda_{2k+1}^2(\tau_{2k+1}^{[2k+4]_2} | \tau_{2k}^{[2k+2]_2}).$$

Multiplying (18) and (19) and drawing on (14) and (16) gives

$$x_{2k+1} = (\underline{s} + x_{2k+2}(1 - \underline{s}))(\underline{s} + x_{2k+1}(1 - \underline{s}))$$

$$x_{2k+2} = (\underline{s} + x_{2k+3}(1 - \underline{s}))(\underline{s} + x_{2k+2}(1 - \underline{s})).$$

Or, equivalently

$$\frac{1}{1 - \underline{s}} \frac{x_{2k+1}(1 - \underline{s}(1 - \underline{s})) - \underline{s}^2}{\underline{s} + x_{2k+1}(1 - \underline{s})} = x_{2k+2} = \underline{s} \frac{\underline{s} + x_{2k+3}(1 - \underline{s})}{1 - (1 - \underline{s})(\underline{s} + x_{2k+3}(1 - \underline{s}))}.$$

This simplifies to the recurrence equation

$$x_{2k+3} = \frac{1 - 2\underline{s}(1 - \underline{s})}{(1 - \underline{s})^2} - \frac{\underline{s}^2}{(1 - \underline{s})^2} \frac{1}{x_{2k+1}}.$$

A similar construction gives

$$x_{2k+2} = \frac{1 - 2\underline{s}(1 - \underline{s})}{(1 - \underline{s})^2} - \frac{\underline{s}^2}{(1 - \underline{s})^2} \frac{1}{x_{2k}}.$$

It remains to verify the starting values. $x_2 = 1 - \frac{\bar{s} - \underline{s}}{\underline{s}(1 - \underline{s})(1 - \underline{s})} \frac{p(\bar{s})}{1 - p(\bar{s}) \frac{\bar{s} - \underline{s}}{1 - \underline{s}}}$ had been determined earlier. x_3 follows from (as before) multiplying (18) and (19) for $k = 1$ and drawing on (14) and (16):

$$x_3 = \frac{1}{1 - \underline{s}} \frac{x_2(1 - \underline{s}(1 - \underline{s})) - \underline{s}^2}{\underline{s} + x_2(1 - \underline{s})}.$$

□

Remark 1. When $\underline{s} > \frac{1}{2}$, then joint investment is the risk-dominant equilibrium (of the low state complete information game). When $\underline{s} < \frac{1}{2}$, then joint divestment is the risk-dominant equilibrium. Results from Kajii and Morris (1997) imply that if $\underline{s} > \frac{1}{2}$, then the infection argument implements joint investment with probability one, but not so when $\underline{s} < \frac{1}{2}$.

For joint divestment to arise with positive probability one requires infinite realizations of an event, namely the draw of a belief which at no order induces investment. Now recall the Borel-Cantelli lemma: a necessary and sufficient condition for divestment with positive probability is that the sum of the elements of λ_∞ is infinite. Careful inspection of the recurrence equation should reveal that this is precisely the case when $\underline{s} > \frac{1}{2}$.

6 Conclusion

This paper proposed a novel approach to information design in supermodular games under greatest equilibrium selection. Central to this approach was a new explicit representation of higher-order beliefs which allows for greater tractability than the representation proposed thus far in the literature. Building on this characterization I managed to solve a number of problems:

First, I proposed a generalized notion of Bayes plausibility. Bayes plausibility is a feasibility constraint which distributions over joint higher-order beliefs must satisfy if beliefs are to be derived from a common prior. As a corollary, I established a characterization of feasibility concerning joint first-order beliefs about the state of nature.

Secondly, I studied the information design problem in supermodular games under greatest equilibrium selection. I established a formal link between higher-order beliefs, level- k rationality and equilibrium play in supermodular games. Here another advantage of the novel approach shows: equilibrium strategies are solely determined by the beliefs held, not by the distribution thereof. I further derived an analytical representation of the sender's problem which generalizes Kamenica and Gentzkow (2011)'s simplifying-the-problem proposition and showed the higher-order implications of recent results in the literature. I showed that the solution to the information design problem could be approximated by considering level- k reasoning players where the depth of reasoning converged to infinity.

Finally, I studied (as an application of the theory) the well-known infection argument; here, one considers a chain of messages alternating between players such that one message inducing a low action is contagious for another player to likewise choose the low action. I introduced a new distinction based on whether the information structure manipulates first-order beliefs only, or creates uncertainty over higher-order beliefs at all orders.

Many (technical) questions remain to be investigated. (i) Does the explicit representation of the sender's problem admit further insight such as a concavification result? Can one deepen the link between higher-order beliefs and action recommendations? For instance, (ii) does a recursive revelation principle obtain whereby it is without loss to require that no two up-to- k -order beliefs are appended by two distinct $k + 1$ -order beliefs which entail the same action recommendation? (iii) What is the sufficient generalization of the dominance state assumption to multi-action games which ensures that the space of implementable outcomes is indeed characterized by sequential obedience? (iv) Under which additional assumptions regarding players' and the sender's utilities does the sender preferred belief-structure admit the perfect coordination property?

A Appendix

A.1 Higher-order Beliefs

I begin by recalling the canonical definitions and results by Mertens and Zamir (1985).

Definition 9 (belief space). *A belief space of order $k + 1$ is defined recursively as*

$$X_1 = \Delta(S) \quad \text{and} \quad X_{k+1} = X_k \times \Delta(S \times (X_k)^{N-1}).$$

A $k+1$ -order belief $\mu_{k+1} = (\mu_k, \nu_k)$ is an element of the belief space $X_{k+1} = X_k \times \Delta(S \times (X_k)^{N-1})$. Here ν_k is a distribution over the state of nature and the other players' beliefs of lower order. Observe that a belief μ_{k+1} also carries a copy of one's own lower order beliefs (μ_1, \dots, μ_k) . In particular

$$X_{k+1} = X_k \times \Delta(S \times (X_k)^{N-1}) = X_k \times \Delta(S \times (X_{k-1})^{N-1} \times \Delta(S \times (X_{k-1})^{N-1})^{N-1})$$

where the latter holds as an equality up to a coordinate permutation. Thus in principal the distribution ν_k allows one to deduce marginal beliefs over $S \times (X_{k-1})^{N-1}$. In order to avoid that higher order beliefs contradict lower order beliefs recall the notion of coherence.

