

From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection*

Jean Tirole[†]

September 24, 2015

Abstract

In a number of interesting environments, dynamic screening involves positive selection: in contrast with Coasian dynamics, only the most motivated remain over time. The paper provides conditions under which the principal's commitment optimum is time consistent and uses this result to derive testable predictions under permanent or transient shocks. It also identifies environments in which time consistency does not hold despite positive selection, and yet simple equilibrium characterizations can be obtained.

Keywords: repeated relationships, screening, positive selection, time consistency, shifting preferences, exit games.

JEL numbers: D82, C72, D42

1 Introduction

The poll tax on non-Muslims that was levied from the Islamic conquest of then-Copt Egypt in 640 through 1856 led to the (irreversible) conversion of poor and least religious Copts to Islam to avoid the tax and to the shrinking of Copts to a better-off minority. To the reader familiar with Coasian dynamics, the fact that most conversions occurred during the first two centuries raises the question of why Muslims did not raise the poll tax over time to reflect the increasing average wealth and religiosity of the remaining Copt population.¹

This paper studies a new and simple class of dynamic screening games. In the standard intertemporal price discrimination (private values) model that has been the object of a voluminous literature, the monopolist moves *down* the demand curve: most eager customers buy or

*The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement no. 249429. The author is grateful to five referees, Daron Acemoglu, Dirk Bergemann, Daniel Garrett, Johannes Hörner, Alessandro Pavan, Di Pei, Jean-Charles Rochet, Mohamed Saleh and to participants at the 3rd CEPR Workshop on Incentives, Management and Organization (Frankfurt), the Jean-Jacques Laffont memorial seminar in Toulouse, the London Business School, the University of Zurich, the Spring Meeting of Young Economists in Ghent, the Nemmers lecture at Northwestern and the Mannheim European Economic Association meetings for helpful discussions and comments and to Arda Gitmez for excellent research assistance.

[†]Toulouse School of Economics (TSE) and Institute for Advanced Study in Toulouse (IAST)

¹A possible conjecture is that the Muslims wanted to preserve a tax base. However, they might also have wanted immediate income and further were also aiming at maximizing conversions to Islam. See Saleh (2013) for an analysis of the impact of the poll tax on the correlation of religious and socio-economic status over these twelve centuries.




Figure 1

consume first, resulting in a right-truncated distribution of valuations or “negative selection”. Similarly, in the dynamic version of Akerlof’s lemons model (common values), the buyer first deals with the most-eager-to-trade seller types –the owners of lemons– and thereafter raises price to account for the information that the seller is less eager to trade than expected. Again there is negative selection and the price setter moves down the demand curve (or rather up the supply curve).

In the poll tax example by contrast, the monopolist (the state) moves *up* the demand curve: Copts who most value their religion or are richer and therefore less sensitive to the pool tax remain in the tax base, while converts and their descendants, under the threat of apostasy forever disappear from that tax base. As we will later note, a number of interesting economic contexts share with the conversion game “positive selection” and, at least approximately, “absorbing exit”.

One might conjecture that the distinction between positive and negative selection is just a matter of sign convention, but this is not the case. To start building some intuition about why this is so, consider a demand curve $D(p) = 1 - F(p)$ obtained by aggregating demands from individual consumers with willingnesses to pay θ distributed according to some cumulative distribution $F(\theta)$. Suppose that the distribution for some reason has been truncated at θ_0 and let $D_0 = 1 - F(\theta_0)$; let η_R and η_L denote the elasticities of the right- and left-truncated residual demands. Then $\eta_R = [-D'(p)p]/[D(p) - D_0]$ (for $D(p) > D_0$) and $\eta_L = [-D'(p)p]/[D(p)]$ (for $D(p) < D_0$). That is, right truncations increase the elasticity of demand, while left ones leave it unchanged. This difference is illustrated in Figure 1, which emphasizes the reduction (invariance) of the set of inframarginal consumers under negative (positive) selection.

This observation has a number of implications.² The most obvious is that the monopoly price

²For instance, auctions of incentive contracts, which amount to a left truncation of the winner’s efficiency’s probability distribution, deliver for the winner the same power of incentive scheme as if the winner were a monopolist; optimal auctions thus only reduce the fixed component of rewards (Laffont and Tirole 1987). Another application of this property of the hazard rate is Niedermayer and Shneyerov (2014), in which a platform matches buyers and sellers. The mechanism that maximizes platform profit can as usual be derived by maximizing total virtual surplus. An interesting result in that paper is that the platform monopoly profit is also attainable in a decentralized manner, where the platform charges membership fees to the participants and then buyers and

is invariant to (moderate) left-, but not to right-truncations of the distribution of willingnesses to pay, a property whose implications for dynamic screening we will investigate. Indeed, this paper focuses on the properties and time consistency of monopoly pricing under left truncations. The dynamics first studied by Coase (1972) have the monopolist move down the demand curve once the cream has been skimmed off and the remaining market is the bottom of the barrel; the intuition is provided by the increasing elasticity of demand under right truncations. The monopolist’s incentive to reduce price over time has been shown to result in the time inconsistency of optimal monopoly pricing and therefore in an erosion of monopoly power. By contrast, the invariance of the elasticity under left truncations suggests that the monopolist will not be tempted to move up the demand curve and so optimal market segmentation is time-consistent.

Consider the following monopoly pricing example: agents – the consumers – have unit demand in each of two periods, $t = 0, 1$, and willingnesses to pay (relative to the outside opportunity) θ distributed on \mathbb{R} according to c.d.f. $F(\theta)$ and density $f(\theta)$ such that the static profit function for marginal cost c and unit mass of consumers, $(p - c)[1 - F(p)]$, is strictly quasi-concave. Let p^m denote the monopoly price ($p^m = c + [1 - F(p^m)]/f(p^m)$). Consumers can consume at date 1 only if they have consumed at date 0. If the principal (the monopolist) can commit at date 0 to a sequence of prices, the optimal price sequence is $p_0 = p_1 = p^m$. Because the consumers are the same at the two dates, the absorbing-exit constraint is not binding. Suppose now that the principal lacks commitment ability. Let the principal charge $p_0 = p^m$ at date 0, and suppose that consumers expect that the date-1 price will not be any lower. Consumers with type $\theta < p^m$ then do not want to consume at date 0 and they exit at that date. Consumers with type $\theta \geq p^m$ benefit from consuming at $t = 0$ regardless of their expectation of p_1 and therefore consume at date 0. Consider then date 1. The monopolist faces survival function $1 - F_1(\theta) = 1$ for $\theta \leq p^m$, and $= [1 - F(\theta)]/[1 - F(p^m)]$ for $\theta \geq p^m$ on its remaining goodwill and so picks price $p_1 = \arg \max_{\{p \geq p^m\}} \{(p - c)[1 - F(p)]/[1 - F(p^m)]\} = p^m$, vindicating consumers’ expectations. So the monopolist’s ability to implement monopoly pricing is not hampered by the lack of commitment.

This paper investigates several questions raised by this toy example. First, how general is the “weak time consistency” result that there exists *one* equilibrium that delivers the commitment payoff for the principal? Second, do *all* equilibria under non-commitment deliver the commitment payoffs? That is, are the commitment payoffs “strongly time consistent”? Third, can one characterize exit dynamics when the environment is uncertain and evolves over time (in the example above, the seller’s cost and the consumers’ willingnesses to pay might follow stochastic processes)? Fourth, does the analysis extend to heterogenous discount factors, to inflows of new agents and to finite re-entry costs? Finally, can exit dynamics still be characterized when time inconsistency prevails?

Section 2 describes the baseline model. Each period, the agent may remain in a relationship

sellers (independently of the platform) make take-it-or-leave-it offers to each other (random proposer model). Buyers with low values and sellers with high values do not join the platform because of the membership fees. For appropriate fees, the membership is the same as under centralization and decentralized trade is efficient: all buyers and sellers who pay the membership fees trade.

with the principal; exit is absorbing. Ignoring the transfers between the two parties, the agent’s flow payoff from the relationship (net of her outside option payoff) is increasing in a privately known type θ and may depend on the composition of the remaining installed base of peers as well as on the current state of nature s_t . Again ignoring transfers, the principal’s flow payoff depends on the state of nature, on the agent’s type and on the composition of the remaining installed base; the model thus accommodates private and common values. At date t , provided that the agent has remained in the installed base (the parties are still in a relationship), the principal offers a (positive or negative) price p_t which the agent accepts or refuses, in which case the game is over.

The paper’s first contribution is to show that provided that types (although not necessarily willingnesses to pay) are permanent, the optimal outcome is indeed time-consistent, and so no commitment ability is required to implement it (Section 3). More precisely, weak time consistency is always satisfied. Strong time consistency obtains if general mechanisms are allowed and there is a continuum of agents; alternatively, it obtains when the principal sequentially sets prices under the assumption of no strongly-positive network externalities, together with, for the case of an infinite horizon, the added requirement that either the equilibrium be Markov perfect or that the principal weakly benefit from a greater clientele at the optimal, non-frontloaded price (we later derive a sufficient condition for this property to obtain). The overall picture is that strong time consistency obtains very broadly with permanent types.³

When the agent’s and principal’s discount factors are allowed to differ, time consistency obtains if and only if the agent is more impatient than the principal. Furthermore, the principal’s payoff is always independent of the agent’s discount factor; this is in contrast with the bargaining (negative selection) literature in which agent impatience benefits the principal, who can then screen at a relatively low cost.

Time consistency transforms the search for a perfect Bayesian equilibrium of the no-commitment game into a simple dynamic stochastic optimization problem. We use this fact in Section 4 to compute the equilibrium in simple cases. One prominent case has an unambiguous aggregate evolution: the consumption becomes either more or less attractive over time; alternatively, the principal over time becomes more or less eager to retain agents. In the borderline case of an invariant environment, equilibrium leads to all exits (“conversions” in the religious example) taking place early on. More generally, when consumption deterministically becomes more attractive, all exit occurs at the initial date. By contrast, when it becomes less attractive over time (perhaps in a stochastic fashion), exit is spread over time and both the principal and the agent in equilibrium behave myopically, i.e., as if this were the last period (Section 4.1). Section 4.2 looks at non-monotone attractiveness and shows that familiar ironing techniques provide the

³Returning to the initial conversion game, this result sheds additional light on why the Muslim rulers did not sell non-expiring worshiping rights, instead of setting up a system with a recurring tax providing only a short-term worshiping right. Two reasons are specific to the context and do not apply to other contexts: First, the payment of an annual poll tax was dictated by an explicit verse in the Koran. Second, the principal in this context is also the contract enforcer (this may also apply to the emigration application discussed below), and so long-term contracts are not credible: A Copt could have paid this “once-and-for-all poll” tax, and the following year or decade been asked to pay again. The third reason is more universal and results from time consistency: Short-term contracting yields the principal’s optimum.

characterization of screening dynamics.

When the economy is subject to transient aggregate shocks (Section 4.3), which in a sense is the polar case of the permanent shocks studied in Section 4.1, shocks have long-lasting effects, the exit volume decreases over time and this volume is serially negatively correlated. At each date t , the participation depends only on the worst shock so far and is given by a simple condition.

Section 5 investigates the robustness of our results to finite re-entry costs and to inflows of new agents under the assumption of monotone attractiveness. Section 5.1 extends the model to finite re-entry costs in the context of monotone attractiveness and provides a lower bound on re-entry costs (equal to 0 in the case of decreasing/constant attractiveness) for re-entry to be irrelevant and therefore for the results obtained previously to apply.

The robustness of the insights to inflows of new agents is analyzed also for monotone-attractiveness environments in Section 5.2. Following the literature on negative selection, we look at whether the absence of price discrimination among identical cohorts impacts the outcome under commitment and non-commitment. Under negative selection, uniform pricing has the potential to restore some of the monopoly power that is eroded by the temptation to lower the price as the principal moves down the demand curve. One may wonder whether, conversely, the combination of the inflow of new cohorts and of uniform pricing might undermine the time consistency of optimal policies and thereby destroy principal value under positive selection.

For the class of monotone-attractiveness games, we obtain two interesting results. First, under commitment, uniform pricing does as well as discriminatory pricing. Second, time consistency obtains for decreasing/constant attractiveness, but not for strictly increasing attractiveness. In the case of decreasing/constant attractiveness, these results stem from the observation that the principal's optimal behavior under price discrimination is myopic and so identical across cohorts; therefore there is no cost for the principal of not being able to make use of cohort information. Neither is there any time consistency issue.

Under strictly increasing attractiveness, the optimal policy under price discrimination is described by invariant, but cohort-specific cut-offs; all exit within a cohort occurs in the cohort's first period of existence, but new cohorts have a lower cut-off, i.e., higher membership, than older ones. We show that under uniform pricing, a specific pattern of frontloaded payments is both necessary and sufficient to achieve the price-discrimination/cohort-specific cut-offs. This frontloading is what makes the optimal outcome time-inconsistent under uniform pricing (while it is time consistent under discriminatory pricing).

Section 6 by contrast studies environments in which time consistency does not hold. Solving for equilibria of principal-agent relationships in the absence of commitment is notoriously difficult as it no longer boils down to solving an optimization problem. Interestingly, though, the simple structure of games with positive selection allows us to provide equilibrium characterizations.

Section 6.1 looks at the possibility that the principal's preferences, and not only the publicly observable environment, change over time. It is well-known that optimal policies are not time consistent in such environments. With a shifting principal type, both the principal and the

agent must factor in the possibility that future principals be more or less eager to pursue the relationship than the current one. We prove two simple propositions. First, if the environment is invariant (except for the principal's type), the date- t remaining installed base is determined by its optimal level for the least eager principal so far. Second, whether attractiveness is monotonically increasing or decreasing, the installed base at any date t is smaller than what the current principal would induce if he and the agent anticipated that the principal's type would no longer change.

Section 6.2 studies the case of multiple principals. Under common agency, the agent's exit pattern is jointly determined by the principals. The analysis thus is a dynamic extension of the standard commons (or moral-hazard-in-team) problem. Despite positive selection, time consistency is invalidated by the principals' desire to influence each other's future policies. Nonetheless, equilibrium in the dynamic retention game can be characterized and shown to have the same qualitative properties as in the case of coordinated principals (i.e., the single-principal case).

We leave the study of environments in which time inconsistency arises on the agent's side for future research, and content ourselves with a few remarks. Section 6.3 allows the agent's type to shift over time. When the agent's type moves in an iid fashion, the equilibrium hazard rate for the termination of the relationship is constant and there is more exit than in the commitment solution. Myerson and Satterthwaite (1983) partnerships, in which exit by one member implies the dissolution of the partnership are briefly discussed in the conclusion and in an on-line Appendix.

Section 7 concludes with suggestions for future research. Omitted proofs can be found in the Appendix.

Economic environments with positive selection

Left truncations are closely related to the economics of incumbency; that is, they arise whenever an authority, firm or technology has an installed base of "customers" that forms a potential "tax base", but may irremediably exit:

Emigration: In an illustration closely related to the conversion game, suppose that emigration is an irrevocable decision (or at least one that is costly to reverse). Then an economically privileged, ethnic or religious group may see some of its members leave as the government levies more taxes or enacts adverse "non-price" policies toward the group. Over time, only the most attached to the country or the least mobile members of the group remain in the country.

Employee retention: A firm or an academic department at any given point in time comprises the subset of legacy employees who are the most committed or immobile. Organization-benefiting policies asking for public service, personal sacrifices or wage moderation create a risk for the organization of losing valuable employees.⁴

⁴Following Burdett and Mortensen (1998)'s work on the impact of on-the-job search on labor markets, a literature has developed that studies firms' retention policies under the threat of incoming outside opportunities. This literature in particular looks at how a steeply rising wage contract can approximate optimal "sell-out contracts" that are enabled by entry (or quitting) fees (Burdett and Coles 2003, Stevens 2004). The focus is rather different

Technology disadoption/licensing: A firm or group of intellectual property owners licensing key patents enabling the implementation of an incumbent technology may be concerned that users might defect for a new technological alternative. Again, the loyal tax base is composed of those with the highest benefits from the incumbent technology or highest switching costs.

Entry of generics: In a closely related example that has been the focus of much empirical literature, generic drugs that enter when a brand-name drug goes off-patent are typically far less expensive (by a factor of 5 or 10) than the brand-name version. Interestingly, the US retail prices of brand-name drugs tend to increase just prior to patent expiration and continue to increase (a bit less) post-patent expiration. This phenomenon, dubbed “price rigidity”, is consistent with the idea of positive selection. The brand-name drug manufacturer is left with the most loyal consumers and has little incentive to lower the price despite a declining market.⁵

Dyad game: Two parties are involved in some relationship. One of the parties is uncertain about the other party (friend, spouse, co-worker)’s commitment to the relationship and therefore does not know how much effort is needed to keep the relationship going. Assuming that the dyad once dissolved does not re-form, the dyad game is an illustration of our framework.⁶

Related literature

Following up on early work on dynamic “Myersonian” mechanism design (e.g., Baron and Besanko 1984, Courty and Li 2000), remarkable progress has recently been made to characterize optimal mechanisms under commitment. The literature has derived generalizations of the envelope characterization of the first-order condition and conditions under which attention can be focused on single, rather than compound deviations (Esö and Szentes 2007, Boleslavsky and Said 2013, and Pavan et al 2014). It has obtained necessary and sufficient conditions for the attainment of the optimal allocation either asymptotically (Battaglini 2005) or overall (Bergemann and Välimäki 2010, Athey and Segal 2013, Skrzypacz and Toikka 2013). Kruse-Strack (2014) studies the optimal commitment policy when the agent must choose a stopping time, the analog of our absorbing exit condition. The agent’s type evolves over time (making the optimum time inconsistent). Under a dynamic single crossing assumption, the paper provides an elegant characterization of optimal stopping rules and their implementation.

The work cited in the previous paragraph presumes “double commitment”: commitment to a long-term mechanism if the initial offer is accepted, and commitment by the principal not to make further offers if the initial offer is rejected. However, commitment is often not

from that in the paper since there is no asymmetric information at the contracting date and thus no dynamic screening issue; furthermore, commitment is assumed, while the present paper studies the time consistency of optimal commitment policies.

