THÈSE de DOCTORAT

de l'UNIVERSITÉ TOULOUSE CAPITOLE

Présentée et soutenue par

Monsieur Emil Tobias MORTENSEN

Le 12 décembre 2025

Essais de macroéconomie

École doctorale : **TSE Toulouse Sciences Economiques**
Spécialité : **Sciences Economiques - Toulouse**
Unité de recherche : **TSE-R - Toulouse School of Economics - Recherche**

Thèse dirigée par *Monsieur Fabrice COLLARD*

Composition du jury

Rapporteur : M. Martial DUPAIGNE
Rapporteur : M. Edouard CHALLE
Examinateur : M. Christian HELLWIG
Examinateur : M. Patrick FèVE
Directeur de thèse : M. Fabrice COLLARD

Essays in Macroeconomics

Ph.D. Thesis

Emil Mortensen[1]

Toulouse School of Economics

August 2025

1. Address: 1, Esplanade de l'Université, 31080 Toulouse, France; email: emitm@outlook.dk

Acknowledgments

I would like to first express my deepest gratitude to Fabrice Collard, Christian Hellwig, and Eugenia Gonzalez-Aguado for their guidance and support as my professors and advisors. Their insightful advice has greatly enhanced the quality of my work, and their patience has been truly remarkable. I owe thanks to Nicolas Werquin and Martial Dupaigne for valuable suggestions and their availability.

I am also grateful to Patrick Fève, Andreas Schaab, Charles Brendon, and Tiziana Assenza for their valuable feedback and stimulating discussions. I would like to warmly thank every member of the macro group at TSE for the support they provided throughout my PhD. I have truly benefited from not just one advisor, but a group of exceptionally gifted individuals who supported me at different stages of my journey. A special thanks to Nour Meddahi for encouragement during my PhD.

I cannot thank my coauthors enough for making the process of writing each chapter so enjoyable. I am deeply grateful to Javier Gonzalez-Morin. While working on my first chapter, I greatly benefited from his sharp intellect, always accompanied by genuine humility. I will fondly remember not only the many hours we spent working together but also the moments spent watching cycling, whether in his office or along the streets of Toulouse. Chenchen Huang also deserves special thanks for his extensive involvement in my second chapter. His assistance in transforming the paper from an idea into a chapter was indispensable.

I would also like to acknowledge all of my fellow Ph.D. students, with particular thanks to my cohort. To Pablo Mileni, for our endless conversations about economics and other weighty matters. I have greatly valued the opportunity to exchange ideas with you. I can think of few people whose beliefs I have tried to challenge more, yet I recognize that your convictions are at least as steadfast as mine. I am also extremely grateful to

Juan Pal and Guillem Foucault for sharing the Ph.D. experience with me, and for using COVID as an excuse to do stupid things. To Giovanni, Pascal, Xi, Hanlin, Yassin, and Valentina, thank you for the many wonderful moments we shared. To Wenxuan Xu, for endless hours of discussions.

Finally, I would like to thank my friends, my family, and En Qi. Everybody has contributed in different ways, but two individuals stand out as especially important in my PhD journey. My granddad has nurtured my curiosity and love of learning from an early age, and taught me that engaging with different viewpoints can deepen understanding and spark even greater curiosity. En Qi's love and support are treasures I cherish every day, and I am profoundly grateful to her for showing me that the world extends far beyond Europe.

Summary

In my PhD thesis, I have explored a variety of issues in macroeconomics. The common thread throughout this work has been my passion for the economic method and interesting economic problems rather than a specific topic.

In Chapter 1, I study, together with Javier Gonzalez-Morin, the effect of forward guidance in the New Keynesian model commonly used to analyze the impact of monetary policy. In the standard model, forward guidance is excessively expansionary : the effect on current output grows without bounds in the promised length of a zero-interest-rate path – the so-called "forward guidance puzzle." We propose augmenting the model with limited commitment and reputation concerns and find that, once these considerations are taken into account, the effect of forward guidance remains bounded but positive. We then show, among other things, that the persistence of demand shocks has a non-monotonic impact on the effectiveness of forward guidance.

In Chapter 2, I jointly with Chenchen Huang study the distributional consequences of a property value tax. We propose a model where households have non-homothetic preferences over housing and non-durable consumption, and we find that a higher property value tax increases wealth inequality while reducing consumption inequality. While this may appear puzzling at first, this is because higher property taxes discourage future home ownership for households with low wealth, and they respond by reducing savings and increasing non-durable consumption.

In Chapter 3, I investigate whether gender inequality in labor markets emerges before parenthood due to expectations about caregiving responsibilities. I develop a theoretical model where women face higher caregiving risk and firms cannot observe workers' future caregiving roles. Under asymmetric information, firms hire fewer women into high-productivity, inflexible jobs, creating an inefficient glass ceiling. The model highlights how statistical discrimination and caregiving expectations jointly reinforce gender disparities.

Résumé

Dans le cadre de ma thèse de doctorat, j'ai exploré diverses questions en macroéconomie. Le fil conducteur de ce travail n'est pas un sujet spécifique, mais plutôt mon intérêt pour la méthode économique et pour l'étude de problèmes économiques stimulants.

Dans le chapitre 1, avec Javier Gonzalez-Morin, nous analysons l'effet de la forward guidance dans le modèle keynésien moderne utilisé pour étudier l'impact de la politique monétaire. Dans sa version standard, la forward guidance produit des effets excessivement expansifs : l'effet sur la production actuelle croît indéfiniment avec la durée promise d'un chemin de taux d'intérêt nul, phénomène connu sous le nom de « puzzle de la forward guidance ». Nous proposons d'enrichir le modèle en y intégrant des contraintes de commitment limité et des préoccupations de réputation. Nous montrons alors que l'effet de la forward guidance reste positif mais borné. Par ailleurs, nous démontrons que la persistance des chocs de demande influence de manière non monotone l'efficacité de cette politique.

Dans le chapitre 2, avec Chenchen Huang, nous étudions les conséquences distributives d'une taxe sur la valeur immobilière. Nous construisons un modèle dans lequel les ménages présentent des préférences non homothétiques pour le logement et la consommation non durable. Nos résultats indiquent qu'une augmentation de la taxe immobilière accroît les inégalités de richesse tout en réduisant les inégalités de consommation. Ce résultat, apparemment paradoxal, s'explique par le fait que des taxes plus élevées découragent l'accession future à la propriété pour les ménages modestes, qui réagissent en réduisant leur épargne et en augmentant leur consommation non durable.

Enfin, dans le chapitre 3, j'étudie l'émergence des inégalités de genre sur le marché du travail avant la parentalité, en lien avec les attentes concernant les responsabilités de soins. Je propose un modèle théorique où les femmes sont confrontées à un risque de charge de soins plus élevé, que les entreprises ne peuvent observer. Dans ce contexte d'information asymétrique, les entreprises recrutent moins de femmes pour des postes

à haute productivité et peu flexibles, générant ainsi un plafond de verre inefficace. Le modèle illustre comment la discrimination statistique et les attentes liées aux soins se renforcent mutuellement, accentuant les disparités de genre.