Definition 10 (coherence). *A $k+1$ -order belief $\mu_{k+1} = (\mu_1, \nu_1, \dots, \nu_k) \in X_{k+1}$ is said to be coherent if*

$$\text{marg}_S \nu_1 = \mu_1 \quad \text{and} \quad \nu_{l-1} = \text{marg}_{S \times (X_{l-1})^{N-1}} \nu_l \quad \forall l \in \{2, \dots, k\}.$$

Accordingly denote Z_{k+1} the space of coherent $k + 1$ -order beliefs. Since higher order beliefs determine lower order beliefs one can simplify $Z_{k+1} \simeq \Delta(S \times (Z_k)^{N-1})$, where \simeq stands for homeomorphic (bijective, continuous with continuous inverse).

The notion of belief spaces lies at the core of the universal type space as introduced by Harsanyi. In particular a player's type is the summary of the player's beliefs about the state of nature as well as her higher order beliefs. Such information is encoded up to order $k + 1$ in the player's belief space of order $k + 1$. In order to encompass all higher order beliefs one takes the projective limit. This is the space of sequences $(\mu_1, \mu_2, \dots) \in \times_{k=1}^{\infty} Z_k \equiv T$, where for every $k \in \mathbb{N}$ the belief hierarchy $\mu_k \in Z_k$ is the projection of the belief hierarchy $\mu_{k+1} \in Z_{k+1}$ on Z_k . An element $t = (\mu_1, \mu_2, \dots) \in T$ is called a type.

I now turn to sequential belief spaces. For convenience I further amend the notation and define the $k + 1$ -support of beliefs $\mathcal{B}_{k+1} = (\mathcal{T}_k)^{N-1}$.

Definition 1' (sequential belief space). *Let $\mathcal{B}_1 = S$, $\mathcal{T}_1 = \Delta(S)$ and define for all $k \geq 1$*

$$\mathcal{B}_{k+1} = (\mathcal{T}_k)^{N-1}, \quad \mathcal{B}_{k+1} = \times_{l=1}^{k+1} \mathcal{B}_l \quad \text{and} \quad \mathcal{T}_{k+1} = \{\tau_{k+1}^j : \mathcal{B}_k \longrightarrow \Delta(\mathcal{B}_{k+1})\}.$$

Then call $\mathcal{T}_{k+1} = \times_{l=1}^{k+1} \mathcal{T}_l$ a sequential belief space of order $k + 1$.

The following result establishes a bijection in between canonical and sequential belief spaces. Here identify sequential belief spaces with their equivalence class, where any two beliefs $\tilde{\tau}_k^j = (\tilde{\tau}_1^j, \dots, \tilde{\tau}_k^j) \in \mathcal{T}_k$ and $\hat{\tau}_k^j = (\hat{\tau}_1^j, \dots, \hat{\tau}_k^j) \in \mathcal{T}_k$ are identical if they agree on their support.

Theorem 1' (equivalence). *Canonical and sequential belief spaces are bijective: $\times_{l=1}^k \mathcal{T}_l \simeq Z_k$ for all $k \in \mathbb{N}$.*

Proof. For $k = 1$ the result $\mathcal{T}_1 \simeq Z_1$ readily holds. For $k = 2$

$$\begin{aligned} Z_2 &\simeq \Delta(S \times (Z_1)^{N-1}) \simeq \Delta(S) \times \Delta((Z_1)^{N-1})^S \\ &= \mathcal{T}_1 \times \{\mathcal{B}_1 \longrightarrow \Delta(\mathcal{B}_2)\} = \mathcal{T}_1 \times \mathcal{T}_2 \end{aligned}$$

Here I used the fact that $\Delta(A \times B) \simeq \Delta(A) \times \Delta(B)^A$. This defines a bijection provided that two properties hold:

- (i) identify any two (α', β') and $(\alpha'', \beta'') \in \Delta(A) \times \Delta(B)^A$ for which $\alpha' = \alpha'' \equiv \alpha$ and $\beta'(\cdot|a) = \beta''(\cdot|a) \in \Delta(B)$ for all a in the support of α ;
- (ii) for any $\beta \in \Delta(B)^A$ and measurable $a, a' \subset A, b \in B$ such that $b = a' \times P_{B \setminus A} b$ (P short for projection) and $a \cap a' = \emptyset, \beta(b|a) = 0$.

When I write $\Delta(A \times B) \simeq \Delta(A) \times \Delta(B)^A$ I will assume those properties to hold throughout the proof. Rewriting the belief space of order $k + 1 \geq 3$ I obtain

$$\begin{aligned} Z_{k+1} &= Z_k \times \Delta(S \times (Z_k)^{N-1}) \simeq \Delta(S \times (Z_k)^{N-1}) \\ &= \Delta(S \times (Z_{k-1})^{N-1} \times (\Delta(S \times (Z_{k-1})^{N-1}))^{N-1}) \\ &\simeq \Delta(S \times (Z_{k-1})^{N-1}) \times \left[\Delta(\Delta(S \times (Z_{k-1})^{N-1})^{N-1}) \right]^{S \times (Z_{k-1})^{N-1}} \\ &\simeq Z_k \times \left[\Delta((Z_k)^{N-1}) \right]^{S \times (Z_{k-1})^{N-1}}. \end{aligned}$$

By induction hypothesis $Z_k \simeq \times_{l=1}^k \mathcal{T}_l$ and

$$\begin{aligned} \left[\Delta((Z_k)^{N-1}) \right]^{S \times (Z_{k-1})^{N-1}} &\simeq \left\{ \tau : S \times \left[\times_{l=1}^{k-1} (\mathcal{T}_l)^{N-1} \right] \longrightarrow \Delta\left(\left[\times_{l=1}^k (\mathcal{T}_l)^{N-1} \right] \right) \right\} \\ &\simeq \left\{ \tau : \mathcal{B}_k \longrightarrow \Delta(\mathcal{B}_k \times \mathcal{B}_{k+1}) \right\} \\ &\simeq \left\{ \tau : \mathcal{B}_k \longrightarrow \Delta(\mathcal{B}_{k+1}) \right\} \simeq \mathcal{T}_{k+1} \end{aligned}$$

where I made use of property (ii). □

A.2 Characterization of believable belief structures

Fix a recursive belief structure $\pi_\infty = (\pi_1, \dots)$ (each π_k admitting a finite support). Denote \mathcal{T}_k^j (the finite set of) player j 's up-to- k order beliefs which lie in the support of $\pi_k = (\pi_1, \dots, \pi_k)$.