⁵In this example, “re-entry” (going back to the brand-name drug after switching to the generic) is relatively costless. However, as will be shown in Section 5.1, re-entry costs need not be large (and actually can be nil in the case of constant/decreasing attractiveness, a reasonable assumption for this application) for the results to hold. Note also that a more satisfactory study of the dynamics of generics entry would take into account specific institutional features such as automatic substitution requirements. Still the considerations developed here would be relevant in this richer model.

⁶For this to be the case, it is important that the informational asymmetry be one-sided. The case of two-sided asymmetric information is discussed in the on-line Appendix.

to be taken for granted. Public policies generally lack commitment, and so do a number of policies in private-sector environments, especially when future policies are hard to describe, let alone contract upon, in advance. More precisely, long-term commitment may be infeasible (“no commitment”); or it may be feasible but renegotiable (i.e., it is renegotiated if the parties involved in the long-term contract all find it advantageous to do so; this is the paradigm of “commitment and renegotiation” studied e.g., in Dewatripont 1989 or Laffont and Tirole 1990); finally, further offers cannot be precluded, as in Coase’s durable good model.

Much less is known for environments in which full commitment is impossible. In specific environments (usually two permanent types), both the commitment-and-renegotiation model and the no-commitment model (without agent anonymity) have been shown to exhibit Coasian dynamics (Hart and Tirole 1988, Maestri 2013, Strulovici 2013). Similar Coasian dynamics have been obtained in the common value counterpart of Coase’s model, the dynamic version of Akerlof’s lemons model (e.g. Daley and Green 2012, Fuchs and Skrzypacz 2013, Gerardi and Maestri 2013, Gerardi et al 2014).

Little is known either regarding general properties of sequential screening in such environments. Bester and Strausz (2001) demonstrate that the cardinality of messages can be confined to that of types for a finite type space. Skreta (2006) looks at the standard risk-neutral seller-buyer game, in which, say, the buyer’s invariant valuation θ is drawn from some distribution with support the interval $[0, 1]$, and the seller’s cost is 0. As long as the buyer is not served, the seller keeps making offers at $t = 0, \dots, T$. These offers need not be prices p_t . Rather they can be full-fledged mechanisms, resulting in date- t probabilities of trade and expected transfers. The central result is that an optimal mechanism is to simply post a price in each period, generalizing the Riley and Zeckhauser (1983) classic result to sequential mechanism design and thereby simplifying the search for equilibria in this class of games.

The literature so far has been concerned with a principal selling goods in a market. The principal’s objective is then to attract consumers without lowering price too much. By contrast, we consider a principal who is trying to retain a customer base while incurring a low cost or charging a high price. In this sense, this model is a mirror image of existing screening models, with the less motivated jumping off ship instead of the most motivating getting on board.

Board and Pycia (2014) study pricing by a durable-good monopolist when the consumers can enjoy an alternative, outside option with positive value. In this framework it is easy to show that the monopoly price (the price p^m that maximizes $[\theta - c][1 - F(\theta)]$, where $F(\theta)$ is the distribution of willingnesses to pay and c is the marginal cost) is time consistent: the monopolist charges at date 0 the monopoly price, those consumers with $\theta \geq p^m$ purchase immediately and the others opt for the outside option. The principal’s (out of equilibrium) beliefs from date 1 on are that remaining consumers have types above p^m . Board and Pycia’s striking result is that this is the only equilibrium.⁷ Our framework differs from theirs in three important and related aspects. First, Board and Pycia’s principal can commit to a long-term contract as is implicit

⁷In Fudenberg et al (1987), the possibility for the seller to consume the good himself or to switch to bargaining with another buyer can restore commitment power but there are multiple equilibria. In Board-Pycia, the outside option is on the buyer side, commitment is fully restored and the equilibrium is unique.

in the durable good framework (selling can be viewed as a commitment to long-term rental). The principal’s commitment assumption implies that the agent’s discount factor affects the principal’s welfare like in pure negative selection models and unlike the pure positive selection model of this paper. Second, their model exhibits both left- and right-truncations (the more motivated buy the durable good and the least motivated buy the outside option), whereas the model studied in this paper has only left truncations and Coase’s traditional model has only right truncations. Third, exit in our model is mechanically induced by the non-continuation/breakup of the relationship, leading to an infinite or finite re-entry cost; in Board-Pycia, non-consumption need not imply exit, and there is a third potential status (“neither in nor out”). This third status is not observed on the equilibrium path of the deterministic world of Board-Pycia, but would become relevant if, as in this paper, stochastic shocks were to affect the attractiveness of the relationship or principal and agent’s intrinsic preferences were to change over time. It would actually be interesting to extend the Board-Pycia analysis to stochastic environments.

2 Model

Time is discrete: $t = 0, 1, \dots, T$, where T is finite or infinite. At the beginning of each period, a state of nature s_t is realized in some set S_t . Let $s^t \equiv (s_0, \dots, s_t) \in S^t \equiv S_0 \times \dots \times S_t$ follow a stochastic process with conditional distribution $G(s^\tau | s^t)$ for $\tau > t$. We will say that $s^\tau \succ s^t$ for $\tau > t$ if there exists (s_{t+1}, \dots, s_τ) such that $s^\tau = (s^t, s_{t+1}, \dots, s_\tau)$.

The players are a principal, who has the bargaining power, and either one agent or a continuum of agents with mass 1, in both cases with a unit demand in each period. Unless otherwise specified (i.e. in Proposition 1(i) and Proposition 6, all results apply to both a single and a continuum of agents). All players have identical discount factor $\delta \in (0, 1)$.

The agent is characterized by a privately-known type $\theta \in [\underline{\theta}, \bar{\theta}]$, distributed according to smooth c.d.f. $F(\theta)$ with density $f(\theta)$. Each period t , the agent consumes ($x_t = 1$) or does not consume ($x_t = 0$); exit is absorbing and so her consumption decision however is relevant only if she has kept consuming in the past. Let $X^t \equiv \prod_{z=0}^t x_z$.

In the non-commitment version of the game, the principal makes an offer in each period, which, together with the agent’s response to the offer, determines the allocation in that period. We will consider two versions, depending on the nature of the offer:

Price-offer version. At each date t , the principal offers a price p_t , and the agent chooses whether to consume ($x_t = 1$) or not ($x_t = 0$).

Mechanism-offer version. More generally, the principal at date t can, as in Skreta (2006), offer a mechanism that determines the date- t allocation. With a single agent, a mechanism is a message space M_t , and for each message $m_t \in M_t$, a transfer $p_t(m_t)$ and a probability $x_t(m_t)$ of date- t consumption. The mechanism also specifies what the principal learns at the end of the period concerning the agent’s message.⁸ With a continuum agents, the allocation $\{p_t^i, x_t^i\}$ of agent i

⁸That is, any coarsening of the message can be transmitted to the principal. To prove strong time consistency below, it will actually suffice to consider mechanisms such that the principal only learns that the agent has

may depend on her message m_t^i , but also on the other agents' date- t messages; in particular, in the continuum-of-agents version of the model, we will allow the principal to make the allocation contingent on the number of agents wanting to keep consuming in the period (this will be used only in part (i) of Proposition 1).

As we will show, the principal gains little traction from using general mechanisms, and so, in the following description of the full game, we will focus on the price-offer version. The mechanism-offer version follows straightforwardly from this description.

Information, timing and strategies. At the beginning of each period t , the state s_t is realized and publicly observed. In (the no-commitment version of) the game, the principal sets a price p_t for date- t membership/consumption, and previously loyal consumers (those for whom $X^{t-1} = 1$) decide whether to consume. Strategies $\{\sigma^P, \sigma^A\}$ are price choices, $p_t = \sigma_t^P(p^{t-1}, s^t) \in \mathbb{R}$ for the principal and consumption choices $x_t = \sigma_t^A(p^t, s^t, \theta) \in \{0, 1\}$ for the agent, where $p^t \equiv (p_0, \dots, p_t)$. We focus on pure strategies and the equilibrium concept is perfect Bayesian equilibrium.

This description of strategies assumes that the price (or mechanism) offered by the principal at date t depends only on the state s^t (and previous prices, and implicitly on the fact that the agent is still in the relationship). This is a natural assumption in the private values case. However, under common values, the principal's flow payoff depends on the agent's type. The measurability assumption can then be justified in one of two ways: The principal may observe only the aggregate performance of the agents; alternatively, the principal is prohibited from discriminating among agents. If none of the three conditions holds (that is, if values are common, individual performance is observable and discrimination is feasible), the outcomes under commitment and non-commitment differ from those described below.⁹

Agents' preferences. Relative to the payoff obtained by not consuming¹⁰, the agent's net surplus from date- t consumption is linear in the date- t transfer $p_t \in \mathbb{R}$ to the principal (a price, or more generally the conditions demanded by the principal for belonging to the consuming group¹¹); his gross surplus from consumption depends on his type θ , on the date- t payoff relevant state s_t and on the set $\Theta_t \subseteq [\underline{\theta}, \bar{\theta}]$ of types who consume at date t .

The dependence of preferences on Θ_t allows for social image/self views and (in the case of a continuum of agents) network externalities to affect consumption decisions. For example, the agent's utility may depend, positively or negatively, on the mass $\mu(\Theta_t)$ of agents in the consumption group; more generally, network externalities may also depend on the identity of members of that group. Allowing for externalities adds an argument in the surplus function, consumed at date t (or learns nothing at all since absorbing exit implies that remaining agents have consumed in the past).

⁹On the other hand, if θ is observable at the end of the period, the commitment outcome is still time consistent: In the absence of commitment, the allocation from date 1 on is efficient conditionally on cutoff θ_0^* . So the issue of time consistency does not arise.

¹⁰The agent's utility is thus defined net of utility upon exit. This normalization of utility upon exit to 0 is without loss of generality as it allows the value of exit to differ for different types.

¹¹The quasi-linearity of preferences is assumed solely for expositional simplicity. Similarly, transfers more generally can involve deadweight losses. The key assumption is positive selection.

but given their relevance in a number of applications, it is worthwhile to investigate whether the results hold when they are present. We will use the “network externalities” terminology whether the dependence on Θ_t arises with a continuum of agents or with a single agent (image concern interpretation).

Skimming property. We assume that higher types have a strictly higher gross surplus for all (Θ_t, s_t) . It is then straightforward¹² to show that, conditionally on having consumed up to date t , if type θ consumes at date t for history (p^t, s^t) (i.e., $x_t(p^t, s^t, \theta) = 1$), then so does type $\theta' > \theta$ (i.e., $x_t(p^t, s^t, \theta') = 1$). Intuitively, this results from the fact that type θ' obtains a strictly higher utility from consumption at date t and that the agent’s continuation valuation at date $(t + 1)$ is weakly increasing in type (as type θ' can always mimic type θ ’s behavior from date $(t + 1)$ on). Thus, incentive compatibility implies the existence of a unique cut-off $\theta_t^*(p^t, s^t)$ such that $\Theta_t = [\theta_t^*(p^t, s^t), \bar{\theta}]$ and $\mu(\Theta_t) = 1 - F(\theta_t^*(p^t, s^t))$. Absorbing exit then implies cutoff monotonicity: $\theta_t^*(p^t, s^t) \geq \theta_{t-1}^*(p^{t-1}, s^{t-1})$.

We can therefore write the agent’s net payoff function as a function of the cutoff:

$$\phi(\theta, \theta_t^*, s_t) - p_t.$$

ϕ is assumed to be strictly increasing in its first argument and differentiable in its first two arguments. The intertemporal utility of a type- θ agent is

$$E_{sT} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) [\phi(\theta, \theta_t^*, s_t) - p_t] \right].$$

Principal’s preferences. The principal also has quasi-linear preferences, with flow payoff

$$\int_{\theta_t^*}^{\bar{\theta}} \psi(\theta, \theta_t^*, s_t) f(\theta) d\theta + p_t [1 - F(\theta_t^*)],$$

and intertemporal utility

$$E_{sT} \left[\sum_{t=0}^{t=T} \delta^t \left[\int_{\theta_t^*}^{\bar{\theta}} \psi(\theta, \theta_t^*, s_t) f(\theta) d\theta + p_t [1 - F(\theta_t^*)] \right] \right].$$

The principal’s objective function deserves some comment as well. Often, the economic model defines the ψ function directly. For example, $\psi(\theta, \theta_t^*, s_t)$ could stand for the marginal productivity of worker θ enjoying (positive or negative) production externalities depending on the set of coworkers $[\theta_t^*, \bar{\theta}]$ in state s_t . Our formalism thus allows for common values: the principal may care about the agent’s type. For example, loyal employees may be loyal because they are enthusiastic about their job and then are highly productive; or they may stay because they are unable to find another job and then are likely to have a low productivity for the firm. Similarly, the principal in the dyad game may exhibit reciprocal altruism and then experience a welfare that depends on (his perception of) θ . Both cases are illustrations of common values.

¹²The proof follows the standard lines (see, e.g., Fudenberg et al 1985).

Sometimes, the economic model gives instead the principal's overall (i.e. summed over $[\theta^*, \bar{\theta}]$) gross surplus directly; one must then define the ψ function accordingly. For example, the cost function of serving a number $1 - F(\theta_t^*)$ of agents might be $C(1 - F(\theta_t^*), s_t)$. One can then define $\psi(\theta, \theta_t^*, s_t) = -C_1(1 - F(\theta), s_t)$ for all θ , where C_1 is the derivative with respect to the first argument, and so $\int_{\theta_t^*}^{\bar{\theta}} \psi(\theta, \theta_t^*, s_t) f(\theta) d\theta \equiv C(0, s_t) - C(1 - F(\theta_t^*), s_t)$. This latter example illustrates the possibility of “network externalities” arising on the principal/cost side. The notion of “externality augmented virtual surplus” introduced below is therefore relevant under non-constant returns even when there are no direct externalities among agents.

Examples. The model embodies the premises of dynamic screening with left truncations: the absorbing-exit condition and non-commitment. Let us provide a few examples.

In the basic *conversion game*, $\phi(\theta, \theta_t^*, s_t) = \theta$ and $\psi(\theta, \theta_t^*, s_t) = -c$. The agent has preferences $X^t[\theta - p_t]$ where θ is the ratio of the agent's religiosity over her marginal utility of income; needless to say, we could enrich this basic set up with in- and out-religious group externalities. The principal is the Muslim rulers, with overall date- t instantaneous payoff $c[1 - \mu(\Theta_t)] + p_t\mu(\Theta_t)$; the parameter c reflects the rulers' intrinsic preference for conversion to Islam, leading to a “markup” $p_t - c$ on poll-tax-paying Copts.

The same payoffs can be used to describe the *dyad game*, where p_t represents (minus) the effort exerted to keep the uncommitted party on board.

With only very slight modifications, the model also accommodates persecutions such as those brought about by the inquisition (against the Albigensian heresy in the 13th century by the Dominicans on behalf of the Pope or against Spanish Jews and Moslems in the late 15th century Spain by Queen Isabella and the Tribunal of the Holy Office of the Inquisition). The screening instrument employed is then purely wasteful, except perhaps for the confiscations, but the results in this paper do not rely on the tax being a pure transfer between the principal and the agent.¹³

In the *technology-disadoption game*, $\phi(\theta, \theta_t^*, s_t) = \theta + \alpha[1 - F(\theta_t^*)] + s_t$ and $\psi(\theta, \theta_t^*, s_t) = -c$, where α is a network-externality coefficient and s_t might stand for shifts in the attractiveness of the challenging technology. The agent has flow preferences $X^t[\theta + \alpha\mu(\Theta_t) + s_t - p_t]$. The principal's flow profit is the $(p_t - c)\mu(\Theta_t)$.

¹³Let the Catholic rulers' objective function at date t be

$$-c[1 - F(\theta_t^*)] - K(i_t)$$

where c is their disutility of non-conversion, i_t is the intensity of inquisition (the probability of detecting non-converts) and K is an increasing and convex cost function. Let the utility of a non-convert with religiosity θ be $\theta - i_t d$ where d is the relative cost of being caught and tried. The remaining installed based at date t is

$$Y^{t-1}[1 - F(\theta_{t-1}^*)]$$

where $Y^{t-1} = (1 - i_0) \cdots (1 - i_{t-1})$. Thus the elasticity at any θ remains the same under left truncations, and the analysis carries over. In particular, varying the rulers' religiosity (c) over time or the impact of the environment (for example, through K), one can as in Sections 4 and 6.1 derive dynamics of inquisition intensity.