Contents

List of Figures

List of Tables

Chapter 1

Forward Guidance with Limited Commitment and Reputation

Javier Gonzalez-Morin [1] & Emil Mortensen [2]

Abstract

We argue that reputational concerns and limited commitment help overcome the forward guidance puzzle. Specifically, we study monetary policy in a New Keynesian model with a recurrent Zero Lower Bound where 1) the central bank is unable to commit to future policy, and 2) has private information about its discount factor giving rise to type-based reputation. Limited commitment imposes an upper bound on the length of sustainable promises while reputational concerns imply that forward guidance remains a useful policy tool at the Zero Lower Bound. Quantitative exploration shows that there is a non-monotonic relationship between the potency of forward guidance and the persistence of the Zero Lower Bound when the central bank cannot commit, while it is monotonically decreasing with full commitment. We find that the optimal length of forward guidance is 5 quarters with limited commitment compared to just 1 quarter with full commitment.

Keywords: The forward guidance puzzle, limited commitment, reputation
JEL classification: E58, E52

1. Toulouse School of Economics. Email: javier.gonzalez-morin@tse-fr.eu
2. Toulouse School of Economics. Email: emitm@outlook.dk. We are grateful to Fabrice Collard, Christian Hellwig and Sara Shahanaghi for providing valuable feedback on earlier versions of the paper as well as for helpful comments. We also thank Patrick Fève, Eugenia Gonzalez-Aguado, Andreas Schaab, and participants at the TSE Macro Workshop, TSE Theory Workshop, 2023 “Challenges of Europe", 2023 EEA-ESEM, and 2023 SMYE conferences.

1.1 Introduction

The welfare cost of the Zero Lower Bound (ZLB) depends on the effectiveness of the unconventional monetary policy tools that the central bank can resort to, among which forward guidance. In the baseline New Keynesian model, forward guidance is excessively expansionary: the effect on current output grows without bounds in the promised length of a zero-interest rate path – the so called "forward guidance puzzle" Del Negro, Giannoni, and Patterson (2012). Hence, any recession can be abated by a sufficiently generous promise about future interest rates McKay, Nakamura, and Steinsson (2016b).

In this paper, we take seriously one important, but overlooked, aspect of forward guidance: time inconsistency. We abandon the strong assumption that central banks can commit to a time-inconsistent policy, and we assume that forward-looking households understand the incentives of the central bank to keep its promise. We argue that this is sufficient to overcome the forward guidance puzzle.

We make two departures from a baseline New Keynesian model with recurrent ZLB events. First, we add limited commitment to capture the idea that central banks have an incentive to promise an expansionary future path of interest rates at the ZLB, but as soon as the economy leaves the ZLB, they would like to fulfill their actual mandate instead of sticking to the promised path of interest rates. Second, we add type-based reputation that we model by introducing private information of the central banker about her preferences. We assume that a central banker is only in charge of the central bank for a finite period, while the central bank is infinitely-lived. The central banker can be of two types: she may either be 1) a forward-looking one who wants to maximize the lifetime utility of infinitely-lived households, or 2) a myopic central banker who only cares about the utility of the households when she is in charge of the central bank. Modeling type-based reputation allows us to study how the effectiveness of forward guidance evolves over time as a function of the central bank's past policy decisions. In this environment, we first derive some theoretical properties of forward guidance under limited commitment. We then develop new insights by exploring the effect of forward guidance under limited commitment quantitatively, and we see how this effect compares to a model with full commitment.

In our model, the key decision faced by a central banker is whether to honor a past promise to keep interest rates low. The central banker trades off the cost of an excess expansion upon exiting the ZLB to honor a past promise against signaling its willingness to keep promises, which builds reputation. When ZLB events are recurrent, reputation is valuable to the central bank as it determines the effectiveness of forward guidance, and hence the severity of recessions in future ZLB episodes. An economy with a myopic central banker will not face an excess expansion upon exiting the ZLB while a forward-looking central banker may be willing to incur this cost to gain reputation.

We prove that the effect of forward guidance is less potent once agents factor in the incentives of the central bank to keep its promise. In particular, we show that there are realizations of the demand shock such that no promise is ever kept. We further prove that there exists a maximum duration of forward guidance that can be sustained in equilibrium. Any promise of a longer duration is not credible and hence futile. As a corollary, the marginal effect of increasing the promised duration of zero-interest is negative at this length.

We then calibrate the model to the US economy to explore the quantitative features of our model. We find that the effect of forward guidance is much mitigated once agents factor in the incentives of the central bank to keep its promises. The effect on output gap of promising one quarter of forward guidance is only half as powerful under limited commitment compared to full commitment. As agents realize that more expansionary promises are harder to keep, this mitigation is increasing in the promised duration of keeping the interest rate at zero. For instance, the effect of promising forward guidance for three quarters is a 0.55 percentage points increase in the output gap under limited commitment while it is 1.4 percentage points under full commitment. For our calibration, the maximum sustainable length of forward guidance is five quarters, which is also the optimal duration of forward guidance, while the optimal length is one quarter with full commitment.

The incentives to keep a promise are linked to the persistence of ZLB episodes, resulting in a non-monotonic relationship between this persistence and the effect of forward guidance. More persistent ZLB episodes decrease the probability of exiting to a more expansionary state where interest rates are still kept at zero. This reduces the effect of forward guidance. Yet, the persistence of the ZLB under limited commitment also determines the value of having a high reputation. The more persistent the ZLB and therefore the more severe the drop in output and inflation, the more important it is for the central bank to have a high reputation implying that they will be more likely to honor its past promise. With full commitment, this latter effect is absent and therefore the effectiveness of forward guidance is monotonically decreasing in the persistence of the ZLB.

We then move on the compare our paper with the discounted Euler literature which is one of the more prominent attempts of solving the forward guidance puzzle. The main difference is that attenuation in our paper works though expected interest rates not responding fully to announcements about forward guidance while fully credible changes in expectations are still as potent as in the standard New Keynesian model. One possible way to distinguish our paper from that literature would be to look at how the yield curve responds to policy announcements. In our framework, we would not expect the yield curve to respond perfectly to forward guidance which is consistent with the results in Campbell et al. (2019) that the FED has only limited influence on expected future interest at long horizons. We also relate our paper to Bhattarai, G. B. Eggertsson, and Gafarov (2023) which show that quantitative easing serves as a commitment device for forward guidance and hence is complementary. We find that these two policies can actually be substitutes as quantitative easing is already reducing the expected future interest rates.

The paper is organized as follows. Section 1.2 relates our paper to the existing literature. Section 1.3 introduces forward guidance in a New Keynesian model augmented with limited commitment and reputation. Section 1.4 presents the main results. Section 1.5 explores some quantitative features of the model. Section 1.6 relates our model to the literature on the discounted Euler equation. Section 1.7 extends the model to study the interaction between forward guidance and quantitative easing. Section 1.8 concludes. Appendix 1.C contains all the proofs.

1.2 Related Literature

Our paper is related to four strands of literature: forward guidance, time inconsistency of monetary policy, reputation, and loose commitment.