Next define the belief updater $\iota_{k|\ell}^j : \mathcal{T}_\ell^j \rightarrow \mathcal{T}_k$ where $\ell \geq k$. The updater interprets an induced player j up-to- ℓ order belief in \mathcal{T}_ℓ^j as a signal and uses Bayes rule to compute the k -order belief. In particular, for $k = \ell$,

$$\begin{aligned} \iota_{1|1}^j(\tau_1^j)(s) &\equiv \frac{p(s)\pi_1^j(\tau_1^j|s)}{\sum_{\tilde{s}} \pi_1^j(\tau_1^j|\tilde{s})p(\tilde{s})} \\ \iota_{k|k}^j(\tau_k^j)(\tau_{k-1}^{-j}|s, \tau_{k-2}^{-j}) &\equiv \frac{\pi_{k-1}(\tau_{k-1}|s, \tau_{k-2})\pi_k^j(\tau_k^j|s, (\tau_{k-2}, \tau_{k-1}))}{\sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k-1}(\tau_{k-1}, \tilde{\tau}_{k-1}^{-j}|s, \tau_{k-2})\pi_k^j(\tau_k^j|s, \tau_{k-2}, \tau_{k-1}^{-j}, \tilde{\tau}_{k-1}^{-j})}. \end{aligned}$$

More generally, for $\ell > k$,

$$\begin{aligned} \iota_{1|\ell}^j(\tau_\ell^j)(s) &= \frac{p(s) \sum_{\tau_{\ell-1}^{-j} \in \mathcal{T}_{\ell-1}} \prod_{m=1}^{\ell-1} \pi_m(\tau_m|s, \tau_{m-1})\pi_\ell^j(\tau_\ell^j|s, \tau_{\ell-1})}{\sum_{\tilde{s}} p(\tilde{s}) \sum_{\tau_{\ell-1}^{-j} \in \mathcal{T}_{\ell-1}} \prod_{m=1}^{\ell-1} \pi_m(\tau_m|\tilde{s}, \tau_{m-1})\pi_\ell^j(\tau_\ell^j|\tilde{s}, \tau_{\ell-1})} \\ \iota_{k|\ell}^j(\tau_\ell^j)(\tau_{k-1}^{-j}|s, \tau_{k-2}^{-j}) &= \frac{\sum_{(\tau_k^{-j}, \dots, \tau_{\ell-1}^{-j}) \in \times_{m=k}^{\ell-1} (\mathcal{T}_m)^{N-1}} \prod_{m=k-1}^{\ell-1} \pi_m(\tau_m|s, \tau_{m-1})\pi_\ell^j(\tau_\ell^j|s, \tau_{\ell-1})}{\sum_{\tau_{k-1}^{-j}} \sum_{(\tau_k^{-j}, \dots, \tau_{\ell-1}^{-j})} \pi_{k-1}(\tau_{k-1}, \tilde{\tau}_{k-1}^{-j}|s, \tau_{k-2}) \prod_{m=k}^{\ell-1} \pi_m(\tau_m|s, \tau_{k-2}, \tau_{k-1}^{-j}, \tilde{\tau}_{k-1}^{-j}, \tau_k, \dots, \tau_{m-1})\pi_\ell^j(\tau_\ell^j|s, \tau_{k-2}, \tau_{k-1}^{-j}, \tilde{\tau}_{k-1}^{-j}, \tau_k, \dots, \tau_{\ell-1})}. \end{aligned}$$

To be believable, a recursive belief structure must satisfy two properties. First, the induced up-to- k order belief must agree with the k order belief computed using Bayes rule. Secondly, the induced up-to- k order belief τ_k^j must be a sufficient statistic to compute the k -order belief τ_k^j .

Definition 11 (believable recursive belief structure). *A recursive belief structure $\pi_\infty = (\pi_1, \dots)$ is believable if, for all $k \in \mathbb{N}$ and $\ell \geq k$, $\iota_{k|\ell}^j(\tau_\ell^j) = \tau_k^j$ for all $\tau_\ell^j = (\tau_{k-1}^j, \tau_k^j, \tau_{k+1}^j, \dots, \tau_\ell^j) \in \mathcal{T}_\ell^j$.*

The proof of theorem 2 goes as follows:

proof of theorem 2. Begin with the implication: given a believable π_∞ , construct $\lambda_1^j \in \Delta(\mathcal{T}_1)$ as in Kamenica and Gentzkow (2011): $\lambda_1^j(\tau_1^j) \equiv \sum_{\tilde{s} \in S} \pi_1^j(\tau_1^j|\tilde{s})p(\tilde{s})$. At higher-orders, construct an auxiliary distribution $\tilde{\lambda}_{k+1}^j : S \times (\mathcal{T}_{k-1})^N \times \mathcal{T}_k \rightarrow \Delta(\mathcal{T}_{k+1})$:

$$\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}^j, \tau_k^j) \equiv \sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k+1}^j(\tau_{k+1}^j|s, (\tau_{k-1}^j, \tau_k^j, \tilde{\tau}_k^{-j})) \frac{\pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1}^j)}{\pi_k^j(\tau_k^j|s, \tau_{k-1}^j)}.$$

The proof of the implication consists of showing that $\tilde{\lambda}_{k+1}^j$ only depends on player j 's up-to- k order beliefs, i.e.,

$$\lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j) \equiv \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s', \tau_{k-1}^{\prime-j}, \tau_{k-1}^j, \tau_k^j) = \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s'', \tau_{k-1}^{\prime\prime-j}, \tau_{k-1}^j, \tau_k^j)$$

for all $s', s'' \in S$ and $\tau_{k-1}^{\prime-j}, \tau_{k-1}^{\prime\prime-j} \in (\mathcal{T}_{k-1})^{N-1}$. If so, one immediately verifies that thus constructed unconditional distribution $\lambda_\infty = (\lambda_1, \dots)$ satisfies conditions (1), (2) and (3).