3 Time consistency

3.1 Optimal mechanism under commitment

Suppose, first, that the principal can commit to an incentive compatible mechanism that specifies for each θ a (present-discounted) payment $P(\theta)$ and a state-contingent consumption pattern $\{x_t(\theta, s^t)\}_{t \in \{0, \dots, T\}}$ (such that $x_t(\theta, s^t) = 0 \Rightarrow x_{t+1}(\theta, s^{t+1}) = 0$ if $s^{t+1} \succ s^t$). Letting $s^t \in S^t$

$$U(\theta) \equiv \max_{\{\tilde{\theta} \in [\underline{\theta}, \bar{\theta}]\}} \left\{ E_{s^T} \left[\sum_{t=0}^{t=T} \delta^t X^t(\tilde{\theta}, s^t) \phi(\theta, \theta_t^*, s_t) \right] - P(\tilde{\theta}) \right\},$$

the participation and incentive constraints require that

$$U(\underline{\theta}) \geq 0$$

and

$$\frac{dU}{d\theta} = E_{s^T} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) \frac{\partial \phi}{\partial \theta}(\theta, \theta_t^*, s_t) \right].$$

Consider an optimal policy under commitment. Let $U(\theta) \geq 0$ denote the ex-ante rent of type θ , and V denote the principal's ex-ante payoff for an arbitrary mechanism. Using the standard decomposition between efficiency and rent, the principal's payoff can be written as:

$$V = E_{\theta} E_{s^T} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) [\phi(\theta, \theta_t^*(s^t), s_t) + \psi(\theta, \theta_t^*(s^t), s_t)] - U(\theta) \right],$$

where (using $U(\underline{\theta}) = 0$)

$$\begin{aligned} E_{\theta}[U(\theta)] &= \int_{\underline{\theta}}^{\bar{\theta}} U(\theta) dF(\theta) = \int_{\underline{\theta}}^{\bar{\theta}} \frac{dU(\theta)}{d\theta} [1 - F(\theta)] d\theta \\ &= E_{s^T} \int_{\underline{\theta}}^{\bar{\theta}} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) \frac{\partial \phi}{\partial \theta}(\theta, \theta_t^*(s^t), s_t) \right] [1 - F(\theta)] d\theta. \end{aligned}$$

And so the principal's payoff can be rewritten in the standard, expected virtual surplus fashion:

$$\begin{aligned} V &= \int_{\underline{\theta}}^{\bar{\theta}} E_{s^T} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) \left[[\phi(\theta, \theta_t^*(s^t), s_t) + \psi(\theta, \theta_t^*(s^t), s_t)] f(\theta) \right. \right. \\ &\quad \left. \left. - \frac{\partial \phi}{\partial \theta}(\theta, \theta_t^*(s^t), s_t) [1 - F(\theta)] \right] d\theta \right] \\ &= \int_{\underline{\theta}}^{\bar{\theta}} E_{s^T} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) \Gamma(\theta, \theta_t^*(s^t), s_t) \right] f(\theta) d\theta \end{aligned}$$

where

$$\Gamma(\theta, \theta^*, s) \equiv \phi(\theta, \theta^*, s) + \psi(\theta, \theta^*, s) - \frac{\partial \phi}{\partial \theta}(\theta, \theta^*, s) \frac{1 - F(\theta)}{f(\theta)}$$

denotes the virtual surplus. Maximizing the principal's payoff amounts to solving the following program:

$$\sup_{\{\theta^*(\cdot)\}} \left\{ E_{s^T} \left[\sum_{t=0}^{T-1} \delta^t W(\theta_t^*(s^t), s_t) \right] \right\} \quad (\text{I})$$

subject to

$$\theta_t^*(s^t) \geq \theta_{t-1}^*(s^{t-1})$$

where

$$W(\theta_t^*(s^t), s_t) \equiv \int_{\theta_t^*(s^t)}^{\bar{\theta}} \Gamma(\theta, \theta_t^*(s^t), s_t) f(\theta) d\theta.$$

The optimization indeed must respect the feasibility constraints (F)¹⁴: For all (t, s^t) ,

$$X^t(\theta, s^t) = 1 \iff \theta \geq \theta_t^*(s^t),$$

or, equivalently, *cutoff monotonicity*:

$$\theta_t^*(s^t) \geq \theta_{t-1}^*(s^{t-1}) \quad \text{if } s^t \succ s^{t-1}. \quad (\text{F})$$

As usual, the policy must be optimal for any subform; that is, for all (t, s^t) , $\{X^\tau(\cdot, \cdot), \theta_\tau^*(\cdot)\}_{\tau \geq t}$ must also maximize :

$$V_t(s^t) = \int_{\theta_{t-1}^*(s^{t-1})}^{\bar{\theta}} E_{s^T | s^t} \left[\sum_{\tau=t}^{T-1} \delta^{\tau-t} X^\tau(\theta, s^\tau) \Gamma(\theta, \theta_\tau^*(s^\tau), s_\tau) \right] \left[\frac{f(\theta)}{1 - F(\theta_{t-1}^*(s^{t-1}))} \right] d\theta$$

subject to the relevant set of feasibility constraints for all (τ, s^τ) such that $\tau \geq t$ and $s^\tau \succeq s^t$.

We let $\{\hat{\theta}_t(s^t)\}_{t=0, \dots, T}$ denote the optimal contingent cutoff sequence under commitment.¹⁵
 $s^t \in S^t$

Next, we show that there is no loss of generality (in the sense that the commitment payoff can be attained—this says nothing about outcome uniqueness) in considering commitments to a sequence of state-contingent prices that the agent accepts or turns down. This observation has two implications: First, restricting the principal to price offers rather than general mechanisms still allows us to compute the commitment welfare. Second, this commitment welfare can be

¹⁴Because the cutoffs are weakly increasing, this condition need only be checked for the last cutoff. Note furthermore that the condition “if $s^t \succ s^{t-1}$ ” in condition (F) can be dispensed with (it does not matter what the cutoff is if s^t is unfeasible given s^{t-1}).

¹⁵We will assume but not investigate the existence of an optimal policy. The optimization boils down to a search for a plan specifying state-contingent cutoffs $\{\theta_t^*(s^t)\}$ so as to solve Program (I). The state of the system at date t is $(s^t, \theta_{t-1}^*(s^{t-1}))$. One can then apply standard results in dynamic programming as stated, say, in Lucas et al (1989).

attained even if the agent is unable to commit. To show this, consider a contingent price sequence

$$\mathbf{p} \equiv \{p_t(s^t)\}_{\substack{t \in \{0, \dots, T\} \\ s^t \in S^t}}$$

so as to implement an arbitrary sequence of contingent, weakly increasing cutoffs $\boldsymbol{\theta}^* \equiv \{\theta_t^*(s^t)\}_{\substack{t \in \{0, \dots, T\} \\ s^t \in S^t}}$.

Cutoffs must satisfy sequential incentive compatibility. Introducing the agent's value function:

$$U_t(\theta, s^t; \mathbf{p}, \boldsymbol{\theta}^*) \equiv \max \left\{ 0, \phi(\theta, \theta_t^*(s^t), s_t) - p_t(s^t) + \delta E[U_{t+1}(\theta, s^{t+1}; \mathbf{p}, \boldsymbol{\theta}^*)] \right\},$$

then

$$x_t(s^t, \theta; \mathbf{p}, \boldsymbol{\theta}^*) = 1 \text{ if and only if } \phi(\theta, \theta_t^*(s^t), s_t) - p_t(s^t) + \delta E[U_{t+1}(\theta, s^{t+1}; \mathbf{p}, \boldsymbol{\theta}^*)] \geq 0. \quad (\text{IC})$$

The principal's commitment payoff in the price-offer game is:

$$\widehat{V} \equiv \max_{\{\mathbf{p}, \boldsymbol{\theta}^* \text{ satisfying (IC)}\}} E_{s^T} \left[\sum_{t=0}^{T-1} \delta^t \left[\int_{\theta_t^*(s^t)}^{\bar{\theta}} \psi(\theta, \theta_t^*(s^t), s_t) f(\theta) d\theta + p_t[1 - F(\theta_t^*)] \right] \right].$$

Note that (for the sake of the definition of \widehat{V} only), we let the principal maximize not only over prices, but also over cutoffs. Indeed, there is no guarantee that a price strategy \mathbf{p} leads to a unique sequence of cutoffs $\boldsymbol{\theta}^*$. Our allowing for network externalities implies a possible multiplicity of static equilibria if externalities are positive and strong; under Assumption 1 below, though, the principal can guarantee himself \widehat{V} by choosing \mathbf{p} only.

Assumption 1 (no strongly positive network externalities). *For all s , the function $\phi(\theta^*, \theta^*, s)$ is strictly increasing in θ^* (i.e., $\phi_1(\theta^*, \theta^*, s) + \phi_2(\theta^*, \theta^*, s) > 0$).*

Assumption 1 is satisfied whenever network externalities are negative or non-existent ($\phi_2 \geq 0$). It is also satisfied for positive network externalities ($\phi_2 < 0$) provided they are not too large. For example, the technology-disadoption game ($\phi(\theta^*, \theta^*, s) = \theta^* + \alpha[1 - F(\theta^*)] + s$) satisfies Assumption 1 provided that $1 - \alpha \sup \{f(\theta)\} > 0$. Assumption 1 prevents multiple equilibria in the static game; in its absence, a “wrong coordination” of agents by itself might induce a payoff for the principal that lies below \widehat{V} . Hence Assumption 1 is a necessary condition for equilibrium uniqueness in the price-offer game; its necessity for the implementation of \widehat{V} is therefore unrelated to the time-consistency issue. In the single-period case for instance, the optimal mechanism $\{x(\theta) = 1 \text{ iff } \theta \geq \theta^*, p(\theta) = \phi(\theta^*, \theta^*) \text{ if } \theta \geq \theta^* \text{ and } = 0 \text{ otherwise}\}$ gives rise to multiple equilibria if there exists θ^{**} such that $\phi(\theta^{**}, \theta^{**}) = \phi(\theta^*, \theta^*)$.

Lemma 1 (irrelevance of agent commitment). *Consider a commitment allocation $\{P(\cdot), x(\cdot, \cdot)\}$ with associated (monotonic) cutoffs $\widehat{\theta}(\cdot)$ and satisfying $U(\underline{\theta}) = 0$. Then the sequence of short-term prices $\mathbf{p} \equiv \{p_t(s^t)\}_{\substack{t \in \{0, \dots, T\} \\ s^t \in S^t}}$ defined by the “cutoff-myopia” property:*

$$\forall s^t: \quad p_t(s^t) = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$$

is such that, if the principal commits to the sequence of state-contingent prices \mathbf{p} , (i) there is an equilibrium that yields the same payoff for the principal and (for all θ) the agent as the commitment allocation; (ii) if Assumption 1 holds, the equilibrium is unique and therefore delivers the commitment allocation.

In this sense, there is no need for commitment by the agent. And the principal does not lose from offering prices.

As we noted in the introduction, while the optimal commitment *outcome* is time consistent, not every optimal commitment *policy* is. Prices can be arbitrarily frontloaded,¹⁶ with the impact of high initial prices being offset by the promise of low prices in the future; however, frontloaded policies in general are time inconsistent as the principal would want to renege on this promise.

3.2 Non-commitment and time consistency

In practice principals may find it difficult to commit to a long-term, state-contingent policy. The absence of commitment is particularly natural either when the principal is a government or when specifying “non-price” dimensions of the future relationship in a contract is complex. This raises the time-consistency issue.

Definition (time consistency). *Weak* time consistency holds if there exists a perfect Bayesian equilibrium of the non-commitment game that delivers expected payoff \widehat{V} for the principal. *Strong* time consistency holds if all perfect Bayesian equilibria of the non-commitment game deliver payoff \widehat{V} for the principal.

For part (iii) of the proposition, we will further either focus on Markov perfect equilibria or make the following assumption:

Assumption 2 (static benefits of a greater clientele). For all (t, s^t) ,

$$\frac{\partial}{\partial \theta^*} \left(\int_{\theta^*}^{\bar{\theta}} \psi(\theta, \theta^*, s_t) f(\theta) d\theta + p_t [1 - F(\theta^*)] \right) \leq 0$$

at any $\theta^* \leq \widehat{\theta}_t(s^t)$ and for $p_t = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$, where $\{\widehat{\theta}_t(s^t)\}$ denotes the sequence of cutoffs along some optimal commitment policy.

Assumption 2 says that the principal does not mind having a greater clientele provided that the price is set at the valuation of the current cutoff’s surplus at the optimal program. In Section 4, we will provide a sufficient condition for Assumption 2 to be satisfied. For instance,

¹⁶A commitment policy $(\mathbf{p}, \boldsymbol{\theta}^*)$ is frontloaded if there exists (t, s^t) such that

$$p_t(s^t) > \phi(\theta_t^*(s^t), \theta_t^*(s^t), s_t).$$

That is, the cutoff type at date t in state s^t must expect some strictly positive surplus (i.e. a price below his gross surplus) at some future date τ in some state s^τ .

it is satisfied by the conversion and technology disadoption games, provided that attractiveness is constant or decreasing as defined in the next section.¹⁷

A *Markov perfect equilibrium* is an equilibrium in which the players' strategies at date t depends only on the previous cutoff θ_{t-1}^* and the part of the state that is a sufficient statistic for the Markov process (i.e., s_t if s follows a first-order Markov process), and, for the agent, on p_t as well.

Proposition 1 (time consistency).

- (i) *In the mechanism-offer version and with a continuum of agents, strong time consistency obtains.*
- (ii) *In the price-offer version,*
 - (a) *Weak time consistency always obtains.*
 - (b) *If $T < +\infty$ and Assumption 1 holds, strong time consistency obtains.*
 - (c) *If $T = +\infty$ and Assumption 1 holds, strong time consistency obtains if either Assumption 2 holds or one focuses on Markov perfect equilibria.*

Proof.

(i) and (ii)(a) The proofs of (i) and (ii)(a) will use very similar constructions. To prove weak time consistency in the price-offer version, consider the following strategies on the equilibrium path:

- The principal sets price $p_t(s^t) = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$ for all (t, s^t) , where $\widehat{\theta}_t(s^t)$ corresponds to an optimal allocation cutoff.
- The agent consumes at date t in state s^t if and only if $\theta \geq \widehat{\theta}_t(s^t)$.

Myopic agent behavior is indeed optimal given the principal's strategy (the current cutoff has zero continuation utility and so do a fortiori all types below the cutoff; higher types strictly benefit from consuming during the period and so do not want to exit). Furthermore, the principal obtains his highest feasible payoff $\widehat{V}_t(s^t)$ starting at any (t, s^t) . And so the principal cannot benefit from deviating from his strategy in any subform (t, s^t) .

Because the principal's strategy is a function of his beliefs, whenever his beliefs are well-defined so is his strategy, which, as verified above, is optimal for any subform. Therefore it remains to consider the subforms in which the principal's beliefs are not uniquely pinned down. There are two possible deviations by the agent that would lead to non-uniquely specified beliefs for the principal. Suppose that $F(\widehat{\theta}_{t-1}(s^{t-1})) = 0$, but the agent has failed to consume at $t-1$ (or earlier). Then beliefs are irrelevant because the game is over due to the no-reentry constraint.

¹⁷Under increasing attractiveness, the principal may want to temporarily "price below marginal cost" so as to keep the clientele. And so a lower cutoff may be (at least temporarily) costly to the principal.

Second, it could be that $F(\widehat{\theta}_{t-1}(s^{t-1})) = 1$, but the agent has always consumed up to $t - 1$ (included). Then, specify that the principal puts all weight on type $\bar{\theta}$ for the rest of the game and sets $p_\tau(s^\tau) = \phi(\bar{\theta}, \bar{\theta}, s_\tau)$ for all (τ, s^τ) with $\tau \geq t$ and $s^\tau \succsim s^t$. The agent therefore cannot obtain a strictly positive continuation utility by deviating. The specified strategies therefore form a perfect Bayesian equilibrium.

Consider now the mechanism-offer version with a continuum of agents. Let the principal offer again in each period price $p_t(s^t) = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$ for all (t, s^t) , but as part of a larger two-stage mechanism within each period: First, agents decide whether to accept. If the total demand, call it q_t , satisfies $q_t \leq 1 - F(\widehat{\theta}_t(s^t))$, then all agents who have accepted are served at price $p_t(s^t)$. Furthermore, if $q_t < 1 - F(\widehat{\theta}_t(s^t))$, the agents who have accepted receive from the principal a large sum of money, and so $q_t \geq 1 - F(\widehat{\theta}_t(s^t))$ necessarily (this participation-enhancing gift is not needed if there are no network externalities, as it is than a dominant strategy for $\theta \geq \widehat{\theta}_t(s^t)$ to accept the offer, or if the externalities are negative). Second, if $q_t > 1 - F(\widehat{\theta}_t(s^t))$, the agents who have accepted pay a small amount $\varepsilon > 0$ to the principal and play a second-period auction¹⁸ in which the best $1 - F(\widehat{\theta}_t(s^t))$ offers are accepted and pay the highest losing bid to the principal (the principal only learns who has won). Importantly, losers do not recoup the payment ε . To see that $q_t = 1 - F(\widehat{\theta}_t(s^t))$, suppose “excess demand”: $q_t > 1 - F(\widehat{\theta}_t(s^t))$. Then, the strict monotonicity of ϕ and the weak monotonicity of the continuation valuation in θ imply that the $1 - F(\widehat{\theta}_t(s^t))$ highest types among those who have paid $p_t(s^t)$ win, while the others predictably lose, and therefore would economize $p_t(s^t)$ by not expressing a demand in the first stage of the mechanism. Hence, strong time consistency obtains.

Assumption 1 is unnecessary in the mechanism-offer game because the principal can organize an auction in case of “excess demand” and offer a large sum of money in case of “excess supply”. This ability to pump in money in such cases allows the principal to correct for miscoordination issues. In the price-offer game, the principal is limited in this regard and miscoordination among the agents could arise, vindicating the use of Assumption 1.

(ii) (b) Suppose that the principal cannot commit and rather sets a price $p_t(s^t)$ in each period. To prove strong time consistency when T is finite and Assumption 1 holds, consider an arbitrary perfect Bayesian equilibrium and let $\delta U_{t+1}(\theta, h^t)$ denote the expected continuation payoff of a type θ that has not yet exited at the end of date t given the entire public history h^t (which includes the realization of s_t and the price p_t). Let $\theta_t^*(h^t) < \bar{\theta}$ denote the equilibrium cutoff given history h^t (if $\theta_t^*(h^t) = \bar{\theta}$, the game is over anyway). Let us show by backward induction that cutoff myopia prevails: $U_{t+1}(\theta_t^*(h^t), h^t) = 0$. Consider date T , with previous cutoff $\theta_{T-1}^* = \theta_{T-1}^*(h^{T-1})$. Suppose that the cutoff enjoys a rent:

$$\phi(\theta_{T-1}^*, \theta_{T-1}^*, s_T) > p_T$$

for some s_T and some optimal price p_T for the principal given θ_{T-1}^* and s_T . The principal’s

¹⁸The exact nature of the auction is not crucial.

date- T payoff is:

$$\int_{\theta_{T-1}^*}^{\bar{\theta}} \psi(\theta, \theta_{T-1}^*, s_T) f(\theta) d\theta + p_T [1 - F(\theta_{T-1}^*)]$$

as all remaining types ($\theta \geq \theta_{T-1}^*$) accept offer p_T . But let the principal offer instead

$$p'_T = \phi(\theta_{T-1}^*, \theta_{T-1}^*, s_T).$$

Then, $p'_T > p_T$ and, from monotonicity and Assumption 1, all types $\theta \geq \theta_{T-1}^*$ accept; and so the principal has increased his date- T payoff by

$$(p'_T - p_T) [1 - F(\theta_{T-1}^*)] > 0,$$

a contradiction.