Forward guidance has been the subject of a large and growing literature. G. Eggertsson and Woodford (2003) shows that the possibility of engaging in forward guidance may significantly reduce the welfare cost of the ZLB which is consistent with the view expressed in Bernanke (2020), while Bilbiie (2019) solves for the optimal level of forward guidance in a New Keynesian model with full commitment. Nevertheless, the large impact of forward guidance predicted by the plain vanilla NK model has been questioned by Del Negro, Giannoni, and Patterson (2012), which has given rise to the well-known "forward guidance puzzle." This puzzle has spurred a literature trying to solve it where a common approach has been to mute the importance of expected future output on current output by assuming some kind of bounded rationality as in Farhi and Werning (2019), Gabaix (2020) and García-Schmidt and Woodford (2019) or lack of common knowledge as in Angeletos and Lian (2018). Another approach is to assume incomplete markets as McKay, Nakamura, and Steinsson (2016b), yet the degree to which market incompleteness provides a solution to the forward guidance puzzle depends on the cyclicality of risk as remarked by Werning (2015), and elaborated further by Acharya and Dogra (2020). One common implication of these papers is that they give rise to a "discounted" Euler equation which attenuates the effect of higher expected future output. Nakata et al. (2019) show that this implies the prediction that to obtain the same effect of forward guidance as in the standard model, interest rates should be kept low for longer time.

Since the seminal contribution of Kydland and Prescott (1977) and Barro and Gordon (1983), it has been known that optimal policies may be time inconsistent, which implies that there are gains to commitment. Forward guidance is a perfect example of this. The ability of committing to a low future interest rate at the ZLB can enhance welfare but the policy is time-inconsistent: once the economy exits the zero lower bound, the central bank would be better off by setting a higher interest rate than promised to stabilize current economic outcomes. Related to this observation, Barthélemy and Mengus (2018) argue that central banks should increase inflation prior to the ZLB to make forward guidance more credible while Bhattarai, G. B. Eggertsson, and Gafarov (2023) shows how quantitative easing can serve as a commitment device to keep future interest rates low. Following Stokey (1991) and Chari and Kehoe (1990), we study forward guidance in a model where central banks cannot commit to future policy, and the usefulness of forward guidance therefore depends on the extent to which it constitutes a sustainable plan in the sense that honoring past promises is sequentially optimal. Similar to Walsh (2018) and Nakata (2018), we exploit the recurrence of ZLB events, which provides incentives to honor past promises.

Another strand of literature that is relevant to our paper is reputation as pioneered by Kreps and Wilson (1982) and Milgrom and Roberts (1982). King and Lu (2022) estimate how the reputation of the FED has evolved assuming that the FED has private information about its discount factor, and Lu, King, and Pasten (2016) solves for optimal reputation building in a New Keynesian model. Dovis and Kirpalani (2020) study a model where the government can be of a commitment type or an optimizing type. The optimizing type chooses policy sequentially which also happens to be the case in our framework for the myopic type. Contrary to them, we don't assume that some type can perfectly commit, and the subject of study is different as they study how the effect of fiscal rules depend on the reputation of the government. Loisel (2008) shows how the inflation bias and the stabilization bias can be overcome by a reputation-concerned central bank if the punishment length is at least a few years. This can induce the central bank to implement otherwise time-inconsistent optimal monetary policy in the NK model.

Finally, our paper is also related to the literature on loose commitment as developed in Debortoli and Nunes (2010), Debortoli and Lakdawala (2016), and Debortoli, Maih, and Nunes (2014). These papers assume that the policy maker re-optimizes each period with some exogenous probability, while in our model the probability that a promise is kept is endogenous to the promise given. This endogeneity is key to solving the forward guidance puzzle.

The paper which is most closely related to ours is Nakata (2018), which uses a trigger strategy to sustain the optimal Ramsey plan in the tradition of Chari and Kehoe (1990). Our main departure from Nakata (2018) is that credibility arises as an equilibrium outcome of a signaling game with reputation *à la* Kreps and Wilson (1982) and Milgrom and Roberts (1982). Consequently, while the effect of forward guidance is constant in Nakata (2018), it varies over time on the equilibrium path in our model. This arises as the probability that a promise is kept is an equilibrium object that varies over time. This probability increases upon promise keeping, which is not the case in Nakata (2018), and it decreases upon deviations from the promise.

1.3 A Model of Forward Guidance

We consider a plain vanilla New Keynesian model where demand shocks can bring the economy to the ZLB. Following G. Eggertsson and Woodford (2003) the central bank engages in forward guidance once the nominal interest rate reaches the lower bound of zero. We first define this policy with full commitment, and we then re-define it under the assumption of limited commitment. Forward guidance under limited commitment is modeled as a reputational game where the central banker can be of two types: myopic or forward-looking. Uncertainty about the discount factor captures the idea that forward guidance is a policy entailing a dynamic trade off and hence valued differently by the two types of central bankers. One natural way of thinking about these types is that there is uncertainty about how accommodative the central banker is to pressure from politicians to stabilize current economic outcomes.

1.3.1 New Keynesian model with a ZLB

We consider a standard New Keynesian model summarized by the New Keynesian Phillips curve and the IS curve following the text book treatment by Woodford (2003) and Galí (2015).

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa y_t \tag{PC}$$

$$y_t = \mathbb{E}_t[y_{t+1}] - \frac{1}{\eta}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \tag{IS}$$

where κ is the slope of the New Keynesian Phillips curve, i.e. it measures the sensitivity of inflation to the output gap, and η is the coefficient of relative risk aversion. The natural rate of interest (r_t^n) follows a persistent stochastic process with bounded support. Following the New Keynesian literature, we interpret a change to this rate as a demand shock.

As is standard in the New Keynesian literature, we need to close the model by making an assumption about how the interest rate is determined. We assume that monetary policy is determined by a central banker with the following quadratic loss function

$$L\big(\{y_t\}_{t=0}^{\infty}, \{\pi_t\}_{t=0}^{\infty}\big) = \sum_{t=0}^{\infty} \delta^t \big(\lambda y_t^2 + \pi_t^2\big), \quad \lambda = \frac{\kappa}{\theta}.$$

where the notation follows Woodford (2003), except that we allow the central banker to have a discount factor δ that may differ from that of the representative household. When $\delta = \beta$, this loss function represents an approximation to the level of expected utility of the household under some regularity conditions for an appropriately chosen weight on the output gap (λ).[3]

We assume that a central banker is only in charge for one period, and that she cannot commit future central bankers. As the cental banker cannot commit to future policy, she faces a sequence of static problems and therefore solves the following problem (MP):

$$\min_{i_t, y_t, \pi_t} \lambda y_t^2 + \pi_t^2 \tag{MP}$$

subject to

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa y_t \tag{PC}$$

$$y_t = \mathbb{E}_t[y_{t+1}] - \frac{1}{\eta}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \tag{IS}$$

$$i_t \geq 0 \tag{ZLB}$$

Where we impose that the central bank is constrained to choose a non-negative nominal interest rate, which gives rise to ZLB events. We notice that the actual discount factor of the central banker does not enter in the loss minimization problem due to the lack of commitment.