Proceed by two nested induction arguments A and B.

A Base case: begin with $\ell = 2$ and $k = 1$ to show that $\tilde{\lambda}_2^j$ does not vary with $s \in S$:

$$\iota_{1|2}^j(\tau_2^j)(s) = \frac{p(s) \sum_{\tau_1^{-j}} \pi_1(\tau_1|s) \pi_2^j(\tau_2^j|s, \tau_1^{-j})}{\sum_{\tilde{s}} p(\tilde{s}) \sum_{\tau_1^{-j}} \pi_1(\tau_1|\tilde{s}) \pi_2^j(\tau_2^j|\tilde{s}, \tau_1^{-j})} = \frac{p(s) \tilde{\lambda}_2^j(\tau_2^j|s, \tau_1^j) \pi_1^j(\tau_1^j|s)}{\sum_{\tilde{s}} p(\tilde{s}) \tilde{\lambda}_2^j(\tau_2^j|\tilde{s}, \tau_1^j) \pi_1^j(\tau_1^j|\tilde{s})}.$$

In particular, for any two $\tau_2^j, \tau_2^{\prime j} : (\tau_1^j, \tau_2^j), (\tau_1^j, \tau_2^{\prime j}) \in \mathcal{T}_2^j$ believability requires that $\iota_{1|2}^j(\tau_1^j, \tau_2^j) = \iota_{1|2}^j(\tau_1^j, \tau_2^{\prime j})$. Whence

$$\frac{\tilde{\lambda}_2^j(\tau_2^j|s, \tau_1^j)}{\tilde{\lambda}_2^j(\tau_2^{\prime j}|s, \tau_1^j)} = \frac{\sum_{\tilde{s}} p(\tilde{s}) \tilde{\lambda}_2^j(\tau_2^j|\tilde{s}, \tau_1^j) \pi_1^j(\tau_1^j|\tilde{s})}{\sum_{\tilde{s}} p(\tilde{s}) \tilde{\lambda}_2^j(\tau_2^{\prime j}|\tilde{s}, \tau_1^j) \pi_1^j(\tau_1^j|\tilde{s})} \quad \text{for all } s \in S.$$

Observe that the RHS does not depend on $s \in S$ so that the LHS ratio is constant across all states of nature. And since $\tilde{\lambda}_2^j(\cdot|s, \tau_1^j)$ is a probability, it follows that not only the ratio, but also the values do not change with the state of nature s : $\tilde{\lambda}_2^j(\cdot|s', \tau_1^j) = \tilde{\lambda}_2^j(\cdot|s'', \tau_1^j)$ for all $s', s'' \in S$.

A Induction step: Fix k . Suppose that (induction hypothesis A) for all $\ell < k + 1$, $\lambda_\ell^j(\tau_\ell^j|\tau_{\ell-1}^j)$ is well-defined. Then show that so is $\lambda_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}^j, \tau_k^j)$, i.e., $\tilde{\lambda}_{k+1}^j$ does not vary with $(s, \tau_{k-1}^j) \in S \times (\mathcal{T}_{k-1})^{N-1}$. To show this, the proof proceeds via an induction argument B (nested within the induction argument A):

B Induction base: Begin by showing that $\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}^j, \tau_k^j)$ does not vary with $\tau_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}$. Consider the belief updater $\iota_{k|k+1}^j(\tau_{k+1}^j)(\tau_{k-1}^{-j}|s, \tau_{k-2}^j)$, which is equal to

$$\begin{aligned} & \frac{\sum_{\tau_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k-1}(\tau_{k-1}|s, \tau_{k-2}^j) \pi_k(\tau_k|s, \tau_{k-1}^j) \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_k^j)}{\sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}} \sum_{\tau_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k-1}(\tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}|s, \tau_{k-2}^j) \pi_k(\tau_k|s, \tau_{k-2}^j, \tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}) \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-2}^j, \tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}, \tau_k^j)} \\ &= \frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}^j, \tau_k^j) \pi_k^j(\tau_k^j|s, \tau_{k-1}^j) \pi_{k-1}(\tau_{k-1}|s, \tau_{k-2}^j)}{\sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-2}^j, \tilde{\tau}_{k-1}^{-j}, \tau_{k-1}^j, \tau_k^j) \pi_k^j(\tau_k^j|s, \tau_{k-2}^j, \tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}) \pi_{k-1}(\tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}|s, \tau_{k-2}^j)}. \end{aligned}$$

And for any two $\tau_{k+1}^j, \tau_{k+1}^{\prime j} : (\tau_k^j, \tau_{k+1}^j), (\tau_k^j, \tau_{k+1}^{\prime j}) \in \mathcal{T}_{k+1}^j$ believability requires that $\iota_{k|k+1}^j(\tau_k^j, \tau_{k+1}^j) = \iota_{k|k+1}^j(\tau_k^j, \tau_{k+1}^{\prime j})$. Whence

$$\begin{aligned} & \frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}^j, \tau_k^j)}{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^{\prime j}|s, \tau_{k-1}^j, \tau_k^j)} \\ &= \frac{\sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-2}^j, \tilde{\tau}_{k-1}^{-j}, \tau_{k-1}^j, \tau_k^j) \pi_k^j(\tau_k^j|s, \tau_{k-2}^j, \tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}) \pi_{k-1}(\tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}|s, \tau_{k-2}^j)}{\sum_{\tilde{\tau}_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^{\prime j}|s, \tau_{k-2}^j, \tilde{\tau}_{k-1}^{-j}, \tau_{k-1}^j, \tau_k^j) \pi_k^j(\tau_k^j|s, \tau_{k-2}^j, \tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}) \pi_{k-1}(\tau_{k-1}^j, \tilde{\tau}_{k-1}^{-j}|s, \tau_{k-2}^j)} \end{aligned}$$

for all $\tau_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}$. Once more, the RHS does not vary as τ_{k-1}^{-j} changes. And since $\tilde{\lambda}_{k+1}^j(\cdot|s, \tau_{k-1}, \tau_k^j)$ is a probability,

$$\tilde{\lambda}_{k+1}^j(\cdot|s, \tau_{k-2}, \tau_{k-1}^{-j}, \tau_k^j) = \tilde{\lambda}_{k+1}^j(\cdot|s, \tau_{k-2}, \tau_{k-1}^{\prime-j}, \tau_k^j) \quad \text{for all } \tau_{k-1}^{\prime-j}, \tau_{k-1}^{-j} \in (\mathcal{T}_{k-1})^{N-1}.$$