Now consider date $T - 1$. The previous cutoff is $\theta_{T-2}^*(h^{T-2})$, the state is s_{T-1} and the principal sets some price $p_{T-1} \in \text{support}(\sigma_{T-1}^P(p^{T-2}, s^{T-1}))$. Either $\theta_{T-1}^* > \theta_{T-2}^*$, and then

$$\begin{aligned} & \phi(\theta_{T-2}^*, \theta_{T-1}^*, s_{T-1}) - p_{T-1} + \delta U_T(\theta_{T-2}^*, h^{T-1}) \\ & < \phi(\theta_{T-1}^*, \theta_{T-1}^*, s_{T-1}) - p_{T-1} + \delta U_T(\theta_{T-1}^*, h^{T-1}) \leq 0 \end{aligned}$$

where the weak inequality (which is an equality if $\theta_{T-1}^* < \bar{\theta}$) results from the fact that θ_{T-1}^* is the cutoff and that the continuation valuation is monotonic in type; and so θ_{T-2}^* exits at date $T - 1$.

Or $\theta_{T-1}^* = \theta_{T-2}^*$. From the induction hypothesis again, type θ_{T-2}^* has no continuation value and net utility from $T - 1$ on therefore equal to:

$$\phi(\theta_{T-2}^*, \theta_{T-2}^*, s_{T-1}) - p_{T-1}.$$

Were this utility to be strictly positive (it cannot be strictly negative, otherwise θ_{T-2}^* and nearby types would exit), the principal would raise price p_{T-1} to $p'_{T-1} = \phi(\theta_{T-2}^*, \theta_{T-2}^*, s_{T-1})$, still inducing no exit and raising payoff.

The same reasoning shows by backward induction that the cutoff type never has a strictly positive continuation utility.

Finally, suppose that at date 0 the principal offers price $p_0 = \phi(\hat{\theta}_0(s^0), \hat{\theta}_0(s^0), s_0)$ that makes the optimal cutoff type myopically indifferent between accepting and rejecting p_0 . The cutoff θ_0^\dagger must necessarily satisfy $\theta_0^\dagger \leq \hat{\theta}_0(s^0)$; for, if $\theta_0^\dagger > \hat{\theta}_0(s^0)$, $\phi(\theta_0^\dagger, \theta_0^\dagger, s_0) > p_0$ and so types θ_0^\dagger and just below should accept p_0 . But if $\theta_0^\dagger < \hat{\theta}_0(s^0)$, type θ_0^\dagger has negative date-0 payoff and has zero continuation utility, a contradiction. Hence $\theta_0^\dagger = \hat{\theta}_0(s^0)$. By the same reasoning, the principal by setting $p_1(s^1) = \phi(\hat{\theta}_1(s^1), \hat{\theta}_1(s^1), s_1)$ uniquely induces cutoff $\theta_1^\dagger = \hat{\theta}_1(s^1)$, and so forth by induction.

(ii) (c) Allow now $T = +\infty$ and make Assumption 2 as well. Suppose that the princi-

pal stubbornly sets a price equal to the flow valuation of the optimal cutoff type $\widehat{p}_t(s^t) = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$ for all t, s^t . There is no commitment to the sequence, nor is this sequence optimal for the principal in every subform. This is simply a strategy choice that reacts to the state of nature, but not to the observed amount of exit.

At date 0, the cutoff type satisfies $\theta_0^\dagger \leq \widehat{\theta}_0(s^0)$ from Assumption 1 and so types $\theta > \widehat{\theta}_0(s^0)$ optimally accept offer $\widehat{p}_0(s^0)$ regardless of the expectation concerning the continuation behavior of the principal. Assumption 2 then implies that the principal's date-0 payoff weakly exceeds

$$\int_{\widehat{\theta}_0(s^0)}^{\bar{\theta}} \psi(\theta, \widehat{\theta}_0(s^0), s_0) f(\theta) d\theta + \widehat{p}_0(s^0) [1 - F(\widehat{\theta}_0(s^0))].$$

Similarly, at date 1, the cutoff in state s^1 will be some $\theta_1^\dagger \leq \widehat{\theta}_1(s^1)$ and yield a weakly higher payoff than $\int_{\widehat{\theta}_1(s^1)}^{\bar{\theta}} \psi(\theta, \widehat{\theta}_1(s^1), s_1) f(\theta) d\theta + \widehat{p}_1(s^1) [1 - F(\widehat{\theta}_1(s^1))]$, and so forth. Thus the principal can guarantee himself the commitment payoff.

Alternatively, we can focus on Markov perfect equilibria and mimic the proof of part (b). Suppose that for some history $p_t < \phi(\theta_{t-1}^*, \theta_{t-1}^*, s_t)$. Then $\theta_t^* = \theta_{t-1}^*$. But then the principal could charge $p_t' = \phi(\theta_{t-1}^*, \theta_{t-1}^*, s_t)$ and still keep all $\theta \geq \theta_{t-1}^*$ on board; and so the cutoff would remain θ_{t-1}^* . The (random) payoff-relevant state in all future periods would be unchanged and so would continuation payoffs in a Markov perfect equilibrium. The principal's payoff would therefore strictly increase, a contradiction. We thus conclude that $p_t \geq \phi(\theta_{t-1}^*, \theta_{t-1}^*, s_t)$, which, together with cutoff monotonicity, implies that the cutoff in a given period never enjoys a strictly positive continuation valuation. In turn, agent myopic behavior implies that stubbornly setting $\widehat{p}_t(s^t) = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$ delivers the commitment payoff. ■

Later on, when we search for a characterization of the non-commitment game when time consistency obtains (Section 4), we will assume that either strong time consistency obtains, or that the equilibrium of the no-commitment game delivering the commitment outcome is selected. As shown in Proposition 1, this is a rather weak requirement.

3.3 Unequal time preference

Suppose, only for the sake of this section, that the principal and the agent have different discount factors (δ_A for the agent and, for comparison purposes, δ for the principal). We assume that the agent cannot commit. Lemma 1 shows that this assumption is innocuous when the two parties are equally patient; by contrast, it is not innocuous under heterogeneous discounting: when the agent is impatient, the agent's payments are optimally backloaded, but, as we will see, the agent's inability to commit prevents this backloading.

Proposition 2 (unequal discount factors). *Suppose that δ_A and δ differ and that the agent cannot commit. Make Assumptions 1 and 2 and assume a finite horizon. Keeping the principal's discount factor δ fixed and varying the agent's discount factor δ_A , let $V^{nc}(\delta_A)$ denote*

the principal's (unique) non-commitment payoff that prevails when the principal cannot commit either.

(i) The principal's payoff does not depend on the agent's discount factor: $V^{nc}(\delta_A) = \widehat{V}$.

(ii) The outcome is time consistent if and only if $\delta_A \leq \delta$ (impatient or equally patient agent).

Intuitively, when the agent is impatient ($\delta_A < \delta$), the principal would like to backload the agent's payments. The agent's lack of commitment however makes this impossible; and so the best policy under commitment satisfies the cutoff myopia property (see Lemma 1). The agent's discount factor is then irrelevant since the current cutoff never has a positive continuation utility.

When the agent is patient ($\delta_A > \delta$) and under commitment, the principal would like to frontload the agent's payments. This policy however is infeasible when the principal cannot commit; by backward induction, the principal never leaves any surplus to the previous cutoff type and again cutoff myopia prevails, making the agent's discount factor irrelevant.

4 Characterization of sequential screening outcomes

Proposition 1 transforms the search for a perfect Bayesian equilibrium of the no-commitment game into a simple optimization problem. This section characterizes the sequential screening outcome in three cases of interest.

Definition 1. Let $s_t \in \mathbb{R}$. Define the externality-augmented virtual surplus as

$$\Lambda(\theta^*, s) \equiv \Gamma(\theta^*, \theta^*, s) - \frac{\int_{\theta^*}^{\bar{\theta}} \frac{\partial \Gamma[\theta, \theta^*, s]}{\partial \theta^*} f(\theta) d\theta}{f(\theta^*)}.$$

$\Lambda(\theta^*, s)f(\theta^*)d\theta^*$ is the loss of aggregate virtual surplus if we exclude the marginal types $[\theta^*, \theta^* + d\theta^*]$ from consumption.¹⁹

In the rest of the paper, we make the following regularity assumption:

Assumption 3 (externality augmented virtual surplus).

¹⁹For example, for the examples given above

$$\Lambda(\theta^*, s) = \theta^* - c - \frac{1 - F(\theta^*)}{f(\theta^*)}$$

in the conversion game, and

$$\Lambda(\theta^*, s) = \theta^* + 2\alpha[1 - F(\theta^*)] + s - c - \frac{1 - F(\theta^*)}{f(\theta^*)}$$

in the technology disadoption game.

(i) For all θ^* and s , $\frac{\partial \Lambda(\theta^*, s)}{\partial \theta^*} > 0$.

(ii) Furthermore, states are ranked: $s \in \mathbb{R}$, and for all θ^* and s , $\frac{\partial \Lambda(\theta^*, s)}{\partial s} > 0$.

Part (ii) of Assumption 3 simply defines higher states as better ones. Part (i), which guarantees the strict quasi-concavity in θ^* of the function $\int_{\theta^*}^{\bar{\theta}} \Gamma(\theta, \theta^*, s) f(\theta) d\theta$, is the generalization of the standard assumption of monotonicity of the virtual surplus to allow for externalities. It is satisfied in the conversion and disadoption games provided that the virtual type $\theta - [1 - F(\theta)]/f(\theta)$ is increasing (for this, it suffices that the Mills ratio $(1 - F)/f$ be decreasing) and network externalities are not too strong if they are positive.

When there are no network externalities and $\frac{\partial \psi}{\partial \theta}(\theta, s) \geq 0$ (high types are valued weakly more by the principal), a sufficient condition for part (i) to be satisfied is:

$$\frac{1 - \left(\frac{1 - F(\theta^*)}{f(\theta^*)} \right)'}{\frac{1 - F(\theta^*)}{f(\theta^*)}} \geq \left(\frac{\frac{\partial^2 \phi}{\partial \theta^2}}{\frac{\partial \phi}{\partial \theta}} \right) (\theta^*, s)$$

for all (θ^*, s) , where the left-hand side is positive for a log-concave distribution.

4.1 Monotone attractiveness

Let us first consider the case in which the consumption offered by the principal becomes (stochastically) more or less attractive over time.

Definition 2 (monotone attractiveness). Suppose that states are ranked as in Assumption 3. Increasing (resp. decreasing) attractiveness holds when $s_{t+1} \geq s_t$ (resp. $s_{t+1} \leq s_t$) for all s_{t+1} in the support of G conditionally on s^t .

Increasing attractiveness for example captures habit formation on the demand side and learning by doing on the supply side. By contrast, decreasing attractiveness may result from a decreasing interest in the incumbent consumption or gradual improvements in the alternative option.²⁰ Obviously, increasing and decreasing attractiveness include as a special case the case of a constant demand.

Proposition 3 (monotone attractiveness). Let $S_t \equiv S \subseteq \mathbb{R}$ for all t .

(i) Under deterministic increasing attractiveness (s^T is a singleton), exit occurs only in the initial period: there exists $\hat{\theta}$ such that at the equilibrium outcome $\hat{\theta}_t(s^t) = \hat{\theta}$ for all t .

(ii) Under decreasing attractiveness, for any (t, s^t, s^{t+1}) such that $s^{t+1} \succ s^t$, then $\hat{\theta}_{t+1}(s^{t+1}) > \hat{\theta}_t(s^t)$ when $s_{t+1} < s_t$. The cutoffs are then given by myopic principal optimization: $\Lambda(\hat{\theta}_t(s^t), s_t) = 0$ for all (t, s^t) .

²⁰One may here have in mind a temporary recession or lack of attractiveness of employer (bad management, scandal).

Proof. For a given allocation $(\mathbf{p}, \boldsymbol{\theta}^*)$, the tree (with generic element (z, s^z)) can be decomposed into the union \mathcal{T} of disjoint subtrees \mathcal{S} of complete subpaths over which the cutoff is constant:

$\mathcal{S} \in \mathcal{T} \iff \exists (t, s^t)$ such that

$$(i) \quad U_t(\theta_t^*(s^t), s^t; \mathbf{p}, \boldsymbol{\theta}^*) = 0$$

$$(ii) \quad U_{t-1}(\theta_t^*(s^t), s^{t-1}; \mathbf{p}, \boldsymbol{\theta}^*) > 0 \text{ (where } s^t \succ s^{t-1}\text{)}$$

$$(iii) \quad \forall (z, s^z) \text{ such that } s^z \succ s^t, \theta_z^*(s^z) = \theta_t^*(s^t) \text{ if and only if } (z, s^z) \in \mathcal{S}.$$

In this definition, (t, s^t) is the vertex of the subtree \mathcal{S} . The cutoff, $\theta_t^*(s^t)$, at the vertex obtains a zero continuation utility (condition (i)), while it received a strictly positive continuation utility (and therefore was not the cutoff) earlier (condition (ii)); so $\theta_t^*(s^t)$ is an interior cutoff at date t . The subtree is the set of (τ, s^τ) with $s^\tau \succ s^t$ that exhibit cutoff $\theta_t^*(s^t)$ (condition (iii)).

The principal maximizes

$$V = \int_{\underline{\theta}}^{\bar{\theta}} E_{s^t} \left[\sum_{t=0}^{t=T} \delta^t X^t(\theta, s^t) \Gamma(\theta, \theta_t^*(s^t), s_t) \right] f(\theta) d\theta,$$

subject to the feasibility constraint (F) yielding first-order condition:²¹

- either the cutoff is constrained by the previous one: $\theta_t^*(s^t) = \theta_{t-1}^*(s^{t-1})$,
- or the expected discounted virtual surplus along a constant-cutoff sub-tree is equal to 0:

$$\tilde{\Lambda}(\theta_t^*(s^t), s_t) \equiv \Lambda(\theta_t^*(s^t), s_t) + E \left[\sum_{\tau=0}^{\tau=T-t} \delta^\tau \mathbb{I}_{\{\theta_{t+\tau}^*(s^{t+\tau}) = \theta_t^*(s^t)\}} \Lambda(\theta_{t+\tau}^*(s^{t+\tau}), s_{t+\tau}) | s^t \right] = 0,$$

- or $\theta_t^*(s^t) = \underline{\theta}$ and $\tilde{\Lambda}(\underline{\theta}, s_t) \geq 0$,
- or $\theta_t^*(s^t) = \bar{\theta}$ and $\tilde{\Lambda}(\bar{\theta}, s_t) \leq 0$.

When the optimal cutoff is not at one of the boundaries of the support of types, either the cutoff monotonicity constraint binds or we have an interior optimum. We will show that the constraint binds at all periods under deterministic increasing attractiveness, and that the optimum is interior in all periods under decreasing attractiveness.

(i) With a deterministic state, we can subsume the dependence of variables on the state through a time index: $\Lambda_t(\hat{\theta}_t)$. A constant-cutoff sub-tree \mathcal{S} is formed by a set of periods $\{t, \dots, z\}$ such that $\hat{\theta}_t = \dots = \hat{\theta}_z$. Suppose that there are at least two such subtrees and so the cutoff is not

²¹To obtain this first-order condition, maximize V over $\{x(\cdot, \cdot), \theta^*(\cdot)\}$ subject to the constraints:

$$\int_{\underline{\theta}}^{\bar{\theta}} x_t(\theta, s^t) f(\theta) d\theta \leq 1 - F(\theta_t^*(s^t))$$

and $\theta_{t+1}^*(s^{t+1}) \geq \theta_t^*(s^t)$ if $s^{t+1} \succ s^t$.

constant over time. Consider the first two, from 0 to $t - 1$ and from t to z say. One has $\widehat{\theta}_0 < \widehat{\theta}_t$ and from the previous characterization:

$$\sum_{\tau=0}^{t-1} \delta^\tau \Lambda_\tau(\widehat{\theta}_0) \geq 0 \geq \sum_{\tau=t}^z \delta^{\tau-t} \Lambda_\tau(\widehat{\theta}_t).$$

Because $\Lambda_\tau(\theta^*)$ is weakly increasing in τ for all θ^* , one necessarily has

$$\Lambda_{t-1}(\widehat{\theta}_0) \geq 0 \geq \Lambda_t(\widehat{\theta}_t),$$

which is inconsistent with $\widehat{\theta}_0 < \widehat{\theta}_t$ and $\Lambda(\cdot)$ being weakly increasing in time and strictly increasing in the cutoff. Hence, the cutoff must be constant over time.

(ii) Under decreasing attractiveness (deterministic or stochastic), the principal-myopic optimum given by (uniquely so from Assumption 3)

$$\forall(t, s^t) : \quad \Lambda(\widehat{\theta}_t(s^t), s_t) = 0,$$

satisfies the feasibility constraints as $\widehat{\theta}_{t+1}(s^{t+1}) \geq \widehat{\theta}_t(s^t)$ for $s^{t+1} \succ s^t$. It can be implemented through prices

$$p_t(s^t) = \phi(\widehat{\theta}_t(s^t), s_t). \quad \blacksquare$$

Remark 1. We need the assumption that the state's evolution is deterministic in part (i) of Proposition 3. That is, it does not suffice that s_t be stochastically increasing with time. To see this, suppose there are only two periods and two states at date 1: $s_0 = s_1^L < s_1^H$. In general, the principal will want to keep the participation high ($\Lambda_0 < 0$) so as to keep an option value of setting $\widehat{\theta}_1(s_1^H) = \widehat{\theta}_0$ low in state s_1^H . If “disappointing” news (s_1^L) accrue, then the principal raises the cutoff to $\widehat{\theta}_1(s_1^L) > \widehat{\theta}_0$. Despite increasing attractiveness, exit is not clustered at date 0.