1.3.2 Forward Guidance with Full Commitment

Once the constraint $i_t \geq 0$ is binding (i.e. the ZLB is hit), we assume that the central bank engages in forward guidance with full commitment which we define as a *binding* promise on future interest rates. This promise specifies a commitment to keeping the interest rate at zero for τ^{FC} periods.

3. The conditions are that the disturbances are small enough, the outcomes are close enough to the allocation around which the expansion is taken, and finally that distortions are small enough Woodford (2003).

Definition 1.3.1. *Forward guidance with full commitment is an exogenous integer $\tau^{FC} \in \mathbb{N}$ that prescribes the binding duration of a zero nominal interest rate path after exiting the ZLB.*

We thus consider a form of Odyssean forward guidance where the central bank gives a promise about future policy rates rather than Delphic forward guidance where the central gives a prediction about future output and inflation. We define the effect of forward guidance as follows:

Definition 1.3.2. *Let $y_{ZLB}(\tau^{FC}, r_t^n)$ be the level of the output gap at the ZLB under forward guidance. The impact response of output to forward guidance is $IR_y(\tau^{FC}, r_t^n) = y_{ZLB}(\tau^{FC}, r_t^n) - y_{ZLB}(0, r_t^n)$. Forward guidance is called effective if $IR_y(\tau^{FC}, r_t^n) > 0$.*

Hence, the impact response is the change in output at the ZLB when there is a promise of keeping the interest at zero for τ^{FC} periods.

A well-known result emphasized by for instance Del Negro, Giannoni, and Patterson (2012), Farhi and Werning (2019) and García-Schmidt and Woodford (2019) is that the response of the output gap grows without bounds in τ^{FC}. This has become known as the "forward guidance puzzle." The IS curve illustrates the root of this puzzle. Future increases in expected output increase one-to-one output today. Hence, any changes in future interest rates that are believed with probability one can have arbitrarily large effects on current output.

1.3.3 Reputation and Limited Commitment

We add type-based reputation and limited commitment to the New Keynesian model with a ZLB. To model type-based reputation, we assume that the central banker may have a different discount factor than households as the central banker faces a finite mandate as she is replaced each period while households live forever. In particular, the discount factor of the central banker is $\delta \in \{0, \beta\}$. We assume that this discount factor is private information to the central banker with common prior $\mu_0 \in \Delta(\{0, \beta\})$.

In words, we imagine that monetary policy is conducted by a central banker who can be of two types: 1) she can be a myopic one who only cares about consumer welfare during her mandate, or 2) she can be a forward-looking central banker who cares about consumer welfare even when she is not in charge. We assume that a mandate lasts for one period causing the myopic central banker to have a discount factor of zero while the forward-looking one has a discount factor of β.

We suppose that the central banker is replaced by someone with the same type as herself with probability $1-\gamma$, while she is replaced by someone with a different type with probability $\gamma < \frac{1}{2}$. This gives rise to stochastic and persistent types, which is meant to capture stability of the central bank as an institution.

Each central banker engages in forward guidance at the ZLB under limited commitment. We define this as follows:

Definition 1.3.3. *Forward guidance with limited commitment is an exogenous integer* $\tau \in \mathbb{N}$ *that prescribes a non-binding promised duration of a zero nominal interest rate path after exiting the ZLB.*

Hence, we consider whether some exogenously given τ can be sustained as an equilibrium. The definition of the impact response to forward guidance under full commitment extends naturally to the case under limited commitment.

As there is limited commitment, the central banker faces the strategic choice of whether to stick to the promised zero interest rate path or deviate to discretion in any of the τ periods upon exiting the ZLB. Following Maskin and Tirole (2001), we restrict attention to stationary Markovian strategies such that the central banker bases her decision on the following state:

Definition 1.3.4. *A state is a triplet* (r_t^n, μ_t, P_t) *such that:*

— $r_t^n \in \mathbb{R}$ *is the natural interest rate at time* t

— $\mu_t \in \Delta(\{0, \beta\})$ *is the private agents' belief distribution over types at time* t

— $P_t \in \mathcal{P} = \{0, 1, \ldots, \tau\}$ *is the number of periods for which forward guidance is honored at calendar time* t.

The first dimension is the natural interest rate r_t^n. The second dimension is the private sector assessment about the distribution of types, μ_t, namely reputation. The third dimension P_t is the number of times that a promise has been kept at calendar time t. For instance, $P_t = 0$ means that so far a promise has not been kept at time t, while $P_t = 3$ means that at time t, the central bank has kept the promise in period $t-2$, $t-1$, and t. The state contains all the payoff relevant variables as inflation and output gap are pinned down by the natural interest rate, reputation and monetary policy.

For each state of the economy, the central banker needs to decide whether or not to honor a past promise. Hence, we define strategies as follows:

Definition 1.3.5. *Fix a promise* τ. *For each calendar period* t *and type* δ, *a strategy* σ *prescribes the probability of transitioning from the promise period* $P_{t-1} \in \mathcal{P}$ *to the promise period* $P_t = P_{t-1} + 1 \in \mathcal{P}$ *with natural real interest rate* r_t^n *and reputation* μ_{t-1}*:*

$$\begin{aligned} \sigma : \{0, \beta\} \times \mathbb{R} \times \Delta\big(\{0, \beta\}\big) \times \mathcal{P} &\longrightarrow [0, 1] \\ (\delta, r_t^n, \mu_{t-1}, P_{t-1}) \qquad &\longmapsto \qquad \mathbb{P}(P_t | P_t - 1, r_t^n, \mu_{t-1}, \delta) \end{aligned}$$

We only consider pure strategies, so the possible transition probabilities are either zero or one. Instead, under full commitment $\sigma \equiv 1$.

All actions are publicly observed, so the private sector will learn about the likelihood of each possible type, and the central bank perfectly tracks the evolution of agent's beliefs. The prior distribution $\mu_0 \in \Delta\big(\{0, \beta\}\big)$ is common knowledge. The private sector is composed of Bayesian agents, such that beliefs after promise keeping evolve according to:

$$\mu_t(\delta; r_t^n, \mu_{t-1}, P_t) = \sum_{\hat{\delta}} \frac{\sigma(\hat{\delta}, r_t^n, \mu_{t-1}, P_t) \cdot \mu_{t-1}(\hat{\delta})}{\sum_{\tilde{\delta}} \sigma(\tilde{\delta}, r_t^n, \mu_{t-1}, P_t) \cdot \mu_{t-1}(\tilde{\delta})} \cdot \Gamma(\hat{\delta}, \delta)$$

where

$$\Gamma(\hat{\delta}, \delta) = \begin{cases} \Gamma(\hat{\delta}, \delta) = 1 - \gamma & \text{if } \hat{\delta} = \delta \\ \Gamma(\hat{\delta}, \delta) = \gamma & \text{if } \hat{\delta} \neq \delta \end{cases}$$

for $\delta, \hat{\delta} \in \{0, \beta\}$, where $\gamma < \frac{1}{2}$.

Whenever there is no promise in place, reputation mechanically evolves following the transition probabilities, as there is no strategic choice being made by the central banker, and hence households do not update from policy.