B Induction step: Suppose that $\tilde{\lambda}_{k+1}^j$ does not vary with $(\tau_m^{-j}, \dots, \tau_{k-1}^{-j}) \in (\mathcal{T}_m)^{N-1} \times \dots \times (\mathcal{T}_{k-1})^{N-1}$ for some $1 < m \leq k-1$ (induction hypothesis B). Then show that $\tilde{\lambda}_{k+1}^j$ does not vary with $\tau_{m-1}^{-j} \in (\mathcal{T}_{m-1})^{N-1}$ neither. To see this, consider the belief updater $\iota_{m|k+1}^j(\tau_{k+1}^j)(\tau_{m-1}^{-j}|s, \tau_{m-2}^{-j})$, which is equal to (drawing on induction hypotheses A and B)

$$\begin{aligned} & \frac{\sum_{(\tau_m^{-j}, \dots, \tau_k^{-j})} \prod_{n=m-1}^k \pi_n(\tau_n|s, \tau_{n-1}) \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_k)}{\sum_{\tilde{\tau}_{m-1}^{-j}} \sum_{(\tau_m^{-j}, \dots, \tau_k^{-j})} \pi_{m-1}(\tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}|s, \tau_{m-2}) \prod_{n=m}^k \pi_n(\tau_n|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m, \dots, \tau_{n-1}) \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m, \dots, \tau_k)} \\ &= \frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-1}, \tau_m^j, \dots, \tau_k^j) \sum_{(\tau_m^{-j}, \dots, \tau_{k-1}^{-j})} \prod_{n=m}^{k-1} \pi_n(\tau_n|s, \tau_{n-1}) \pi_k^j(\tau_k^j|s, \tau_{k-1})}{\sum_{\tilde{\tau}_{m-1}^{-j}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m^j, \dots, \tau_k^j) \sum_{(\tau_m^{-j}, \dots, \tau_{k-1}^{-j})} \prod_{n=m}^{k-1} \pi_n(\tau_n|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m, \dots, \tau_{n-1}) \pi_k^j(\tau_k^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m, \dots, \tau_{k-1})} \\ &= \frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-1}, \tau_m^j, \dots, \tau_k^j) \prod_{n=m+1}^{k-1} \lambda_n^j(\tau_n^j|s, \tau_{n-1}^j) \pi_m^j(\tau_m^j|s, \tau_{m-1})}{\sum_{\tilde{\tau}_{m-1}^{-j}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m^j, \dots, \tau_k^j) \prod_{n=m+1}^{k-1} \lambda_n^j(\tau_n^j|s, \tau_{n-1}^j) \sum_{\tilde{\tau}_{m-1}^{-j}} \pi^j(\tau_m^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j})} \\ &= \frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-1}, \tau_m^j, \dots, \tau_k^j) \pi_m^j(\tau_m^j|s, \tau_{m-1})}{\sum_{\tilde{\tau}_{m-1}^{-j}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m^j, \dots, \tau_k^j) \pi_m^j(\tau_m^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j})}. \end{aligned}$$

In particular, for any two $\tau_{k+1}^j, \tau_{k+1}^{\prime j} : (\tau_k^j, \tau_{k+1}^j), (\tau_k^j, \tau_{k+1}^{\prime j}) \in \mathcal{T}_{k+1}^j$ believability requires that $\iota_{m|k+1}^j(\tau_k^j, \tau_{k+1}^j) = \iota_{m|k+1}^j(\tau_k^j, \tau_{k+1}^{\prime j})$. Whence

$$\frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-1}, \tau_m^j, \dots, \tau_k^j)}{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^{\prime j}|s, \tau_{m-1}, \tau_m^j, \dots, \tau_k^j)} = \frac{\sum_{\tilde{\tau}_{m-1}^{-j}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m^j, \dots, \tau_k^j) \pi_m^j(\tau_m^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j})}{\sum_{\tilde{\tau}_{m-1}^{-j}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^{\prime j}|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j}, \tau_m^j, \dots, \tau_k^j) \pi_m^j(\tau_m^j|s, \tau_{m-2}, \tau_{m-1}^j, \tilde{\tau}_{m-1}^{-j})}$$

for all τ_{m-1}^{-j} . As before, observe that the RHS does not depend on τ_{m-1}^{-j} so that the ratio on the LHS must be constant in τ_{m-1}^{-j} . And since $\tilde{\lambda}_{k+1}^j(\cdot|s, \tau_{m-1}, \tau_m^j, \dots, \tau_k^j)$ is a probability, it follows that $\tilde{\lambda}_{k+1}^j$ does not vary with τ_{m-1}^{-j} .

The induction argument B thus establishes that $\tilde{\lambda}_{k+1}^j$ does not depend on other players' beliefs at all. Therefore write $\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_k^j)$. It remains to show that it neither conditions on the state of nature s . This follows from considering the belief updater $\iota_{1|k+1}^j$. Invoking believability,

one obtains

$$\frac{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_k^j)}{\tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|s, \tau_k^j)} = \frac{\sum_{\tilde{s}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|\tilde{s}, \tau_k^j)p(\tilde{s})}{\sum_{\tilde{s}} \tilde{\lambda}_{k+1}^j(\tau_{k+1}^j|\tilde{s}, \tau_k^j)p(\tilde{s})} \quad \text{for all } s.$$

Then conclude with identical arguments as before.