As pointed out by a referee, a weaker result then obtains. If $s^{t+1} \succ s^t$, then let $\bar{s}^{t+1} = (s^t, \bar{s}_{t+1})$ where \bar{s}_{t+1} is the supremum in the support of s_{t+1} conditional on s^t . Then for stochastic increasing attractiveness,

$$\widehat{\theta}_{t+1}(\bar{s}^{t+1}) = \widehat{\theta}_t(s^t).$$

Corollary 1 (sufficient condition for Assumption 2 to hold). *Assumption 2 (the principal benefits from a greater clientele at price $p_t = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t)$) holds provided that:*

- (a) $\int_{\theta^*}^{\bar{\theta}} \psi(\theta, \theta^*, s) f(\theta) d\theta = -C(1 - F(\theta^*), s)$ with $C_{11} \leq 0$ (constant or increasing returns),
- (b) ϕ is separable in θ and θ^* : $\phi(\theta, \theta^*, s) = \xi(\theta, s) + \nu(\theta^*, s)$ and Assumption 1 holds, and
- (c) decreasing (including constant) attractiveness holds.

4.2 Non-monotone attractiveness: ironing techniques

In general, attractiveness need not be monotonic. For example, a relationship may need to be build up over time to become attractive, and is later threatened by new, external opportunities.

In this section, we will assume for simplicity that the evolution of attractiveness is deterministic; so functions will be indexed only by time, and not by state s_t . Also for expositional simplicity, we assume that time is continuous on $[0, T]$.²² Let r denote the interest rate.

With this notational simplification, the optimization program becomes:

$$\begin{aligned} \sup_{\{\theta_t^*\}} \left\{ \int_0^T e^{-rt} W_t(\theta_t^*) dt \right\} & \quad (II) \\ \text{s.t. } \frac{d\theta_t^*}{dt} \geq 0. & \end{aligned}$$

This program is highly reminiscent of that in static adverse selection models, in which a one-dimensional allocation must be monotonic in type for incentive compatibility reasons. Here “type” is replaced by “time” and “allocation” by “cutoff”. One can push the analogy further: the decreasing attractiveness case of Section 4.1 corresponds to the “regular case” of Myerson (1981) while the increasing attractiveness one is an extreme “non-regular case” of Myerson, called “non-responsiveness” by Guesnerie and Laffont (1984) and giving rise to “full pooling”. More generally, the optimal allocation can be obtained through the convex analysis used in Myerson and Guesnerie-Laffont. Because the only contribution of this section is to identify the formal analogy between the two problems and use it to characterize dynamic screening, we content ourselves with stating the result.

Proposition 4 (ironing for non-monotone deterministic attractiveness). *Assume that attractiveness is deterministic; that W_t is strictly concave in θ_t^* and is C^2 ; that the C^1 solution $\hat{\theta}_t$ of $dW_t/d\theta_t^* = \Lambda_t(\hat{\theta}_t) = 0$ is such that $d\hat{\theta}_t/dt$ changes sign a finite number of times.*

Then the solution to program (II) exists, is unique and exhibits a finite number of points of discontinuity. $\hat{\theta}_t$ coincides with $\hat{\theta}_t$ except on a finite numbers of intervals (t_k, t_{k+1}) such that $\hat{\theta}_t$ is constant on each of these intervals and

$$\int_{t_k}^{t_{k+1}} e^{-rt} \Lambda_t(\hat{\theta}_t) dt = 0.$$

The evolution of the cut-off is depicted in Figure 2(a), which is also familiar from static incentive theory.

²²See the beginning of the proof of Proposition 3 for the discrete-time version of the ironing condition.

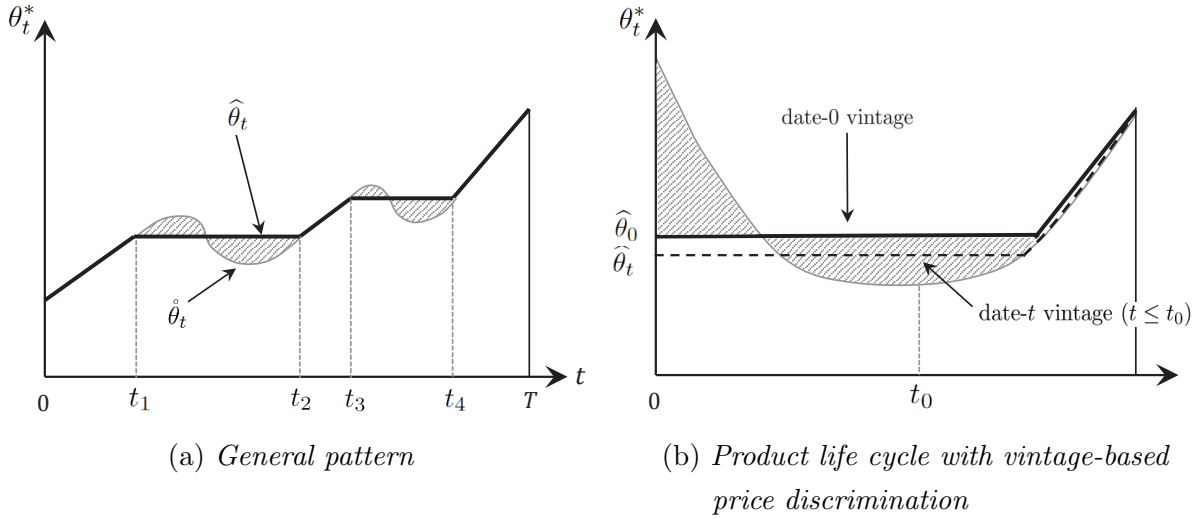


Figure 2

Under deterministic attractiveness, one would thus expect episodes of gradual exit alternating with exit-less periods.

Remark (product life cycle). In some environments, attractiveness is likely to conform to a simple U -shaped pattern. For example, a new product may face the following life cycle. At introduction it may be relatively unattractive: bugs still need to be fixed, network externalities are still limited by the small goodwill, etc. The firm must then offer very advantageous conditions to attract customers. Then, the product becomes more and more attractive, until time t_0 , at which a viable competition enters and slowly erodes the product's installed base of consumers.

The prediction for such life cycles is that the clientele is built through very advantageous conditions and remains steady until it declines over time. One implication is unpalatable for this industrial organization example: the clientele is built overnight. To obtain a more gradual build-up of clientele, though, one may generalize the model and assume that consumers are made aware or arrive in the market in different, but similar cohorts, as in Section 4.3 below. Proposition 4 holds all the same as long as the firm can practice vintage-based price discrimination (Figure 2(b)), for which the aggregation of cohorts yields an increasing and then decreasing demand. When vintage-based price discrimination is infeasible, things become more complex (see part (ii) of Proposition 7), but backward induction techniques from the date t_0 at which the product becomes less attractive can still be employed to yield a similar pattern of demand evolution.

4.3 Transient shocks

Suppose now that the realizations of the shocks s_t are identically and independently distributed over time. Let $g(s_t)$ and $G(s_t)$ denote the density and the cumulative distribution of the shock. We assume that the horizon is infinite, so as to provide a simpler characterization of the equilibrium outcome.²³

²³A similar characterization is available by backward induction for a finite T , but the strategy then depends on the length of the remaining horizon.

Proposition 5 (transient shocks). *Suppose that shocks are identically and independently distributed with density $g(\cdot)$ and that $T = +\infty$. The optimal outcome is characterized by: for all (t, s^t)*

$$\widehat{\theta}_t(s^t) = \theta^u(\min_{\tau \leq t} s_\tau),$$

where $\theta^u(s)$, a decreasing function of s , is uniquely defined by $H(\theta^u(s)) = 0$ where

$$H(\theta^*, s) \equiv \Lambda(\theta^*, s) + \frac{\delta}{1 - \delta[1 - G(s)]} \int_s^\infty \Lambda(\theta^*, \tilde{s})g(\tilde{s})d\tilde{s},$$

if $H(\underline{\theta}, s) \leq 0 \leq H(\bar{\theta}, s)$, and $\theta^u(s) = \underline{\theta}$ if $H(\underline{\theta}, s) \geq 0$ and $\theta^u(s) = \bar{\theta}$ if $H(\bar{\theta}, s) \leq 0$.

Proof. Maximizing V with respect to $x_t(\theta, s^t)$ yields the first-order condition

$$x_t(\theta, s^t) = 1 \iff \Lambda(\theta, s_t) + E_{s^\infty | s^t} [\sum_{\tau=t+1}^\infty \delta^{\tau-t} X_{t+1}^\tau(\theta, s^\tau) \Lambda(\theta, s_\tau)] \geq 0$$

where $X_{t+1}^\tau(\theta, s^\tau) \equiv x_{t+1}(\theta, s^{t+1}) \cdots x_\tau(\theta, s^\tau)$. Note that if $X^{t-1}(\theta, s^{t-1}) = 0$, then the first-order condition is irrelevant, but one can still impose this date- t first-order condition without loss of generality. Because Λ is strictly increasing in θ , there is a unique threshold $\widehat{\theta}_t(s^t)$ in $[\underline{\theta}, \bar{\theta}]$ for each s^t such that this condition is satisfied if and only if $\theta \geq \widehat{\theta}_t(s^t)$. The stationarity of the problem then suggests looking for a strictly decreasing cutoff $\widehat{\theta}_t(s^t) = \theta^u(s_t)$ (where “ u ” stands for “unconstrained by the previous exit pattern”). Noting that $X_{t+1}^\tau(\theta^u(s_t), s^\tau) = 0$ if $\min\{s_{t+1}, \dots, s_\tau\} < s_t$, such a cutoff then satisfies (if interior, i.e., in $(\underline{\theta}, \bar{\theta})$):

$$\Lambda(\theta^u(s), s) + \delta \left[\int_s^\infty \Lambda(\theta^u(s), \tilde{s})g(\tilde{s})d\tilde{s} + [1 - G(s)]\delta \left[\int_s^\infty \Lambda(\theta^u(s), \tilde{s})g(\tilde{s})d\tilde{s} + \dots \right] \right] = 0$$

or

$$\Lambda(\theta^u(s), s) + \frac{\delta}{1 - \delta[1 - G(s)]} \int_s^\infty \Lambda(\theta^u(s), \tilde{s})g(\tilde{s})d\tilde{s} = 0.$$

Differentiating this condition and using it to eliminate two terms, one obtains:

$$\frac{\partial \Lambda}{\partial s} + \left[\frac{\partial \Lambda}{\partial \theta} + \frac{\delta}{1 - \delta[1 - G(s)]} \int_s^\infty \frac{\partial \Lambda}{\partial \theta} g(\tilde{s})d\tilde{s} \right] \frac{d\theta^u}{ds} = 0$$

and so

$$\frac{d\theta^u}{ds} < 0.$$

The tentative solution

$$\widehat{\theta}_t(s^t) = \max_{\tau \leq t} \{\theta^u(s_\tau)\} = \theta^u(\min_{\tau \leq t} s_\tau)$$

indeed satisfies the first-order condition above. ■

Corollary 2 (testable predictions for transient shocks). *With transient shocks and an*

infinite horizon, let for all t, s^{t-1}

$$\mathcal{V}_t = F\left(\theta^u\left(\min_{\tau \leq t} s_\tau\right)\right) - F\left(\theta^u\left(\min_{\tau \leq t-1} s_\tau\right)\right)$$

denote the volume of exit at date t . Then the following properties hold for all $\tau > t$:

(i) Decreasing exit: $E[\mathcal{V}_t | s^{t-1}] \geq E[\mathcal{V}_\tau | s^{t-1}]$.

(ii) Negative serial correlation: $\frac{\partial}{\partial \mathcal{V}_t} E_t[\mathcal{V}_\tau | s^{t-1}, \mathcal{V}_t] < 0$.

Proof. Let $\hat{s} \equiv \min_{\tau \leq t-1} s_\tau$; and let $g_n(s) \equiv ng(s)[1-G(s)]^{n-1}$ denote the density of the distribution of the minimum realization over n periods.

(i) Then

$$\begin{aligned} E[\mathcal{V}_\tau | s^{t-1}] &= \int_{\hat{s}}^{\infty} \left[\int_{-\infty}^{\hat{s}} \left[F(\theta^u(\tilde{s})) - F(\theta^u(\hat{s})) \right] g(\tilde{s}) d\tilde{s} \right] g_{\tau-t}(s) ds \\ &\quad + \int_{-\infty}^{\hat{s}} \left[\int_{-\infty}^s \left[F(\theta^u(\tilde{s})) - F(\theta^u(s)) \right] g(\tilde{s}) d\tilde{s} \right] g_{\tau-t}(s) ds \\ &\leq \int_{-\infty}^{\hat{s}} \left[F(\theta^u(\tilde{s})) - F(\theta^u(\hat{s})) \right] g(\tilde{s}) d\tilde{s} = E[\mathcal{V}_t | s^{t-1}]. \end{aligned}$$

(ii) Note that $E[\mathcal{V}_\tau | s^{t-1}]$ depends only on, and is increasing with \hat{s} ; and that $E[\mathcal{V}_\tau | s^{t-1}, \mathcal{V}_t] = E[\mathcal{V}_\tau | \min\{\hat{s}, s_t\}]$. Because \mathcal{V}_t is (weakly) decreasing in s_t , then $E[\mathcal{V}_\tau | s^{t-1}, \mathcal{V}_t]$ is weakly decreasing in \mathcal{V}_t , and strictly so when $\mathcal{V}_t > 0$. ■

5 Robustness

The section provides several robustness results by allowing inflows of new agents and finite re-entry costs.

5.1 Finite re-entry costs

We have assumed so far that re-entry costs were infinite. In practice, though, except in the case of apostasy, reestablishing a relationship is costly, but not infinitely so. Although the very reasons why the relationship broke up in the first place and the investments made in the meantime in alternative relationships hinder such reintegration, we see customers returning to brands they dropped, spouses remarrying their divorced partner or employees being reemployed by their abandoned employer.

Intuitively, our results should carry over for large re-entry costs. We may then wonder, how large is “large”? This section sheds some light on this question. We generalize the model by considering an unchanged flow payoff for the agent while the principal’s flow payoff is reduced by the expected re-entry cost:

$$\int_{\underline{\theta}}^{\bar{\theta}} \max \{0, (x_t(\theta) - x_{t-1}(\theta))r\} f(\theta) d\theta,$$

where $r \geq 0$ is the re-entry cost incurred at date t whenever $x_t = 1$ and $x_{t-1} = 0$. For expositional simplicity, we assume that r represents a cost that is borne by the principal (this assumption also allows us to abstract from the possibility that the principal lowers the price substantially in order to attract re-entrants, who might then fake re-entry – i.e., not spend r – and thereby “take the money/surplus and run”). We assume that there is a continuum of agents, and that the principal cannot price discriminate (i.e., he charges a uniform price p_t in each period).²⁴

Proposition 6 (re-entry). *Suppose that exit is not necessarily definitive: re-entry involves cost $r \geq 0$.*

- (i) *Under decreasing/constant attractiveness, for all $r \geq 0$, payoff \widehat{V} without re-entry is still an equilibrium payoff with re-entry, (the unique one if T is finite and Assumption 1 is satisfied).*
- (ii) *Under deterministic,²⁵ strictly increasing attractiveness, equilibrium payoff \widehat{V} without re-entry is still an equilibrium payoff with re-entry, provided that $r \geq \underline{r}$ for some \underline{r} computed in Appendix D (again, the unique one if T is finite and Assumption 1 is satisfied).*

Intuitively, the principal wants to shed goodwill when the relationship becomes less attractive, and so would not want to incur a re-entry cost, however small, to bring back agents who have quit in the past. By contrast, under increasing attractiveness, the principal might want the agent to re-enter later on. The re-entry cost must then be sufficiently large for the absorbing-exit solution to prevail.

5.2 Inflow of new agents

We have so far assumed that all agents are present at date 0. Suppose by contrast that at each date t , a new cohort of agents enters, that (as in Conlisk et al (1984) and Sobel (1991)’s durable goods models with negative selection) has the same type distribution $F(\theta)$ as previous ones.²⁶ Newcomers, who live from date t through date T , have a chance to interact with the principal at date t . Non-membership at the entry date is, like exit, an absorbing state.

²⁴With a single agent and, say, r small, a ratchet effect would arise: an early rejection by the agent would lead to a price cut relative to an early acceptance. Similarly, ratcheting might occur with a continuum of agents and price discrimination between the installed base and re-entrants.

²⁵The analysis can be generalized to a stochastic environment, but at the expense of increased notational complexity.

²⁶We want to abstract from the standard issues associated with the impact of third-degree price discrimination under heterogeneous submarkets (see, e.g., Aguirre et al 2010 for a recent entry on this topic).

To abstract from direct interactions among cohorts, assume that there are no cross-cohorts network externalities (there can be within-cohort network externalities) and that returns to scale are constant; otherwise the virtual surplus for a given cohort would depend on the number of retained agents in the other cohorts. To capture these requirements, we thus assume that the function Λ is invariant to the presence of other cohorts, and so the only interaction among cohorts is through pricing.²⁷ If the principal is able to price discriminate among cohorts, each cohort is then treated in isolation and the previous analysis, including that of time consistency, applies. Rather, we ask whether time consistency still holds when the principal is constrained to practise uniform pricing.²⁸

Proposition 7 (inflow of new agents). *Suppose that each period $t = 0, 1, \dots, T$, a new cohort of arbitrary mass and type distribution $F(\theta)$ arrives. Suppose that Λ is invariant to the presence of other cohorts (no cross-cohorts network externalities and no returns to scale). For the class of monotone-attractiveness games considered in Proposition 3:*

- (i) *Under commitment, uniform pricing does as well for the principal as discriminatory pricing.*
- (ii) *Weak²⁹ time consistency obtains for decreasing/constant attractiveness, but not for deterministic, strictly increasing attractiveness.*

Proof. It will be convenient to consider sequentially the cases of decreasing/constant attractiveness and of strictly increasing attractiveness.

Under decreasing/constant attractiveness, the optimal policy for cohort t when price discrimination is feasible (see Proposition 3) is given by myopic optimization:

$$\Lambda(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t) = 0 \quad \text{for all } (t, s^t).$$

Thus the cutoff $\widehat{\theta}_t(s^t)$ is independent of the cohort and can be implemented by cohort independent price

$$p_t(s^t) = \phi(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t).$$

Furthermore, the function Λ is left invariant by left truncations,³⁰ and so the price path just defined is time consistent.