1.3.4 Equilibrium

An equilibrium in our limited commitment-reputation model is to constitute both a macroeconomic equilibrium, in the sense that the macroeconomic variables are given by the New Keynesian model, and a strategic equilibrium such that the strategies of the central bankers together with the private sector's beliefs constitute a Perfect Bayesian Equilibrium.

Definition 1.3.6. *An equilibrium is a tuple (σ, μ, y, π, i) such that:*

- *Given (y, π, i), the pair (σ, μ) is a Perfect Bayesian Equilibrium:*
  - *Beliefs μ are consistent with the strategies σ of the central bankers on the equilibrium path.*
  - *Strategies σ of the central bankers are sequentially rational given the beliefs μ.*
- *Given strategies and beliefs (σ, μ), output gap y, inflation π, and nominal interest rates i satisfy (IS) and (PC).*

1.3.5 Computation of the Equilibrium

To solve for the macroeconomic equilibrium, we first notice that for each pair (σ, μ) there is a subset $Z_{\sigma,\mu} \subset \mathbb{R}$ such that the ZLB binds if and only if $r_t^n \in Z_{\sigma,\mu}$. At the

ZLB, output gap and inflation are given by the New Keynesian Phillips curve and the IS curve, imposing that the nominal interest rate is zero. We report the full set of equations characterizing the macroeconomic equilibrium in Appendix 1.B. Below we report the expression for the IS curve at the ZLB.

$$y_t^Z = \mathbb{E}[y_{t+1}] + \frac{1}{\eta}(\mathbb{E}[\pi_{t+1}] + r_t^n)$$

where expected output gap is given by [4]

$$\mathbb{E}[y_{t+1}] = \mathbb{E}\left[\underbrace{y_{t+1}^Z \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\}}_{(a)} + \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\}\left(\underbrace{y_{t+1}^P \mathbb{E}_\delta^{\mu_t}[\sigma(\delta)]}_{(b)} + \underbrace{y_{t+1}^D \mathbb{E}_\delta^{\mu_t}[1-\sigma(\delta)]}_{(c)}\right)\right]$$

The operator $\mathbb{E}$ denotes the expectation over r_{t+1}^n given r_t^n, $\mathbb{E}_\delta^{\mu_t}$ denotes the expectation with respect to δ using distribution μ_t. Z denotes the *zero lower bound*, P denotes *promise keeping* and D denotes *discretion*. Hence, y_t^P is output under promise keeping, and y_t^D is output in discretion. We emphasize that output gap and inflation depend on the state of the economy but we have suppressed it for the sake of exposition.

To better describe how to calculate the expectation, we decompose expected output gap into three terms. The term (a) corresponds to output gap whenever the ZLB is binding again. Terms (b) and (c) refer to a realization of the natural rate of interest such that the ZLB is not binding. Term (b) is output gap under promise keeping, i.e. the output gap that prevails if the central bank chooses $i_{t+1} = 0$, multiplied by the probability that the central bank keeps its promise, $\mathbb{E}_\delta^{\mu_t}[\sigma(\delta)]$. For notational convenience, we have suppressed the dependence of this probability on the state of the economy. Finally, term (c) is output gap in normal times, which is multiplied by the probability that the central bank breaks the promise, $\mathbb{E}_\delta^{\mu_t}[1-\sigma(\delta)]$.

Having explained how to solve for the macroeconomic equilibrium at the ZLB, we

4. Expected inflation can be obtained similarly.

now explain how to solve for the macroeconomic equilibrium in normal times, i.e. for any shock $r^n_t \notin Z_{\sigma,\mu}$ and $P_t = 0$. As the central bank follows optimal discretion, monetary policy is set to minimize the instantaneous loss function such that the first order condition reads:

$$-\kappa\pi_t = \lambda y_t$$

Output gap and inflation are then given by the New Keynesian Phillips curve and the IS curve. The full system of equations is reported in Appendix 1.B.

Finally, the economy can also be in promise keeping, i.e. it is away from the ZLB ($r^n_t \notin Z_{\sigma,\mu}$) and there is a promise in place ($P_t > 0$). The central banker then decides whether to keep the promise or not for another period as captured by the strategy σ. Finally, the economy can be in the last period of promise keeping out of the ZLB, in which $P_t = \tau$ and no strategic choice is made. This follows as the economy simply transitions to either discretion or the ZLB depending on the realization of the natural interest rate. If the economy were to hit the ZLB in the sense that the desired interest rate is negative, we assume that the central bank renews the promise of length τ. This helps us gain a recursive structure.

Assumption 1.3.1. *If $P_t \geq 0$, and $r^n_{t+1} \in Z_{\sigma,\mu}$, then $P_{t+1} = 0$.*

The problem of the central bank given the state of the economy at the ZLB is given by

$$L^Z_t(\delta) = \lambda(y^Z_t)^2 + (\pi^Z_t)^2 + \delta\mathbb{E}\left[\underbrace{L^Z_{t+1}(\delta)\mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\}}_{(a)}\right.$$
$$\left. + \min_{x\in\{0,1\}}\left\{\underbrace{xL^P_{t+1}(\delta)}_{(b)} + \underbrace{(1-x)L^D_{t+1}(\delta)}_{(c)}\right\}\mathbb{1}\{r_{t+1} \notin Z_{\sigma,\mu}\}\right]$$

The loss at the ZLB depends on the current realization of output and inflation and the discounted expected future loss. The future expected loss can be decomposed into 3 terms. Term (a) reflects the loss of being at the ZLB, while term (b) reflects the loss of promise keeping. Finally, term (c) is the loss from discretion. The problems in normal times and in promise keeping are carefully described in Appendix 1.B.

The above intertemporal loss function plays an important role once the economy exits the ZLB, and the central banker decides whether to honor a promise or to break it and conduct discretionary monetary policy. The trade off she faces is the following: on the one hand, by keeping the promise, inflation and the output gap are too expansionary compared to what could have been achieved by setting the interest rate according to optimal discretion. All else equal, this provides the central banker with an incentive to renege on her promise. On the other hand, by keeping the promise the central banker signals its ability to keep promises in future ZLB events. Consequently, the private sector's assessment about her type evolves over time as a function of the behavior of the central banker and the state of the economy.

1.4 Solution to the Forward Guidance Puzzle

A New Keynesian model augmented with limited commitment and reputation does not suffer from the forward guidance puzzle. While some level of forward guidance at the ZLB remains useful, we find that an unreasonably expansionary promise cannot be sustained as an equilibrium strategy.

When the central banker decides whether to keep the promise or conduct optimal discretionary policy, she faces the following trade off: on the one hand, keeping a promise causes an excess expansion today. On the other hand, keeping the promise helps the central bank gain reputation which is useful to stabilize output gap and inflation in future ZLB events. A first result is the following:

Proposition 1.4.1. *There exist* $\bar{\lambda} > 0$ *such that if the weight on output gap* $\lambda \in [0, \bar{\lambda}]$ *the myopic central banker never keeps any promise* $\tau > 0$*. I.e, in any equilibrium we have* $\sigma(\delta = 0, r_t^n, \mu_t, P_t) = 0$ *for any* r_t^n*,* μ_t *and* P_t*.*

For a large $\lambda > 0$, even the myopic central banker may keep her promise as reputational gains have a direct impact on current economic outcomes. In fact, in this case it is possible to find values of the parameters such that the myopic central banker will be willing to keep the promise.[5] Going forward we assume that λ is small enough such that the myopic central banker never keeps her promise.