Conversely, given a Bayes plausible λ_∞ , the recursive belief structure is given by

$$\begin{aligned} \pi_1^j(\tau_1^j|s) &= \tau_1^j(s)\lambda_1^j(\tau_1^j)/p(s) \\ \pi_k^j(\tau_k^j|s, \tau_{k-1}^j) &= \frac{\tau_k^j(\tau_{k-1}^{-j}|s, \tau_{k-2}^{-j})\lambda_k^j(\tau_k^j|\tau_{k-1}^j)\pi_{k-1}^j(\tau_{k-1}^j|s, \tau_{k-2}^j)}{\pi_{k-1}^j(\tau_{k-1}^j, \tau_{k-1}^{-j}|s, \tau_{k-2}^j)} \\ \pi_k(\tau_k^j, \tau_k^{-j}|s, \tau_{k-1}^j) &= \left[\sum_{\tilde{\tau}_{k+1}^j \in \mathcal{T}_{k+1}^j} \tilde{\tau}_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j})\lambda_{k+1}^j(\tilde{\tau}_{k+1}^j|s, \tau_k^j) \right] \pi_k^j(\tau_k^j|s, \tau_{k-1}^j). \end{aligned}$$

It remains to prove that the recursive belief structure is believable. First observe that

$$\begin{aligned} & \sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}^j, \tau_k^j, \tilde{\tau}_k^{-j}) \frac{\pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1}^j)}{\pi_k^j(\tau_k^j|s, \tau_{k-1}^j)} \\ &= \lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j)\pi_k^j(\tau_k^j|s, \tau_{k-1}^j) \sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \frac{\tau_{k+1}^j(\tilde{\tau}_k^{-j}|s, \tau_{k-1}^{-j})}{\pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1}^j)} \frac{\pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1}^j)}{\pi_k^j(\tau_k^j|s, \tau_{k-1}^j)} \\ &= \lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j) \end{aligned}$$

for arbitrary s and τ_{k-1}^{-j} . In particular, higher-order beliefs can not change the Bayesian update regarding lower-order beliefs, because the underlying distribution of higher-order beliefs does not condition on lower-order beliefs.

To conclude, prove that induced beliefs can be computed using Bayes rule. At order one the Bayesian belief updater gives

$$\iota_{1|1}^j(\tau_1^j) = \frac{p(s)\pi_1^j(\tau_1^j|s)}{\sum_{\tilde{s}} \pi_1^j(\tau_1^j|\tilde{s})p(\tilde{s})} = \frac{p(s)\tau_1^j(s)\lambda_1^j(\tau_1^j)/p(s)}{\lambda_1^j(\tau_1^j)} = \tau_1^j(s).$$

For higher order beliefs the Bayesian belief updater gives

$$\begin{aligned}
\ell_{k+1|k+1}^j(\tau_{k+1}^j)(\tau_k^{-j}|s, \tau_{k-1}^{-j}) &= \frac{\pi_k(\tau_k|s, \tau_{k-1})\pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k)}{\sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1})\pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k^j, \tilde{\tau}_k^{-j})} \\
&= \tau_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j}) \frac{\pi_k^j(\tau_k^j|s, \tau_{k-1})\lambda_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k^j)}{\sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1})\pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k^j, \tilde{\tau}_k^{-j})} \\
&= \tau_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j}) \frac{\pi_k^j(\tau_k^j|s, \tau_{k-1}) \sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k^j, \tilde{\tau}_k^{-j}) \frac{\pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1})}{\pi_k^j(\tau_k^j|s, \tau_{k-1})}}{\sum_{\tilde{\tau}_k^{-j} \in (\mathcal{T}_k)^{N-1}} \pi_k(\tau_k^j, \tilde{\tau}_k^{-j}|s, \tau_{k-1})\pi_{k+1}^j(\tau_{k+1}^j|s, \tau_{k-1}, \tau_k^j, \tilde{\tau}_k^{-j})} \\
&= \tau_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j}).
\end{aligned}$$

□

Proof of corollary 2. Proceed by induction: fix λ_2^j and use condition (3) to deduce π_1

$$\pi_1(\tau_1|s) = \tau_2^j(\tau_1^{-j}|s)\lambda_2^j(\tau_2^j|\tau_1^j)\pi_1^j(\tau_1^j|s) = \tau_2^j(\tau_1^{-j}|s)\lambda_2^j(\tau_2^j|\tau_1^j)\frac{\tau_1^j(s)}{p(s)}\lambda_1^j(\tau_1^j),$$

and λ_1^i for any $i \neq j$:

$$\begin{aligned}
\lambda_1^i(\tau_1^i) &= \sum_s \sum_{\tau_1^{-j}} \pi_1(\tau_1|s)p(s) \\
&= \sum_s \sum_{\tau_1^j} \sum_{\tau_2^j} \sum_{\tau_1^{-i \wedge j}} \tau_2^j(\tau_1^i, \tau_1^{-i \wedge j})\lambda_2^j(\tau_2^j|\tau_1^j)\pi_1^j(\tau_1^j|s)p(s) \\
&= \sum_{\tau_1^j} \sum_s \tau_1^j(s)\lambda_1^j(\tau_1^j) \sum_{\tau_2^j} \sum_{\tau_1^{-i \wedge j}} \tau_2^j(\tau_1^i, \tau_1^{-i \wedge j}|s)\lambda_2^j(\tau_2^j|\tau_1^j).
\end{aligned}$$

$(\tau_1^{-i \wedge j} \in (\mathcal{T}_1)^{N-2})$ refers to a profile of first-order beliefs which excludes players i and j .)

Induction step: fix $(\lambda_1^j, \dots, \lambda_{k+1}^j)$ and $(\pi_1, \dots, \pi_{k-1})$. Condition (3) implies that

$$\begin{aligned}
\pi_k(\tau_k|s, \tau_{k-1}) &= \sum_{\tau_{k+1}^j} \tau_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j})\lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j)\pi_k^j(\tau_k^j|s, \tau_{k-1}) \\
&= \sum_{\tau_{k+1}^j} \tau_{k+1}^j(\tau_k^{-j}|s, \tau_{k-1}^{-j})\lambda_{k+1}^j(\tau_{k+1}^j|\tau_k^j)\tau_k^j(\tau_{k-1}^{-j}|s, \tau_{k-2}^{-j})\lambda_k^j(\tau_k^j|\tau_{k-1}^j) \frac{\pi_{k-1}^j(\tau_{k-1}^j|s, \tau_{k-2})}{\pi_{k-1}(\tau_{k-1}|s, \tau_{k-2})}
\end{aligned}$$

which proves (4). To compute λ_k^i , note that due to corollary 1 and the preceding,