²⁷More precisely, we assume that the principal's intertemporal payoff V is separable across cohorts. Let $c \in \{0, 1, \dots, T\}$ denote a cohort, with mass α_c . Let $X_c^t(\theta, s^t)$ denote the probability that type θ of cohort c has not exited yet at $t \geq c$ in state s^t . Similarly, $\theta_{t,c}^*(s^t)$ is the date- t cutoff for cohort c . Then

$$V = \sum_{c=0}^{c=T} \alpha_c \int_{\underline{\theta}}^{\bar{\theta}} E_{s^T} \left[\sum_{t=c}^{t=T} \delta^t X_c^t(\theta, s^t) \Gamma(\theta, \theta_{t,c}^*(s^t), s_t) \right] f(\theta) d\theta,$$

²⁸There has been substantial interest in the literature on negative selection regarding the impact of the arrival of new cohorts under uniform pricing (Conlisk et al 1984 and Sobel 1991 are classic references here).

²⁹Strong time consistency can be obtained through further assumptions as in Proposition 1. We focus on weak time consistency for conciseness.

³⁰For example, when cohorts have equal sizes, at date t , the posterior cumulative over the $(t+1)$ existing cohorts is:

$$F_t(\theta) = \frac{F(\theta)}{1 + t[1 - F(\theta_{t-1}^*)]} \quad \text{for } \theta \leq \theta_{t-1}^* \quad \text{and} \quad F_t(\theta) = 1 - \frac{(t+1)[1 - F(\theta)]}{1 + t[1 - F(\theta_{t-1}^*)]} \quad \text{for } \theta \geq \theta_{t-1}^*.$$

Next, we consider the case of deterministic, strictly increasing attractiveness. We first want to show that uniform pricing does as well as discriminatory pricing. Given that attractiveness increases over time, the generation t cutoff will enjoy future rents under uniform pricing. To cancel these rents, the principal ought to frontload the payment pattern. More precisely, under price discrimination, the optimal policy for each cohort t consists in a constant cutoff $\widehat{\theta}(t)$ defined by (using the notation Λ_τ of the proof of Proposition 3):³¹

$$\sum_{\tau=t}^{\tau=T} \delta^{\tau-t} \Lambda_\tau(\widehat{\theta}(t)) = 0.$$

Because s_t is strictly increasing, $\widehat{\theta}(t)$ is strictly decreasing.³²

Let

$$P_t \equiv \sum_{\tau=t}^{\tau=T} \delta^{\tau-t} \phi(\widehat{\theta}(t), \widehat{\theta}(t), s_\tau)$$

denote the present discounted value of the cohort- t marginal type's surplus. P_t also represents what cohort t will have to pay for membership from t through T . Let p_t be defined by $P_t = p_t + \delta P_{t+1}$. Then

$$p_t \equiv \phi(\widehat{\theta}(t), \widehat{\theta}(t), s_t) + \sum_{\tau=t+1}^{\tau=T} \delta^{\tau-t} \left[\phi(\widehat{\theta}(t), \widehat{\theta}(t), s_\tau) - \phi(\widehat{\theta}(t+1), \widehat{\theta}(t+1), s_\tau) \right].$$

The second term on the right-hand side of this expression of p_t is the present discounted rent of the cohort- t marginal type and is strictly positive. The difference $p_t - \phi(\widehat{\theta}(t), \widehat{\theta}(t), s_t)$ thus measures the required frontloading of the payment that delivers cutoff $\widehat{\theta}(t)$ for generation t . The price sequence $\{p_t\}$ generates cutoff sequence $\{\widehat{\theta}(t)\}$.³³

Finally, suppose that there is no commitment and that the principal charges uniform prices. Is the sequence $\{p_t\}$ defined above an equilibrium of the non-commitment game? To see that this is not the case, consider the two-period version of the model: $t = 0, 1$ and assume for notational simplicity that there are no network externalities, even within a cohort. Necessarily, for time consistency to obtain,

$$p_1 = \phi(\widehat{\theta}(1), s_1)$$

and

$$p_0 + \delta p_1 = \phi(\widehat{\theta}(0), s_0) + \delta \phi(\widehat{\theta}(0), s_1).$$

Furthermore (recalling the assumption of no network externalities and constant returns to scale, so $\Lambda = \Gamma$)

$$\phi(\widehat{\theta}(1), s_1) + \psi(\widehat{\theta}(1), s_1) - \frac{\partial \phi}{\partial \theta}(\widehat{\theta}(1), s_1) \frac{1 - F(\widehat{\theta}(1))}{f(\widehat{\theta}(1))} = 0$$

³¹For conciseness, we assume interior solutions ($\underline{\theta} < \widehat{\theta}(t) < \bar{\theta}$). The result however does not hinge on this assumption.

³² $\sum_{\tau=t}^{\tau=T} \delta^{\tau-t} \Lambda_\tau(\widehat{\theta}(t)) = 0 \implies \sum_{\tau=t+1}^{\tau=T} \delta^{\tau-(t+1)} \Lambda_\tau(\widehat{\theta}(t)) > 0$. Hence if $\widehat{\theta}(t+1) \geq \widehat{\theta}(t)$,

$$\sum_{\tau=t+1}^{\tau=T} \delta^{\tau-(t+1)} \Lambda_\tau(\widehat{\theta}(t+1)) > 0,$$

a contradiction.

³³It is unique if T is finite. We conjecture that it is also unique if T is infinite.

and

$$\sum_{\tau=0}^T \delta^\tau \left[\phi(\hat{\theta}(0), s_\tau) + \psi(\hat{\theta}(0), s_\tau) - \frac{\partial \phi}{\partial \theta}(\hat{\theta}(0), s_\tau) \frac{1 - F(\hat{\theta}(0))}{f(\hat{\theta}(0))} \right] = 0.$$

Because $\hat{\theta}(0) > \hat{\theta}(1)$, the principal's date-1 payoff in the neighborhood of $p_1 = \phi(\hat{\theta}(1), s_1)$ is, letting α_t denote the weight of cohort t and $\theta_1^*(p_1)$ be defined by $p_1 = \phi(\theta_1^*(p_1), s_1)$:

$$p_1 \left[\alpha_1 [1 - F(\theta_1^*(p_1))] + \alpha_0 [1 - F(\hat{\theta}(0))] \right] \\ + \alpha_1 \int_{\theta_1^*(p_1)}^{\bar{\theta}} \psi(\theta, s_1) f(\theta) d\theta + \alpha_0 \int_{\hat{\theta}(0)}^{\bar{\theta}} \psi(\theta, s_1) f(\theta) d\theta.$$

The derivative at $p_1 = \phi(\hat{\theta}(1), s_1)$ is strictly positive, reflecting the fact that the demand of cohort 0 is locally inelastic. Thus the principal cannot obtain the commitment payoff. ■

6 Time inconsistency

This section is devoted to the analysis of dynamic screening with positive selection in environments that do not satisfy the conditions for time consistency.

For the remainder of the paper, we will alleviate notation by making

Assumption 4 (no network externalities). $\frac{\partial \Gamma}{\partial \theta^*}(\theta, \theta^*, s) = 0$ for all (θ, s) .

By an abuse of notation, we omit the variable θ^* as an argument of ϕ and ψ . Note also that $\Gamma = \Lambda$ under Assumption 4. The absence of network externalities plays no major role in the results to come.

6.1 Shifting principal type

Sometimes the principal's preferences may change over time. Indeed, in the conversion game, Muslim rulers exhibited varying degrees of piousness, altering the trade-off between tax receipts and adherence to the Muslim faith. Let $\gamma_t \in \mathbb{R}$ denote the date- t principal's type, which is assumed to affect only the principal's objective function ψ and not the agent's utility ϕ . We assume that ψ is strictly increasing in γ , which will imply that a high γ principal prefers a lower cutoff compared with a low γ one.

We assume that the realizations of s_t and γ_t are public information at the beginning of date t ; otherwise the principal's price might signal his type. The parameters s_t and γ_t follow independent stochastic processes and differ in that the date- t principal's payoff from date- τ agent participation for $\tau > t$ is (under Assumption 4) $\psi(\theta, s_\tau, \gamma_t)$ as opposed to $\psi(\theta, s_\tau, \gamma_\tau)$ for the date- τ principal. That is, temporal variations in γ_t capture the change in the principal's preferences over time and will be the source of conflict among principals; by contrast, s_t is the

mere evolution of the part of the rest of the state. Thus, the date- t principal's objective function is:

$$V_t(s^t, \gamma^t) = \int_{\theta_{t-1}^*(s^{t-1}, \gamma^{t-1})}^{\bar{\theta}} E_{s^T, \gamma^T} \left[\sum_{\tau=t}^T \delta^{\tau-t} X^\tau(\theta, s^\tau, \gamma^\tau) \Gamma(\theta, s_\tau, \gamma_t) \left[\frac{f(\theta)}{1 - F(\theta_{t-1}^*(s^{t-1}, \gamma^{t-1}))} \right] d\theta \right]$$

Obviously, the principal's commitment policy in general will not be time consistent. Nonetheless simple equilibrium solutions again are available. Let us assume that the state $(s, \gamma) \in \mathbb{R}^2$, follows a first-order Markov process $G^s(s_{t+1}|s_t) \times G^\gamma(\gamma_{t+1}|\gamma_t)$ with full support. Suppose that $\partial\psi/\partial s > 0$ and that the virtual surplus

$$\Gamma(\theta, s, \gamma) = \phi(\theta, s) + \psi(\theta, s, \gamma) - \frac{\partial\phi(\theta, s)}{\partial\theta} \frac{1 - F(\theta)}{f(\theta)}$$

is strictly increasing in θ (this is Assumption 3 applied to this context). It is also increasing in γ because Γ is.

Suppose, first, that the “time-consistent” part of the state, s_t , is constant (only the principal's type varies over time). Intuitively, when inducing a cut-off at date t , type γ_t constrains, but is not affected by future choices of types $\gamma' > \gamma_t$. By contrast, he is affected by future choices of types $\gamma' < \gamma_t$, but cannot do anything about it (altering these choices would require making future cutoffs even higher, while they are already too high). For example, a muslim ruler cares about the exact religiousity of future rulers who will be more pious than he is, but not about that of less religious ones.

Proposition 8 (shifting principal type, invariant environment). *Suppose that only the principal's type changes over time: $s_t = s$ for all t , that the virtual surplus Γ is strictly increasing in θ , and make Assumption 4.*

Let θ_γ^* be defined by $\theta_\gamma^* = \underline{\theta}$ if $\Gamma(\underline{\theta}, s, \gamma) \geq 0$, $\theta_\gamma^* = \bar{\theta}$ if $\Gamma(\bar{\theta}, s, \gamma) \leq 0$ and $\Gamma(\theta_\gamma^*, s, \gamma) = 0$ otherwise. That is, θ_γ^* is the optimal cutoff for principal type γ . There exists a Markov perfect equilibrium of the game such that on the equilibrium path the cutoff is at each point of time the optimal cutoff for the least eager principal so far:

$$\theta_t^* = \theta_{\min_{\tau \leq t} \{\gamma_\tau\}}^*$$

Proof. Note that

$$M(\theta^*, \gamma) \equiv \int_{\theta^*}^{\bar{\theta}} \Gamma(\theta, s, \gamma) f(\theta) d\theta$$

is strictly quasi-concave with maximum at θ_γ^* . Furthermore θ_γ^* is weakly decreasing in γ .

Consider an arbitrary date τ and history $h^{\tau-1} \equiv (\gamma_0, \dots, \gamma_{\tau-1}, p_0, \dots, p_{\tau-1})$ at that date. Let $\hat{\theta}(p_0, \dots, p_{\tau-1})$ be defined by the solution to $\phi(\theta, s) = \max\{p_0, \dots, p_{\tau-1}\}$ (if interior; otherwise $\hat{\theta} = \underline{\theta}$ if $\phi(\underline{\theta}, s) \geq \max\{p_0, \dots, p_{\tau-1}\}$ and $\hat{\theta} = \bar{\theta}$ if $\phi(\bar{\theta}, s) \leq \max\{p_0, \dots, p_{\tau-1}\}$) Suppose

that at date τ , principal γ_τ sets

$$p_\tau = \phi \left(\max \{ \theta_{\gamma_\tau}^*, \hat{\theta}(p_0, \dots, p_{\tau-1}) \}, s \right),$$

and that the agent behaves myopically ($x_\tau(\theta, h^\tau) = 1$ if and only if $\phi(\theta, s) \geq p_\tau$). Consider the date- t principal, with type γ_t . Then

(i) either $\theta_{\gamma_t}^* \geq \hat{\theta}(p_0, \dots, p_{t-1})$ and then for all τ such that $\min \{ \gamma_{t+1}, \dots, \gamma_\tau \} \geq \gamma_t$, $M(\theta_{\gamma_t}^*, \gamma_t) \geq M(\theta_z^*, \gamma_t)$ for $z \in \{t+1, \dots, \tau\}$ and so $\theta_{\gamma_t}^*$ provides a higher utility for such realizations. By contrast, consider $(\tau > t, \gamma_\tau)$ such that $\gamma_\tau < \gamma_t$. Then any alternative cutoff $\theta_t^* \geq \theta_{\gamma_\tau}^*$ would have no impact on the date- τ cutoff. And if $\theta_t^* < \theta_{\min\{\gamma_{t+1}, \dots, \gamma_z\}}^*$ raising θ_t^* at the margin improves type γ_t 's welfare from the quasi-concavity of M .

(ii) or $\theta_{\gamma_t}^* < \hat{\theta}(p_0, \dots, p_{t-1})$. If $\hat{\theta}(p_0, \dots, p_t, p_{t+1}, \dots, p_\tau) = \hat{\theta}(p_0, \dots, p_t) < \theta_{\gamma_\tau}^*$ the quasi-concavity of M and the fact that $\hat{\theta}$ is weakly increasing in p_t implies that any increase in p_t above $\hat{\theta}(p_0, \dots, p_t)$ would reduce profit not only at date t but also at dates $t+1, \dots, \tau$. And if $\hat{\theta}(p_0, \dots, p_t, p_{t+1}, \dots, p_\tau) = \theta_{\min\{\gamma_{t+1}, \dots, \gamma_z\}}^*$, we are back to case (i). We thus conclude that the proposed strategies indeed form an equilibrium. ■

The next result allows the non-principal-related part of the state, s_t , to evolve over time, making the relationship either increasingly attractive or increasingly unattractive.

Let $\hat{\theta}_{t, \gamma_t}$ denote the optimal date- t cutoff for principal γ_t as characterized in Proposition 3. That is, the cutoff is that which would prevail in the thought experiment in which (a) the principal's type remains γ_t for the rest of the game and (b) this principal is unconstrained by previous truncations of the distribution ($\theta_{t-1}^* = \underline{\theta}$, say): Thus, $\hat{\theta}_{t, \gamma_t}$ is the cutoff that would prevail in a different game in which both the principal and the agent both believed that $\gamma_\tau = \gamma_t$ for all $\tau > t$.

Proposition 9 (shifting principal type, deterministic monotone attractiveness). *Suppose that $s_t \in \mathbb{R}$, $\partial\Gamma/\partial s > 0$, $\partial\Gamma/\partial\theta > 0$ and Assumption 4 holds. Then there exists an equilibrium and a sequence θ_{t, γ_t}^* such that the cutoff θ_t^* induced by principal γ_t at date t is $\max \{ \theta_{t-1}^*, \theta_{t, \gamma_t}^* \}$ where*

$$\theta_{t, \gamma_t}^* \geq \hat{\theta}_{t, \gamma_t} \quad \text{under} \quad \begin{aligned} &\text{either increasing attractiveness (} s_t \text{ increasing)} \\ &\text{or decreasing attractiveness (} s_t \text{ decreasing),} \end{aligned}$$

where $\hat{\theta}_{t, \gamma_t}$ is the cutoff that would be selected by a date- t principal with type γ_t , were the future principals also to have type γ_t .

Proposition 9 says that there is too little retention going forward from the point of view of all successive principals. Interestingly, the reason why this is so is not the same for increasing and decreasing demand. Under increasing attractiveness the principal exerts cutoff moderation when having a constant type as he expects that he will prefer wider participation in the future.

Cutoff moderation is like an investment, but with changing type, the investment has a lower value as the cutoff in future periods may be raised by less eager (lower γ) types. Under decreasing attractiveness, the date- t principal would myopically raise the cutoff over time if he were permanent (Proposition 3). Increasing the cutoff a bit above the myopic optimum is beneficial as this commits future, more eager (higher γ) types.

6.2 Multiple principals: retention as a dynamically provided public good

The agent's decision to exit often depends on the behavior of multiple principals rather than a single one. Retention in a work, volunteering, sports or religious community relies on the joint efforts by its members to make staying a comfortable option for the member. Immigration decisions similarly may be guided by a mixture of policies enacted by local and national authorities, workplace atmosphere, overall society openness, and so forth.

This section studies environments in which n principals each set a "price" every period for that period, and the agent's continuation decision is guided by the sum of those prices. Such environments are not conducive to time consistency since under commitment each principal might want to commit to relatively high prices in order to force other principals to bear the brunt of the retention effort in the future.³⁴

Suppose that there are n symmetrical principals with surplus $\psi(\theta, s_t)/n$ each (so as to keep total surplus the same). At date t , the principals simultaneously set prices p_t^i ; principal i 's flow payoff given resulting cutoff θ_t^i is then

$$p_t^i [1 - F(\theta_t^*)] + \int_{\theta_t^*}^{\bar{\theta}} \frac{\psi(\theta, s_t)}{n} f(\theta) d\theta.$$

Provided that he does not exit, the agent's flow payoff is

$$\phi(\theta, s_t) - \sum_{i=1}^n p_t^i.$$

We will be focusing on *symmetric Markov perfect equilibria in which the agent behaves myopically*:

$$\begin{aligned} \theta_t^* &= \theta_{t-1}^* & \text{if } \phi(\theta_{t-1}^*, s_t) &\geq \sum_{i=1}^n p_t^i \\ \theta_t^* &= \bar{\theta} & \text{if } \phi(\bar{\theta}, s_t) &\leq \sum_{i=1}^n p_t^i \end{aligned}$$

or, if the solution is interior:

$$\phi(\theta_t^*, s_t) = \sum_{i=1}^n p_t^i.$$

Markov behavior means that the vector of prices charged at date t , $\{p_t^i\}_{i=1}^n$ depends only on the previous cutoff θ_{t-1}^* and on the current state s_t (provided that the state follows a first-order

³⁴This environment is different from that studied by Admati and Perry (1991) and the literature they initiated. Admati and Perry consider a cumulative-contribution game in which n players make sequential commitments toward assembling a fixed amount needed to implement a project. There is no strategic agent involved, and a fortiori no screening of the agents' information.