Another result is that there are states of the world where the reputation gains are not important enough to outweigh the cost of an excess expansion:

Proposition 1.4.2. *There exists $\hat{R} < \infty$ such that in any equilibrium, for all $r_t^n > \hat{R}$, μ_t, P_t and δ we have $\sigma(\delta, r_t^n, \mu_t, P_t) = 0$.*

The intuition is that for a very positive demand shock, the expansion that the economy would undergo with a zero interest rate would be too costly. It would therefore be preferable to deviate to discretion to mitigate the boom in output gap and inflation. In fact, while the loss of the central banker is always bounded under optimal discretion for any realization of the demand shock, the loss goes to infinity as the realization increases if the interest rate is kept at zero.

We further find that the existence of even a minimal cost of keeping the promise is enough to ensure the existence of a pooling equilibrium where neither type keeps the promise:

Proposition 1.4.3. *For any promise τ, there exist equilibrium strategies σ such that for any δ, r_t^n, μ_t and P_t, we have that $\sigma(\delta, r_t^n, \mu_t, P_t) = 0$.*

This is due to the fact that pessimistic off-path beliefs render deviations to promise keeping futile as the reputation channel is muted.
Combined with Proposition 1.4.1, we can therefore restrict attention to semi-separating equilibria where the forward-looking central banker keeps her promise in some states of the world only, and the myopic central banker never keeps her promise. There can be

5. However, we find that with our preferred parameterization, the myopic central banker never keeps her promise.

no other equilibrium in which forward guidance is effective.

While an excess expansion provides the central banker with a reason to renege on past promises, the reputation channel can be strong enough to provide her with an adequate incentive to honor them. To better understand how this channel works, suppose that the forward-looking central banker keeps her promise in some states of the world. In that case, the effect of forward guidance increases in the probability that agents assign to the central banker being of the forward-looking type. This, in turn, provides the central banker with a reason to indeed keep the promise. For a strong enough reputation channel, there exists a semi-separating equilibrium where the forward-looking central banker keeps her promise in some states of the world. We verify that, with our preferred parameterization of the model, such an equilibrium indeed exists, i.e. promises are kept in states of the world where the realization of the demand is sufficiently moderate. We are now ready to state the main result of the paper: for a too generous promise, the effect of forward guidance is zero.

Proposition 1.4.4. *There exists $\bar{\tau} \in \mathbb{N}$ such that for any variable X, for all $\tau > \bar{\tau}$, $IR_X(\tau) = 0$.*

We thus find that there is a maximum level of forward guidance that is sustainable in the sense that promises are kept in some states of the world. Any promise with a longer duration will not be believed and hence it will fail to stabilize economic outcomes at the ZLB.

Intuitively, the cost of keeping a promise is increasing in its length. A generous promise, if believed, implies a significant expansion once the economy exists the ZLB. Hence, agents understand that a promise cannot be kept if it is too generous, and therefore keeping it does not constitute a strategic equilibrium. The only equilibrium is therefore pooling where promises are never kept, rendering forward guidance ineffective. From Proposition 1.4.4, we immediately have the following corollary:

Corollary 1.4.1. *The marginal impact of forward guidance is negative at* $\tau = \bar{\tau}$*:*

$$IR_y(\bar{\tau}+1) - IR_y(\bar{\tau}) < 0$$

Once a promise goes from being sustainable to unsustainable, the effect of forward guidance drops to zero, and hence the marginal effect of increasing the duration is negative. This corollary is an example of how our model is different from the literature giving rise to a "discounted Euler equation" where the marginal effect of forward guidance is always weakly positive. We further discuss differences with that literature in section 1.6.

We provide a graphical representation of Proposition 1.4.4 and Corollary 1.4.1 in Figure 1.1. We find that there exist semi-separating equilibria where the forward-looking central banker keeps her promise for moderate demand shocks for a duration of the promise being no higher than five periods corresponding to one year and a quarter.

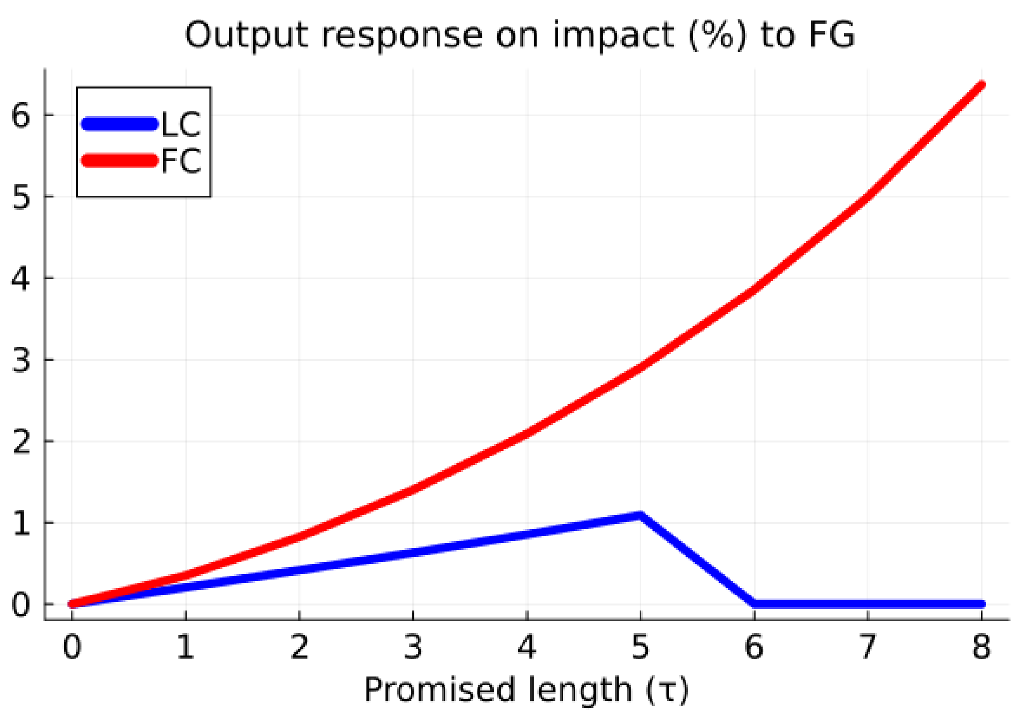

Figure 1.1 – Mitigation of forward guidance effect for $\mu_0(\beta) = 0.6$.

Despite forward guidance being a useful policy tool in the semi-separating equilibrium, the effect is much mitigated compared to the full commitment case. This happens for two reasons. First, households realize that the myopic central banker will never keep her promise. Second, they realize that there are states of the world where it is too costly for any central banker to keep her promise regardless of her type.