$$\begin{aligned}
\lambda_k^i(\tau_k^i | \tau_{k-1}^i) &= \sum_{\tilde{\tau}_{k-1}^{-i} \in (\mathcal{T}_{k-1})^{N-1}} \pi_k^i(\tau_k^i | s, \tau_{k-2}, \tau_{k-1}^i, \tilde{\tau}_{k-1}^{-i}) \frac{\pi_{k-1}(\tau_{k-1}^i, \tilde{\tau}_{k-1}^{-i} | s, \tau_{k-2})}{\pi_{k-1}^i(\tau_{k-1}^i | s, \tau_{k-2})} \\
&= \sum_{\tilde{\tau}_{k-1}^{-i}} \sum_{\tau_k^{-i}} \sum_{\tau_{k+1}^j} \tau_{k+1}^j(\tau_k^{-j} | s, \tau_{k-2}, \tau_{k-1}^i, \tilde{\tau}_{k-1}^{-i \wedge j}) \lambda_{k+1}^j(\tau_{k+1}^j | \tau_{k-2}, \tilde{\tau}_{k-1}^j, \tau_k^j) \pi_k^j(\tau_k^j | s, \tau_{k-2}, \tilde{\tau}_{k-1}^{-i}, \tau_{k-1}^i) \frac{\pi_{k-1}(\tau_{k-1}^i, \tilde{\tau}_{k-1}^{-i} | s, \tau_{k-2})}{\pi_{k-1}^i(\tau_{k-1}^i | s, \tau_{k-2})} \\
&= \sum_{\tilde{\tau}_{k-1}^{-i}} \sum_{\tau_k^{-i}} \sum_{\tau_{k+1}^j} \tau_{k+1}^j(\tau_k^{-j} | s, \tau_{k-2}, \tau_{k-1}^i, \tilde{\tau}_{k-1}^{-i \wedge j}) \lambda_{k+1}^j(\tau_{k+1}^j | \tau_{k-2}, \tilde{\tau}_{k-1}^j, \tau_k^j) \\
&\quad \tau_k^j(\tau_{k-1}^i, \tilde{\tau}_{k-1}^{-i \wedge j} | s, \tau_{k-2}) \lambda_k^j(\tau_k^j | \tau_{k-2}, \tilde{\tau}_{k-1}^j) \frac{\pi_{k-1}^j(\tilde{\tau}_{k-1}^j | s, \tau_{k-2})}{\pi_{k-1}^i(\tau_{k-1}^i | s, \tau_{k-2})}.
\end{aligned}$$

□

Bibliography

- Anderson, E. T. and Dana, J. (2009). When is price discrimination profitable? *Management Science*, 55(6):980–989.
- Arieli, I. and Babichenko, Y. (2019). Private bayesian persuasion. *Journal of Economic Theory*, 182:185 – 217.
- Arieli, I., Babichenko, Y., Sandomirskiy, F., and Tamuz, O. (2020). Feasible joint posterior beliefs.
- Arrow, K. J. (1965). *Aspects of the theory of risk-bearing*. Yrjö Jahnessonin Säätiö.
- Atakan, A. E. (2006). Assortative matching with explicit search costs. *Econometrica*, 74(3):667–680.
- Aumann, R. J. and Maschler, M. B. (1995). *Repeated Games with Incomplete Information*. MITPress. With the collaboration of R. E. Stearns.
- Becker, G. S. (1973). A theory of marriage: Part i. *Journal of Political economy*, 81(4):813–846.
- Benjamin, D. J., Brown, S. A., and Shapiro, J. M. (2013). Who is ‘behavioral’? cognitive ability and anomalous preferences. *Journal of the European Economic Association*, 11(6):1231–1255.
- Bergemann, D. and Morris, S. (2016a). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.
- Bergemann, D. and Morris, S. (2016b). Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review*, 106(5):586–91.
- Boldrin, M., Kiyotaki, N., and Wright, R. (1993). A dynamic equilibrium model of search, production, and exchange. *Journal of Economic Dynamics and Control*, 17(5-6):723–758.
- Bonneton, N. and Sandmann, C. (2019). Non-stationary search and assortative matching.
- Bonneton, N. and Sandmann, C. (2021). Probabilistic assortative matching under nash bargaining. Toulouse School of Economics.
- Brandenburger, A. and Dekel, E. (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1):189–198.

- Brocas, I. and Carrillo, J. D. (2007). Influence through ignorance. *The RAND Journal of Economics*, 38(4):931–947.
- Burdett, K. and Coles, M. G. (1997). Marriage and class. *The Quarterly Journal of Economics*, 112(1):141–168.
- Burdett, K. and Coles, M. G. (1998). Separation cycles. *Journal of Economic Dynamics and Control*, 22(7):1069–1090.
- Carlsson, H. and van Damme, E. (1993). Global Games and Equilibrium Selection. *Econometrica*, 61(5):989–1018.
- Chade, H., Eeckhout, J., and Smith, L. (2017). Sorting through search and matching models in economics. *Journal of Economic Literature*, 55(2):493–544.
- Chade, H. and Swinkels, J. (2020). The no-upward-crossing condition, comparative statics, and the moral-hazard problem. *Theoretical Economics*, 15(2):445–476.
- Chade, H. and Swinkels, J. M. (2019). The no-upward-crossing condition, comparative statics, and the moral-hazard problem. *Theoretical Economics*.
- Chiappori, P.-A. and Reny, P. J. (2016). Matching to share risk. *Theoretical Economics*, 11(1):227–251.
- Chistyakov, V. V. and Tretyachenko, Y. V. (2010a). Maps of several variables of finite total variation. i. mixed differences and the total variation. *Journal of Mathematical Analysis and Applications*, 370:672–686.
- Chistyakov, V. V. and Tretyachenko, Y. V. (2010b). Maps of several variables of finite total variation. ii. e. helly-type pointwise selection principles. *J. Math. Anal. Appl.*, 369:82–93.
- Choo, E. and Siow, A. (2006). Who marries whom and why. *Journal of political Economy*, 114(1):175–201.
- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259.
- Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review*, 100(3):1238–60.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3):522–550.
- Echenique, F. and Edlin, A. (2004). Mixed equilibria are unstable in games of strategic complements. *Journal of Economic Theory*, 118(1):61–79.
- Eeckhout, J. and Kircher, P. (2010). Sorting and decentralized price competition. *Econometrica*, 78(2):539–574.