Markov process, or more generally on a statistics for the history of states that is a sufficient statistics for describing current and future payoffs). Furthermore, $p_t^i = p_t$ for all i and all histories of the game.

The following assumption is the counterpart of Assumption 3 in the common agency context:

Assumption 5. For all (θ, s)

$$\frac{\partial}{\partial \theta} \left(\phi(\theta, s) + \psi(\theta, s) - n \frac{1 - F(\theta)}{f(\theta)} \frac{\partial \phi}{\partial \theta}(\theta, s) \right) > 0$$

We furthermore assume that states are ordered: $s \in \mathbb{R}$ and

$$\frac{\partial}{\partial s} \left(\phi(\theta, s) + \psi(\theta, s) - n \frac{1 - F(\theta)}{f(\theta)} \frac{\partial \phi}{\partial \theta}(\theta, s) \right) > 0.$$

Proposition 10 (common agency). Under Assumptions 4 and 5, a symmetric Markov perfect equilibrium with myopic agent behavior exists and has the following properties:

(i) Under deterministic increasing attractiveness, exit occurs only in the initial period: there exists θ^* such that $\theta_t^*(s^t) = \theta^*$ for all t . Furthermore

$$\sum_{t=0}^{t=T} \delta^t \left[\phi(\theta^*, s_t) + \psi(\theta^*, s_t) - n \frac{1 - F(\theta^*)}{f(\theta^*)} \frac{\partial \phi}{\partial \theta}(\theta^*, s_t) \right] = 0.$$

(ii) Under (possibly stochastic) decreasing attractiveness, the cutoffs $\theta_t^* = \theta_t^*(s_t)$ are increasing over time and satisfy for all (t, s_t) :

$$\phi(\theta_t^*, s_t) + \psi(\theta_t^*, s_t) = n \frac{1 - F(\theta_t^*)}{f(\theta_t^*)} \frac{\partial \phi}{\partial \theta}(\theta_t^*, s_t).$$

Proposition 10 can be viewed as a generalization of Cournot n^{th} marginalization and more generally the static common-agency-with-private-information literature³⁵ to dynamic games of exit/retention. When demand is increasing, all exit occurs in the first period, like in the single-principal case; retention is a collective investment and free riding implies that there is less retention than if the principals coordinated their price choices. When demand is decreasing by contrast, exit occurs progressively and the remaining installed base is determined by the static Cournot n^{th} marginalization condition. It again involves insufficient retention.

Proof. Let us first consider the “unconstrained optimization” at date t ; that is, one considers the thought experiment in which no exit has yet occurred at date t ($\theta_{t-1}^* = \underline{\theta}$). Of course, cutoff monotonicity is imposed from date t on. Let $\hat{\theta}_t(s^t)$ be defined like in Proposition 10, but for the game starting at t with no exit prior to date t ; $\hat{\theta}_t(s^t)$ satisfies:

$$\sum_{\tau=t}^{\tau=T} \delta^\tau \left[\phi(\hat{\theta}_t(s^t), s_\tau) + \psi(\hat{\theta}_t(s^t), s_\tau) - n \frac{1 - F(\hat{\theta}_t(s^t))}{f(\hat{\theta}_t(s^t))} \frac{\partial \phi}{\partial \theta}(\hat{\theta}_t(s^t), s_\tau) \right] = 0,$$

³⁵See Martimort and Stole (2015) for a state-of-the-art contribution to this literature (the retention game roughly corresponds to a dynamic extension of their “congruent preferences” case).

under deterministic increasing attractiveness, and

$$\phi(\hat{\theta}_t(s^t), s_t) + \psi(\hat{\theta}_t(s^t), s_t) - n \frac{1 - F(\hat{\theta}_t(s^t))}{f(\hat{\theta}_t(s^t))} \frac{\partial \phi}{\partial \theta}(\hat{\theta}_t(s^t), s_t) = 0$$

under decreasing attractiveness.

Consider the following Markov strategies. Principals all charge price $p_t^i = \phi(\theta_t, s_t)/n$ where $\theta_t = \max\{\theta_{t-1}^*, \hat{\theta}_t\}$. And the agent behaves myopically, as described just prior to the statement of the Proposition. Assumption 5 guarantees that provided that other principals charge $\phi(\theta_t, s_t)$, each principal's flow payoff is strictly quasi-concave in the cutoff.

Consider first decreasing attractiveness. The strict quasi-concavity of the instantaneous payoff in the cutoff implies that a deviation from the presumed price reduces the principal's current payoff; the deviation has no impact at date $\tau > t$ in state s_τ provided that the induced cutoff θ_t^* satisfies $\theta_t^* \leq \hat{\theta}_\tau(s_\tau)$. Furthermore if $\hat{\theta}_\tau(s_\tau) < \theta_t^*$, then the date- t deviation also reduces the date- τ payoff. The proof for increasing attractiveness follows the same steps. ■

6.3 Time inconsistency arising on agent side

Section 6.1 and 6.2 investigated environments in which time inconsistency originates on the principal's side: Either the principal's preferences change over time or there are multiple principals in each period. We leave the more complex cases of time inconsistency originating on the agent side for future research; we point out, however that the simplicity of positive selection environments makes us hopeful that interesting and tractable characterizations will be available.

The factors of time-inconsistency on the agent's side mirror those on the principal's side: the agent's tastes may change (here in an unobservable way) over time; and there may be multiple agents forming a team that will be dissolved if any of them quits.

Let us start this discussion with the former. The framework so far rules out the case of a "shifting type" for the agent (for which asymmetries of information may be reduced over time), on which many of the recent advances on dynamic mechanisms design have focused. Time consistency of the optimal commitment policy is then not to be expected. Intuitively, the principal might want to promise low (and efficient) future prices in exchange of a higher price today. However, price frontloading is not conducive to time consistency.³⁶

One therefore can no longer rely on solving an optimal control problem to obtain a perfect Bayesian equilibrium of the no-commitment environment. Nonetheless, explicit derivations may be available. One simple such case arises when the agent's type is redrawn in each period in an i.i.d. manner from a known distribution. Both the principal and the agent then have (stationary if $T = +\infty$ ³⁷) continuation option values. Appendix F shows that the derivation of the (unique)

³⁶Pavan et al (2014) stress, "because of the serial correlation of types, it is optimal to distort allocations not only in the initial period, but at every history at which the agent's type is responsive to his initial type, as measured by the impulse response function." This memorization of past, now-payoff-irrelevant types in the optimal commitment allocation makes the commitment solution time-inconsistent.

³⁷Appendix F focuses on the case of a Markov perfect equilibrium.

Markov perfect equilibrium is straightforward. Much work remains to be done, though, to study the more likely case of imperfect, but non-zero correlation over time, which has been the focus of the commitment literature.

Second, one can analyze the case of partnerships. A partnership dissolves if any of its members quits. The “Myerson-Satterthwaite platform”, whether benevolent or for profit, that arranges the conditions for the team of agents to operate, must then keep all members on board. Membership, though, is no longer a 0/1 decision as in the rest of the paper. From Myerson and Satterthwaite (1983), we know that efficient partnerships correspond to a set Θ of types that is not rectangular: whether type θ_i should be in depends on the other agents’ types. And therefore, each agent’s information about other agents’ types drawn from the mere continuation of the relationship necessarily varies with the agent’s own type.

The on-line appendix proves a limited, but nonetheless interesting result: there is no efficient and time-consistent allocation such that the agents learn at the end of each period only whether the relationship continues or not. The analysis of repeated Myerson-Satterthwaite relationships is a major item on the research agenda.

7 Alleys for future research

This paper provides first insights on repeated relationships with positive selection. The main ones were summarized in the introduction. This conclusion therefore focuses on future research. The first front is empirical: environments with positive selection are as untested as they are theoretically investigated; yet, this paper and subsequent research provides (will provide) clear empirical patterns that ought to guide empirical research in this area. Second, at least six broad areas of research on the theoretical front seem worth pursuing.

First, we saw that shifting types, common agency and partnerships all disconnect the resolution of the dynamic screening game from the simple optimization problem associated with commitment. While we showed that the simple structure of screening with positive selection allows for interesting characterizations, much work remains to be done in order to obtain general predictions for these environments. In particular, we have hardly scratched the surface when discussing environments in which time inconsistency originates on the agent’s side.

Second, the model should be generalized to allow for competition among principals. Firms compete for employees and consumers, department for professors, religions for followers, municipalities and countries for plants and headquarters, languages for speakers, and so forth, and principals and their agents are engaged as in this paper in relationships of endogenous lengths. A richer model would formalize not only the retention policies studied here (human resource management, customer relationship management, evolution of financial and non-financial terms), but also how mobility affects the policies of competing principals and the agents’ reservation utilities attached to splitting from their principal.

Third, the model could be enriched in several dimensions. It is hard to predict without

further inquiry whether these extensions will deliver insights that go beyond a mere combination of existing insights. But it seems for example worthwhile to add private information held by the principal. The commitment case would involve mechanism design by an informed principal, and the non-commitment case repeated signaling. One could then study the role of commitment in this enlarged framework. Similarly, we assumed (finite or infinite) re-entry costs to be exogenous. While this assumption may be reasonable in a number of contexts, one could also allow the principal or the agent to impact this re-entry cost.

Fourth, our model captures the dynamics of relationships when exit is absorbing due to large re-entry costs (as we have seen, “how large” re-entry costs must be depends on the evolution of the relationship’s attractiveness). More generally, one might have in mind that relationships must be cultivated; shared routines, a common history and understanding, learning by doing make relationships that have been activated in the past more attractive, although not in a discontinuous way as in this paper’s model. This could be captured by a state variable s_t that would depend not only on exogenous events, but also on the discounted intensity of previous relationships between the principal and the agent. This approach would yield an interesting characterization of second-degree price-discrimination dynamics.

Fifth, the extension of the analysis to multi-sided markets would enhance our theoretical and empirical understanding of platform dynamics. For example, platforms’ life-cycle (the two-sided extension of Section 4.2) would be worth of investigation. I conjecture, but have not verified, that as long as the platform can charge membership fees, the platform’s commitment outcome is time consistent under conditions similar to those derived in Proposition 1. It would also be interesting to analyse the case of pure usage fees and investigate whether cutoff myopia would still prevail then.

Finally, the paradigm should be enlarged to accommodate political economy considerations. In a number of environments (such as religions or firms), the principal’s preferences can be taken as exogenous in a first approximation. However, as Dewatripont and Roland (1992) stress, the principal’s preferences may result from a vote or power relationships, and therefore change with the composition of the in- and out-groups; for instance, religious conversions may affect the balance of political power and quits may have a long-lasting effect on the orientation of an academic department. Political economy considerations add a new form of (positive or negative) network externalities, which are intrinsically dynamic rather than contemporaneous.

References

- Admati, A., and M. Perry (1991) “Joint Projects without Commitment,” *Review of Economic Studies*, 58(2): 259–276.
- Aguirre, I., Cowan, S. and J. Vickers (2010) “Monopoly Price Discrimination and Demand Curvature,” *American Economic Review* 100(4): 1601–1615.
- Athey, S., and I. Segal (2013) “An Efficient Dynamic Mechanism,” *Econometrica*, 81(6): 2463–2485.
- Baron, D., and D. Besanko (1984) “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1: 267–302.
- Battaglini, M. (2005) “Long-Term Contracting with Markovian Consumers,” *American Economic Review*, 95(3): 637–658.
- Benoit, J.P., and V. Krishna (1985) “Finitely Repeated Games,” *Econometrica*, 53: 890–904.
- Bergemann, D., and J. Välimäki (2010) “The Dynamic Pivot Mechanism,” *Econometrica*, 78: 771–789.
- Bester, H., and R. Strausz (2001) “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case,” *Econometrica*, 69: 1077–1098.
- Board, S., and M. Pycia (2014) “Outside Options and the Failure of the Coase Conjecture,” *American Economic Review*, 104(2): 656–71.
- Boleslavsky, R., and M. Said (2013) “Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process,” *Review of Economic Studies*, 80: 1–34.
- Burdett, K., and M. Coles (2003) “Equilibrium Wage-Tenure Contracts,” *Econometrica*, 71: 1377–1404.
- Burdett, K., and D. Mortensen (1998) “Equilibrium Wage Differentials and Employer Size,” *International Economic Review*, 39: 257–276.
- Chatterjee, K., and W. Samuelson (1983) “Bargaining under Incomplete Information,” *Operations Research*, 31: 835–851.
- Coase, R. (1972) “Durability and Monopoly,” *Journal of Law and Economics*, 15: 143–49.
- Conlisk, J., Gerstner, E., and J. Sobel (1984) “Cyclic Pricing by a Durable Good Monopolist,” *Quarterly Journal of Economics*, 98: 489–505.
- Courty, P., and H. Li (2000) “Sequential Screening,” *Review of Economic Studies*, 67: 697–717.
- Cremer, J., and R. McLean (1985) “Optimal Selling Strategies under Uncertainty for a Discriminating Monopoly when Demands Are Interdependent,” *Econometrica*, 53: 345–361.

- Daley, D., and B. Green (2012) “Waiting for News in the Market for Lemons,” *Econometrica*, 80: 1433–1504.
- Dewatripont, M. (1989) “Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts,” *Quarterly Journal of Economics*, 104: 589–619.
- Dewatripont, M., and G. Roland (1992) “Economic Reforms and Dynamic Political Constraints,” *Review of Economic Studies*, 59: 703–730.
- Esö, P., and B. Szentes (2007) “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74: 705–731.
- Fuchs, W., and A. Skrzypacz (2013) “Costs and Benefits of Dynamic Trading in a Lemons Market,” mimeo.
- Fudenberg, D., Levine, D., and J. Tirole (1985) “Infinite Horizon Models of Bargaining with One Sided Incomplete Information,” in A. Roth (ed.), *Game Theoretic Models of Bargaining*, Cambridge University Press, pp. 73–98.
- Fudenberg, D., Levine, D., and J. Tirole (1987) “Incomplete-Information Bargaining with Outside Opportunities,” *Quarterly Journal of Economics*, 52: 37–50.
- Gerardi, D., and L. Maestri (2013) “Bargaining over a Divisible Good in the Market for Lemons,” mimeo.
- Gerardi, D., J. Hörner and L. Maestri (2014) “The Role of Commitment in Bilateral Trade,” mimeo.
- Hart, O., and S. Grossman (1986) “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy*, 94(4): 691–719 .
- Guesnerie, R. and J. Laffont (1984) “A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm,” *Journal of Public Economics*, 25(3): 329–369.
- Hart, O., and J. Tirole (1988) “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55: 509–540.
- Kruse, T., and P. Strack (2014) “Optimal Stopping with Private Information,” mimeo, <http://ssrn.com/abstract=2291937>
- Laffont, J.J., and J. Tirole (1987) “Auctioning Incentive Contracts,” *Journal of Political Economy*, 95: 921–937.
- Laffont, J.J., and J. Tirole (1990) “Adverse Selection and Renegotiation in Procurement,” *Review of Economic Studies*, 57: 597–625.
- Lucas, R., Prescott, E., and N. Stokey (1989) *Recursive Methods in Economic Dynamics*. Harvard University Press.

- Maestri, L. (2013) “Dynamic Contracting under Adverse Selection and Renegotiation,” mimeo, TSE.
- Martimort, D. and L. Stole (2015) “Menu Auctions and Influence Games with Private Information,” mimeo, PSE and University of Chicago.
- Myerson, R. (1981) “Optimal Auction Design,” *Mathematics of Operations Research*, 6(1): 58–73.
- Myerson, R., and Satterthwaite (1983) “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory*, 28: 265–281.
- Niedermayer, A., and A. Shneyerov (2014) “For-Profit Search Platforms,” *International Economic Review*, (55)3: 765–789.
- Osband, K. (1989) “Optimal Forecasting Incentives,” *Journal of Political Economy* 97: 1091–1112.
- Pavan, A., Segal, I., and J. Toikka (2014) “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82: 601–653.
- Riley, J., and R. Zeckhauser (1983) “Optimal Selling Strategies: When to Haggle, When to Hold Firm,” *Quarterly Journal of Economics*, 98, 267–287.
- Saleh, M. (2013) “On the Road to Heaven: Self-Selection, Religion, and Socioeconomic Status,” mimeo, IAST and TSE.
- Skreta, V. (2006) “Sequentially Optimal Mechanisms,” *Review of Economic Studies*, 73: 1085–1111.
- Skrzypacz, A. and J. Toikka (2013) “Mechanisms for Repeated Bargaining,” mimeo, Stanford and MIT.
- Sobel, J. (1991) “Durable Goods Monopoly with Entry of New Consumers.” *Econometrica*, 59(5): 1455–85.
- Stevens, M. (2004) “Wage-Tenure Contracts in a Frictional Labor Market: Firms’ Strategies for Recruitment and Retention,” *Review of Economic Studies*, 71: 535–551.
- Strulovici, B. (2013) “Contract Negotiation and the Coase Conjecture,” mimeo, Northwestern.

Appendix A. Proof of Lemma 1

Consider the price sequence \mathbf{p} defined in the statement of the lemma.

(i) Note first that this sequence leaves no rent to the lowest type: $U(\underline{\theta}) = 0$. We must show that if the cutoff delivered by the sequence of short-term prices are exactly the cutoffs $\widehat{\theta}$ that obtain under commitment it is indeed in individual agents' interest to behave in a way that gives rise to cutoffs $\widehat{\theta}$. Note first that for types $\theta > \widehat{\theta}_t(s^t)$, not exiting at date t in state s^t is a dominant strategy as they enjoy a strictly positive instantaneous surplus and can always exit later on. Now consider a type $\theta < \widehat{\theta}_t(s^t)$. Because the cutoff sequence is necessarily monotonic, $p_\tau > \phi(\theta, \widehat{\theta}_\tau(s^\tau), s_\tau)$ for all $\tau \geq t$ and so not exiting delivers a strictly negative payoff.