1.5 The Effect of Forward Guidance

For the quantitative exploration of the model, we will assume that the loss function takes the following form:

$$L\big(\{y_t\}_{t=0}^{\infty}, \{\pi_t\}_{t=0}^{\infty}\big) = \sum_{t=0}^{\infty} \delta^t (\lambda(y_t - \bar{y})^2 + \pi_t^2), \quad \delta \in \{0, \beta\}$$

where $\bar{y} > 0$ is the targeted output gap. The parameter $\bar{y}$ captures that there may be inefficiencies in the economy such that the central seeks to stabilize output at a level which is higher than the natural one and thereby target a positive output gap. We include a positive output gap target to counteract the deflationary forces caused by the ZLB.

We assume that the natural rate of interest follows an AR(1) process:

$$r_t^n = (1-\rho)\bar{r} + \rho r_t^n + \varsigma \xi_t, \quad \xi_t \sim \mathcal{N}(0,1)$$

We discretize the AR(1) process by the Tauchen (1986) method where we choose three values for the shock to the natural interest rate. We choose three to have a realization such that the ZLB binds, a realization where the forward-looking central banker keeps the promise, as well as a realization where it is excessively costly for any central banker to keep the promise. This is sufficient to illustrate the main points of the paper.

We choose to parameterize the model such that one period corresponds to one quarter. We calibrate the persistence of the demand shock (ρ), the variance of the demand shock (ς^2), and the targeted output gap ($\bar{y}$), while we choose values that we consider standard in the New Keynesian literature for the remaining parameters. We display all of the values in Table 1.1.

Parameter	Value	Description	Target
β	0.99	Discount factor	4% annual interest rate
κ	0.02	Slope of Phillips Curve	Walsh (2018)
$\bar{y}$	1%	Output gap target	Average output gap of zero
η	2.0	Relative risk aversion	McKay, Nakamura, and Steinsson (2016b)
λ	0.003	Weight on output gap	Walsh (2018)
$\mu_0(\beta)$	0.6	Prior probability of high type	King and Lu (2022)
γ	0.03	Transition probability of types	8 years mandate
ς	0.014	Standard deviation of ξ_t	Std. deviation of output gap
ρ	0.62	Auto correlation of r_t^n	Auto correlation of inflation
$\bar{r}$	1%	Natural real interest rate	$-\log(\beta)$

Table 1.1 – Parameter values

We calibrate the standard deviation of the innovation (ς) such that the model with no forward guidance matches the standard deviation in US output gap in the period 1987-2018. Similarly, we calibrate the persistence (ρ) of the demand shock to the auto correlation of US inflation. We retrieve data from FRED on real GDP in the US over the period 1987-2018. We obtain a standard deviation of output gap of 1.87, and an auto correlation of 0.54, at a quarterly level in the period 1987-2018. The implied values for the standard deviation for the innovation (ς) and for the persistence of the demand shock (ρ) are 1.4 and 0.62, respectively.

We choose to calibrate the output target to match a zero average output gap. If we did not include a positive target for the output gap, the model would predict a too low average output gap due to the deflationary bias introduced by the ZLB as explained by G. Eggertsson (2006).

For the remaining parameters, we choose values that we consider standard in the literature. First, as we consider a semi annual model, we choose β, the discount factor of private agents, to be 0.99, and hence the value of $\bar{r}$, i.e. the natural rate of interest is equal to $-\log(\beta)$, which corresponds to a steady state interest rate of around 4 % a year. We choose a value of the slope of the Phillips curve (κ) to be 0.02. For the coefficient of relative risk aversion (η), we choose a value of 2.

1.5.1 Reputation

The existence of a semi-separating equilibrium implies that the effect of forward guidance depends on reputation. The higher the probability agents assign to the central banker being the forward-looking type, the larger the effect of forward guidance on output as shown in Figure 1.2. We also report the effect of forward guidance under full commitment in red for comparison.

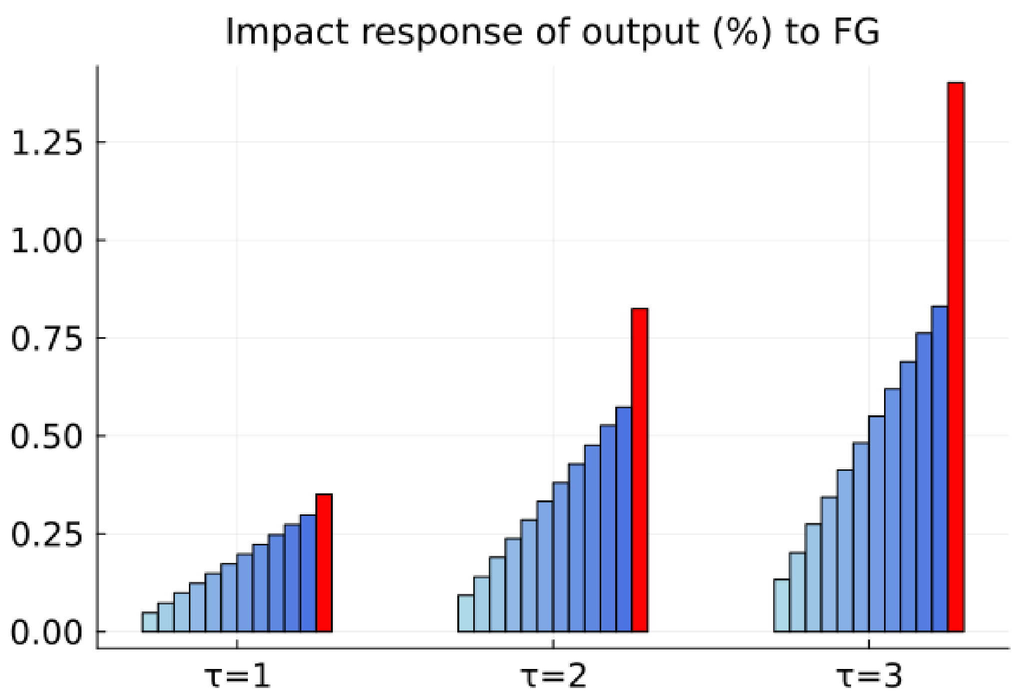

Figure 1.2 – Reputation and forward guidance. In red, the effect of forward guidance under full commitment; in blue, the effect of forward guidance under limited commitment. The darker the blue, the higher the belief $\mu_0(\beta)$, ranging from 0 to 1 with increments of 0.1.

From Figure 1.2 we see that for $\mu_0(\beta) = 0$, forward guidance is almost futile as a policy tool as the households expect the current central banker to not keep the promise, and assign very low probability ($\gamma = 0.03$) to a change in the type of the central banker. It is then monotonically increasing in the belief that the central banker is the forward-looking type. The effect of forward guidance with limited commitment is mitigated compared to full commitment, even if households believe that the central banker is the forward-looking type (i.e. $\mu_0(\beta) = 1$). This occurs as households realize that (1) there are realizations of the natural interest rate such that it is too costly to keep the promise for the forward-looking type, and (2) types are mean-reverting. This mitigation effect is not very strong for $\tau = 1$. The intuition is that given that the economy is at the ZLB, it is considered very likely that the economy will not be in an excessive boom in the following period. That is, the realization of the demand shock is likely to be only moderate once the economy exits the ZLB. We see a stronger mitigation effect for larger values of τ when we compare limited commitment to full commitment. The reason is that for higher values of τ, it is more likely that the economy will be hit by a very positive demand shock in which case even the forward-looking central banker will have to deviate from her promise rendering forward guidance less credible.