- Eeckhout, J. and Lindenlaub, I. (2019). Unemployment cycles. *American Economic Journal: Macroeconomics*, 11(4):175–234.
- Ely, J. C. and Peski, M. (2006). Hierarchies of belief and interim rationalizability. *Theoretical Economics*, 1(1):19 – 65.
- Forges, F. (1993). Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310.
- Francis, C. (2013). *Functional Analysis, Calculus of Variations and Optimal Control*. Springer.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4):25–42.
- Gollier, C. (2004). *The Economics of Risk and Time*, volume 1. The MIT Press, 1 edition.
- Guiso, L. and Paiella, M. (2004). The role of risk aversion in predicting individual behaviors.
- Hagedorn, M., Law, T. H., and Manovskii, I. (2017). Identifying equilibrium models of labor market sorting. *Econometrica*, 85:29–65.
- Halac, M., Lipnowski, E., and Rappoport, D. (2021). Rank uncertainty in organizations. *American Economic Review*, 111(3):757–86.
- Harsanyi, J. (1967,1968). Games with incomplete information played by bayesian players, i-iii. *Management Science*, 14:159–182,320–334, 486–502.
- Hörner, J. and Skrzypacz, A. (2016). Selling information. *Journal of Political Economy*, 124(6):1515–1562.
- Idczak, D. and Walczak, S. (1994). On helly’s theorem for functions of several variables and its applications to variational problems. *Optimization*, 30(4):331–343.
- Kajii, A. and Morris, S. (1997). The robustness of equilibria to incomplete information. *Econometrica*, 65(6):1283–1309.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Palgrave Macmillan.
- Königsberger, K. (2004). *Analysis 2*. Springer-Verlag.
- Laffont, J.-J. and Martimort, D. (2002). *The Theory of Incentives: The Principal-Agent Model*. Princeton.
- Lauermann, S. and Nöldeke, G. (2015). Existence of steady-state equilibria in matching models with search frictions. *Economics Letters*, 131(C):1–4.

- Lauer mann, S., Nöldeke, G., and Tröger, T. (2020). The balance condition in search-and-matching models. *Econometrica*, 88(2):595–618.
- Legros, P. and Newman, A. F. (2007). Beauty is a beast, frog is a prince: Assortative matching with nontransferabilities. *Econometrica*, 75(4):1073–1102.
- Leonov, A. S. (1996). On the total variation for functions of several variables and a multidimensional analog of helly’s selection principle. *Mathematical Notes*, 63(1):61–71.
- Levy, G., Moreno de Barreda, I., and Razin, R. (2021). Feasible joint distributions of posteriors: A graphical approach.
- Li, F., Song, Y., and Zhao, M. (2019). Global manipulation by local obfuscation.
- Lise, J. and Robin, J.-M. (2017). The macrodynamics of sorting between workers and firms. *American Economic Review*, 107(4):1104–35.
- Liu, Q. (2009). On redundant types and bayesian formulation of incomplete information. *Journal of Economic Theory*, 144(5):2115 – 2145.
- Mariotti, T., Meier, M., and Piccione, M. (2005). Hierarchies of beliefs for compact possibility models. *Journal of Mathematical Economics*, 41:303–324.
- Maschler, M., Zamir, S., and Solan, E. (2013). *Game Theory*. Cambridge University Press.
- Maskin, E. and Diamond, P. (1979). An equilibrium analysis of search and breach of contracts, i: Steady states. *Bell Journal of Economics*, 10:282–316.
- Maskin, E. and Riley, J. (1984). Monopoly with incomplete information. *RAND Journal of Economics*, 15(2):171–196.
- Mathevet, L., Perego, J., and Taneva, I. (2019). On information design in games. *Journal of Political Economy*.
- McNamara, J. M. and Collins, E. (1990). The job search problem as an employer–candidate game. *Journal of Applied Probability*, 27(4):815–827.
- Mertens, J.-F. and Zamir, S. (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14(1):1–29.
- Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277.
- Morgan, P. B. (1994). *A model of search, coordination and market segmentation*. Department of Economics, State University of New York at Buffalo.
- Moriya, F. and Yamashita, T. (2020). Asymmetric-information allocation to avoid coordination failure. *Journal of Economics & Management Strategy*, 29(1):173–186.

- Morris, S. (2019). Presidential address econometric society. Coordination and Incomplete Information.
- Morris, S., Oyama, D., and Takahashi, S. (2020). Implementation via Information Design in Binary-Action Supermodular Games. Technical report.
- Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies*, 61(3):397–415.
- Munkres, J. (2015). *Topology*.
- Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317.
- Noussair, C. N., Trautmann, S. T., and Van de Kuilen, G. (2013). Higher order risk attitudes, demographics, and financial decisions. *Review of Economic Studies*, 81(1):325–355.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1-2):122–136.
- Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987.
- Rochet, J. (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200.
- Rubinstein, A. (1989). The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge. *American Economic Review*, 79(3):385–391.
- Sandmann, C. (2019). Price discrimination of non-rival goods. Toulouse School of Economics.
- Schulhofer-Wohl, S. (2006). Negative assortative matching of risk-averse agents with transferable expected utility. *Economics Letters*, 92(3):383–388.
- Shimer, R. and Smith, L. (2000). Assortative matching and search. *Econometrica*, 68(2):343–369.
- Shimer, R. and Smith, L. (2001). Nonstationary search. *University of Chicago and University of Michigan mimeo*, 112:122.
- Shmaya, E. and Yariv, L. (2016). Experiments on decisions under uncertainty: A theoretical framework. *American Economic Review*, 106(7):1775–1801.
- Smith, L. (2006). The marriage model with search frictions. *Journal of political Economy*, 114(6):1124–1144.
- Smith, L. (2011). Frictional matching models. *Annu. Rev. Econ.*, 3(1):319–338.
- Stokey, N. L. (1979). Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93(3):355–371.

- Strausz, R. (2006). Deterministic versus stochastic mechanisms in principal–agent models. *Journal of Economic Theory*, 128(1):306 – 314.
- Taneva, I. (2019). Information design. *American Economic Journal: Microeconomics*, 11(4):151–85.
- Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.
- Van Zandt, T. (2010). Interim Bayesian Nash equilibrium on universal type spaces for super-modular games. *Journal of Economic Theory*, 145(1):249–263.
- Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19(3):305 – 321.
- Ziegler, G. (2020). Bilateral adversarial information design.