(ii) To demonstrate uniqueness under Assumption 1, suppose that there exists (t, s^t) such that $\theta_t^*(s^t) \neq \widehat{\theta}_t(s^t)$ (the equilibrium cutoff differs from the optimal cutoff). If $\theta_t^*(s^t) < \widehat{\theta}_t(s^t)$, then from Assumption 1, $\phi(\theta_t^*(s^t), \theta_t^*(s^t), s_t) - p_t(s^t) < 0$. Hence there must exist $(\tau > t, s^\tau \succ s^t)$ such that $x_\tau(\theta_t^*(s^t), s^\tau) = 1$ (and so $\theta_t^*(s^t)$ is still the cutoff at (τ, s^τ)) and

$$\phi(\theta_t^*(s^t), \theta_t^*(s^t), s_\tau) - p_\tau(s^\tau) > 0;$$

and so (from Assumption 1), $\theta_t^*(s^t) > \widehat{\theta}_\tau(s^\tau)$, contradicting the monotonicity of the optimal cutoff sequence. If $\theta_t^*(s^t) > \widehat{\theta}_t(s^t)$, then from Assumption 1,

$$\phi(\theta_t^*(s^t), \theta_t^*(s^t), s_t) - p_t(s^t) > 0.$$

Then there must exist $(\tau \leq t, s^\tau \preceq s^t)$ such that $\theta_{\tau-1}^*(s^{\tau-1}) < \theta_t^*(s^t)$ and $\theta_\tau^*(s^\tau) = \theta_{\tau+1}^*(s^{\tau+1}) = \dots = \theta_t^*(s^t)$. Furthermore,

$$\phi(\theta_\tau^*(s^\tau), \theta_\tau^*(s^\tau), s_\tau) - p_\tau(s^\tau) > 0$$

from Assumption 1 and the monotonicity of the sequence $\widehat{\theta}(\cdot)$. But then $\theta_\tau^*(s^\tau)$ cannot be the cutoff at (τ, s^τ) since it is a strictly dominant strategy for $\theta_\tau^*(s^\tau) - \varepsilon$ (for ε small and positive) to stay at (τ, s^τ) . ■

Appendix B. Proof of Proposition 2

For notational simplicity only we will assume that the path of the state s_t is deterministic.

(i) *Impatient agent* ($\delta_A \leq \delta$). Consider the optimal allocation when the principal can commit and the agent cannot. Let $\{\theta_t^*\}_{t=0, \dots, T}$ denote the sequence of optimal cutoffs and $\{p_t\}_{t=0, \dots, T}$ the contributions. Because the agent cannot commit, these cutoffs must satisfy

$$\begin{aligned} \phi(\theta_T^*, \theta_T^*, s_T) - p_T &\geq 0, \\ \phi(\theta_{T-1}^*, \theta_{T-1}^*, s_{T-1}) - p_{T-1} + \delta_A \max\{0, \phi(\theta_{T-1}^*, \theta_T^*, s_T) - p_T\} &\geq 0 \\ \dots \end{aligned}$$

Fix the cutoffs and optimize the principal's payoff (whose expression is the same as earlier, because we have taken δ to be the principal's discount factor) with respect to payments. Suppose that $\phi(\theta_T^*, \theta_T^*, s_T) > p_T$. Then, consider new payments $(\hat{p}_{T-1}, \hat{p}_T)$ such that

$$\hat{p}_{T-1} + \delta_A \hat{p}_T = p_{T-1} + \delta_A p_T$$

and

$$\hat{p}_T = \phi(\theta_T^*, \theta_T^*, s_T).$$

The new cutoffs satisfy $\hat{\theta}_T^* = \theta_T^*$ and $\hat{\theta}_{T-1}^* \leq \theta_{T-1}^*$. From Assumption 2 and $\delta_A \leq \delta$, the principal's payoff is increased. Repeat this reasoning; backward induction then shows that at the optimal allocation, the contributions satisfy:

$$p_t = \phi(\theta_t^*, \theta_t^*, s_t) \quad \text{for all } t.$$

Finally, the agent's discount factor is irrelevant under cutoff myopia as the cutoff's continuation valuation is always equal to 0. Hence $V^{nc}(\delta_A) = \hat{V}$.

(ii) *Patient agent* ($\delta_A > \delta$). Suppose that the principal cannot commit. Then any price $p_T < \phi(\theta_T^*, \theta_T^*, s_T)$ is strictly suboptimal; hence $p_T = \phi(\theta_T^*, \theta_T^*, s_T)$. By backward induction, cutoff myopia prevails on the equilibrium path, and so again δ_A is irrelevant. But the commitment solution is not time-consistent: The principal would like to frontload payments, which requires commitment, as already noted in the text. ■

Appendix C. Proof of Corollary 1

Suppose that $\psi(\theta, \theta^*, s) = -C_1(1 - F(\theta), s)$ and so:

$\int_{\theta^*}^{\bar{\theta}} \psi(\theta, \theta^*, s) f(\theta) d\theta = -C_1(1 - F(\theta^*), s) + C_1(0, s)$. Then Assumption 2 is equivalent to

$$\phi(\hat{\theta}_t(s^t), \hat{\theta}_t(s^t), s_t) \geq C_1(1 - F(\theta^*), s_t)$$

for all $\theta^* \leq \hat{\theta}_t(s^t)$. Now if $C_{11} \leq 0$, this is satisfied provided that

$$\phi(\hat{\theta}_t(s^t), \hat{\theta}_t(s^t), s_t) \geq C_1(1 - F(\hat{\theta}_t(s^t)), s_t).$$

When ϕ is separable ($\phi(\theta, \theta^*, s) \equiv \xi(\theta, s) + \nu(\theta^*, s)$), the condition $\Lambda(\widehat{\theta}_t(s^t), \widehat{\theta}_t(s^t), s_t) = 0$ (see part (ii) of Proposition 2) takes the following form at $\theta^* = \widehat{\theta}_t(s^t)$ and $s = s_t$:

$$\begin{aligned} \phi(\theta^*, \theta^*, s) - C_1(1 - F(\theta^*), s) &= \frac{1 - F(\theta^*)}{f(\theta^*)} \xi_1(\theta^*, s) + \frac{\int_{\theta^*}^{\bar{\theta}} \nu_1(\theta^*, s) f(\theta) d\theta}{f(\theta^*)} \\ &= \frac{1 - F(\theta^*)}{f(\theta^*)} [\xi_1(\theta^*, s) + \nu_1(\theta^*, s)] \geq 0 \end{aligned}$$

under (a) and (b). ■

Appendix D. Proof of Proposition 6

Let $\widehat{\theta}_t(s^t)$ denote the optimal cutoff for infinite re-entry costs.

(i) Under decreasing/constant attractiveness, when the principal charges $p_t = \phi(\widehat{\theta}_t(s_t), \widehat{\theta}_t(s_t), s_t)$ for all (t, s_t) with $\Lambda(\widehat{\theta}_t(s_t), s_t) = 0$ ($\widehat{\theta}_t$ here depends on s^t only through s_t) and so $\widehat{\theta}_t(s_t) \geq \widehat{\theta}_t(s_{t-1})$, the agent has no incentive to exit to later re-enter. The principal cannot obtain more than \widehat{V} , because even for $r = 0$ the commitment outcome is $\{\widehat{\theta}_t(s_t)\}_{t, s_t}$. To show uniqueness when T is finite and Assumption 1 is satisfied, note that the strict quasi-concavity of the instantaneous profit function (Assumption 3) and backward induction from the end of the horizon imply that for any finite re-entry cost, for all (t, s^t) , the principal induces cutoff $\theta_t^*(s^t) = \max\{\widehat{\theta}_t(s^t), \theta_{t-1}^*(s^{t-1})\}$ and charges $p_t = \phi(\theta_t^*(s^t), \theta_t^*(s^t), s_t)$.

(ii) By contrast, with increasing attractiveness, for r small, the principal might want the agent to exit and re-enter later on. The solution takes the following form (we write it for simplicity in the case of no network externalities; and as earlier we omit the state in the deterministic case and just index functions by time). Let $T(\theta)$ be defined by (if not equal to 0 or $+\infty$):

$$\Gamma_{T(\theta)}(\theta) \geq (1 - \delta)r > \Gamma_{T(\theta)-1}(\theta).$$

In words, $T(\theta)$ is type θ 's optimal re-entry date, if any. Re-entry however can be strictly optimal only if two conditions hold. First re-entry must be profitable relative to exit at date 0 and no re-entry:

$$\sum_{t=T(\theta)}^{t=T} \delta^t \Gamma_t(\theta) > \delta^{T(\theta)} r.$$

Second, it must be profitable relative to no exit at date 0:

$$-\sum_{t=0}^{t=T(\theta)-1} \delta^t \Gamma_t(\theta) > \delta^{T(\theta)} r.$$

So if

$$r \geq \max_{\{\theta\}} \left[\min \left\{ \sum_{t=T(\theta)}^{t=T} \delta^{t-T(\theta)} \Gamma_t(\theta), -\sum_{t=0}^{t=T(\theta)-1} \delta^{t-T(\theta)} \Gamma_t(\theta) \right\} \right] \equiv \underline{r},$$

the optimum involves no re-entry. The first term in the min is increasing in θ and the second

term decreasing.³⁸ Thus $\underline{r} \equiv \sum_{t=T(\hat{\theta})}^{t=T} \delta^{t-T(\hat{\theta})} \Gamma_t(\hat{\theta})$ where $\hat{\theta}$ is uniquely defined by $\sum_{t=0}^{t=T} \delta^t \Gamma_t(\hat{\theta}) = 0$. Note finally that even if $r < \underline{r}$, the commitment solution still resembles that of the absorbing exit paradigm: Types $\theta \geq \theta^*$ never exit (for some θ^*); types $\theta < \theta^*$ exit at date 0, and may re-enter, with higher types re-entering earlier. We assume that the discounted virtual surplus is uniformly bounded, otherwise the commitment payoff \widehat{V} may not exist. And so \underline{r} is well defined. ■

Appendix E. Proof of Proposition 9 (sketch)

(i) *Decreasing attractiveness.* Recall that the time-consistent cutoffs under decreasing attractiveness are given by myopic optimization, i.e., $\hat{\theta}_{t,\gamma_t}$ is given by

$$\Gamma(\hat{\theta}_{t,\gamma_t}, s_t, \gamma_t) = 0;$$

and that the sequence $\hat{\theta}_{t,\gamma_t}$ is monotonically increasing in t for a given γ_t .

Let the agent behave myopically: $x_t = 1$ iff $p_t \leq \phi(\theta, s_t)$. Then setting prices is equivalent to setting cutoffs (subject to the cutoff being no smaller than the previous one). We look for an equilibrium in which for all (t, γ_t)

$$\theta_{t,\gamma_t}^* \geq \max \{ \theta_{t-1}^*, \hat{\theta}_{t,\gamma_t} \}.$$

Given this, type γ_t setting cutoff $\theta_t^* < \hat{\theta}_{t,\gamma_t}$ at date t (assuming this is allowed by previous cutoffs) reduces the principal's date- t payoff from the strict concavity of M (the fact that Γ is increasing in θ). At a future date $\tau > t$, either θ_t^* is locally irrelevant ($\theta_t^* < \theta_\tau^*$) or $\theta_t^* = \theta_\tau^*$. Because $\hat{\theta}_{\tau,\gamma_t} \geq \hat{\theta}_{t,\gamma_t}$ and by strict quasi-concavity, raising θ_t^* slightly would also raise type γ_t 's payoff at date τ in such events.

The existence of cutoff $\{\theta_t^*\}$ is obtained through a fixed-point argument.

(ii) *Increasing attractiveness.* The strategy of proof is identical to that of case (i). Again, let the agent behave myopically, and the principal set a cutoff

$$\theta_{t,\gamma_t}^* \geq \max \{ \theta_{t-1}^*, \hat{\theta}_{t,\gamma_t} \}$$

where the time-consistent cutoff $\hat{\theta}_{t,\gamma_t}$ is no longer given by a myopic optimization (see Proposition 3). The strategy of proof again consists in using the strict monotonicity of the Γ function to show that setting a cutoff $\theta_t^* < \hat{\theta}_{t,\gamma_t}$ is strictly suboptimal for type γ_t at date t . ■

³⁸For the first term and using the envelope theorem, we know that $\sum_{t=T(\theta)}^{t=T} \delta^{t-T(\theta)} \Gamma(\theta, s_t) - r \delta^{T(\theta)}$ is increasing in θ and that $T(\theta)$ is decreasing in θ . Similarly, for the second term, $T(\theta)$ can be seen as the time minimizing $\sum_{t=0}^{t=T(\theta)-1} \delta^{t-T(\theta)} \Gamma(\theta, s_t) + r \delta^{T(\theta)}$.

Appendix F. Transient agent type

Proposition 11 (transient agent types). *Assume that $\phi(\theta, s) = \theta$ and $\psi(\theta, s) = 0$. Suppose that at $t = 0, 1, \dots, \infty$, the agent's type is drawn in an i.i.d. manner from density $f(\theta)$ and c.d.f. $F(\theta)$ on $[\underline{\theta}, \bar{\theta}]$ where $\underline{\theta} \geq 0$ and the virtual surplus $\theta - [(1 - F(\theta))/f(\theta)]$ is strictly increasing. Any Markov Perfect Equilibrium is characterized by a (uniquely defined and increasing in δ) threshold θ^* given by the following generalized virtual surplus:*

$$J(\theta^*) \equiv \theta^* - \frac{1 - F(\theta^*)}{f(\theta^*)} + \frac{\delta \int_{\theta^*}^{\bar{\theta}} \theta dF(\theta)}{1 - \delta[1 - F(\theta^*)]} = 0 \quad (1)$$

(if interior; $\theta^* = \underline{\theta}$ if $J(\underline{\theta}) \geq 0$ and $\theta^* = \bar{\theta}$ if $J(\bar{\theta}) \leq 0$)

The principal in each period sets price $p_t = \theta^* + \delta \int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) dF(\theta)$ conditional on the agent not having exited yet. Letting θ^m denote the monopoly price (i.e., $\theta^m = [1 - F(\theta^m)]/f(\theta^m)$), then $\theta^* \in [0, \theta^m)$.

Proof. Let U , V and W denote the continuation payoffs (U for the agent, V for the principal and $W \equiv U + V$). These are constant in a Markov Perfect Equilibrium, since the only payoff-relevant state variable is that the agent has not exited yet.

We treat only the case of an interior solution (the treatment of the corner solutions $\theta^* = \underline{\theta}$ or $\bar{\theta}$ is analogous). Price p^* at date t induces a cutoff θ^* given by

$$\theta^* - p^* + \delta U = 0.$$

The principal solves

$$\max_{\theta^*} \{[1 - F(\theta^*)][(\theta^* + \delta U) + \delta V]\}$$

which yields

$$\theta^* - \frac{1 - F(\theta^*)}{f(\theta^*)} + \delta W = 0$$

with the continuation welfare given by

$$W = \int_{\theta^*}^{\bar{\theta}} \theta dF(\theta) + [1 - F(\theta^*)]\delta W.$$

Simple computations show that $J'(\theta^*) > 0$ whenever $J(\theta^*) = 0$. Hence the solution θ^* is unique. The price p^* is given by

$$p^* = \theta^* + \delta U$$

where

$$U = \int_{\theta^*}^{\bar{\theta}} [\theta - p^* + \delta U] f(\theta) d\theta = \int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) f(\theta) d\theta.$$

■

The commitment solution for the environment described in Proposition 11 can be implemented by a commitment to a sequence of prices:

$$p_t^c = 0 \quad \text{for all } t \geq 1 \quad \text{and} \quad p_0^c = \theta^c + \frac{\delta E[\theta]}{1 - \delta},$$

where either

$$\theta^c - \frac{1 - F(\theta^c)}{f(\theta^c)} + \frac{\delta E[\theta]}{1 - \delta} = 0 \quad \text{or} \quad \theta^c = \underline{\theta} \quad \text{and} \quad \frac{\delta E[\theta]}{1 - \delta} \geq \frac{1}{f(\underline{\theta})} - \underline{\theta}.$$

Note that

$$\theta^c \leq \theta^*$$

(with equality only when $\theta^* = \underline{\theta}$) and that for each $t \geq 1$

$$p_t^c \leq p^* \leq p_0^c$$

with strict inequalities whenever $\theta^* > \underline{\theta}$.

Relatedly, it is important for the analysis in Sections 3 and 4 that, in each period t , the state of nature be public knowledge prior to price setting (or, if not, that the principal be able to offer a state-contingent price p_t). Suppose by contrast that at each date t the principal first learns the realization of s_{t-1} (either directly or through the date- $(t-1)$ realized demand); the principal sets a price p_t ; the agents then observe s_t and decide whether to consume. The date- t shock then plays a role similar to that of a transient shock to the agent's type, in that it confers an informational advantage to the agent for exactly one period.³⁹

³⁹To make the basic point in the simplest manner, suppose that $t = 0, 1$; that $s_t \in \mathbb{R}$; that $\phi = s_t$ (homogeneous preferences and no network externality); that $\psi = 0$ (costless production); and that s_t is i.i.d. with distribution G .

At date $t = 1$, the principal sets $p^m = \arg \max_p [1 - G(p)] \equiv \pi(p)$. Let $S(p) \equiv \int_p^\infty (s - p) dG(s)$.

In the absence of commitment, the principal chooses p_0 so as to solve:

$$\max_{\{p_0\}} \left\{ [1 - G(p_0 - \delta S(p^m))] [p_0 + \delta \pi(p^m)] \right\} = \max_{\{s_0^*\}} \left\{ [1 - G(s_0^*)] [s_0^* + \delta [\pi(p^m) + S(p^m)]] \right\}.$$

By contrast, the commitment outcome corresponds to the solution of

$$\max_{\{s_0^*\}} \left\{ [1 - G(s_0^*)] [s_0^* + \delta S(0)] \right\}$$

since $S(0) = \max_{\{p\}} [\pi(p) + S(p)]$. Under commitment, the principal charges a higher date-0 price and has a larger date-0 clientele.