To illustrate the dynamics of the model, we consider a situation where households initially hold a belief of 0.5 that the central banker is forward-looking, and the economy is in normal times for the first 4 periods while in period 5 the ZLB hits and the central banker engages in forward guidance with $\tau = 1$. For the purpose of illustration, we assume that the ZLB only binds in period 5. In Figure 1.3, we show how beliefs evolve depending on whether the central banker keeps her promise or not.

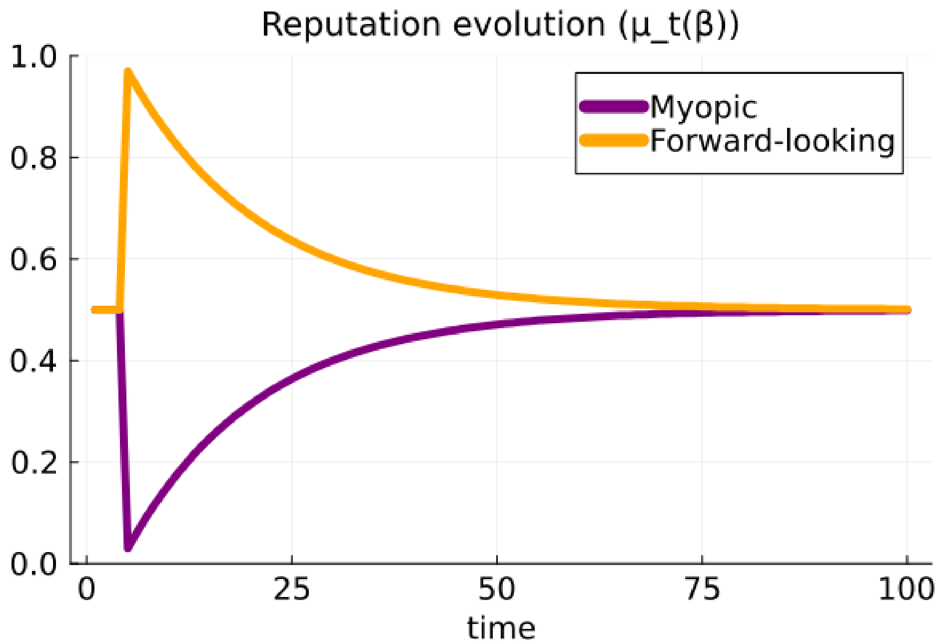

Figure 1.3 – Reputation convergence with $\mu_0(\beta) = 0.5$.

We see that immediately after the promise keeping decision of the central banker, there is a large change in beliefs. Following this change, beliefs then converge back to the stationary distribution in the absence of further ZLB periods.

We simulate our model for 100 periods in a path where the type of the central bank happens to not change with the aim of showing how reputation evolves depending on the type of the central banker, and how the reputation of the central banker influences economic outcomes. In particular, we plot the value of inflation and how it evolves together with reputation in a model with $\tau = 3$ and $\mu_0(\beta) = 0.5$.

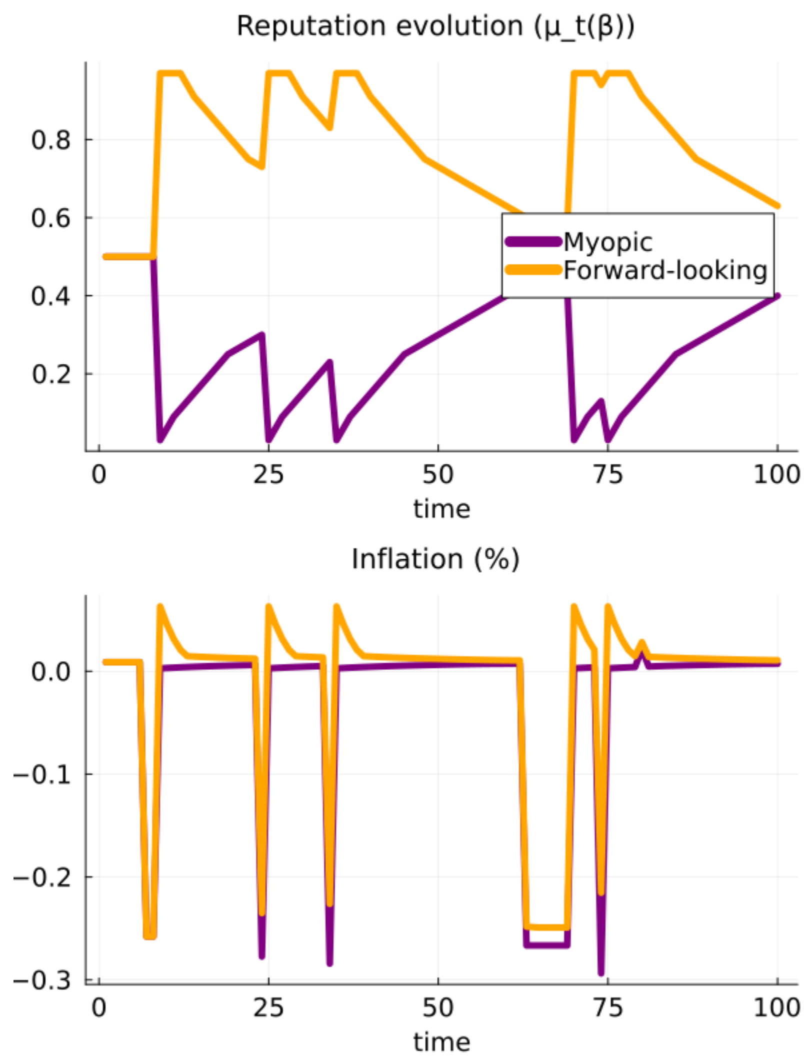

Figure 1.4 – A simulation of reputation and inflation with $\tau = 3$ and $\mu_0(\beta) = 0.5$.

We see that having a high reputation is associated to a smaller drop in inflation at the ZLB. Interestingly, if we consider inflation at the ZLB that hits the economy in periods 60 and 75, respectively, we notice that the drop is mitigated in period 75 for the forward-looking type (yellow line) while it is exacerbated for the myopic type (purple line). This happens as the occurrence of a ZLB in the recent past has given agents a good idea of the type of the central banker.

1.5.2 Forward guidance and the persistence of the ZLB

We find that the persistence of the demand shock plays an important role in determining the sustainability of forward guidance as well as its effect on the output gap at the ZLB. To show that, we vary the autoregressive coefficient (ρ) while at the same time adjusting the variance of the innovation ς^2 such that the variance of the natural interest rate is kept fixed.[6]

In Figure 1.5 we show the effect on the output gap at the ZLB of promising a zero interest rate for a time horizon of τ periods as a function of the persistence of the demand shock. We also show the effect of forward guidance with full commitment in red for comparison.

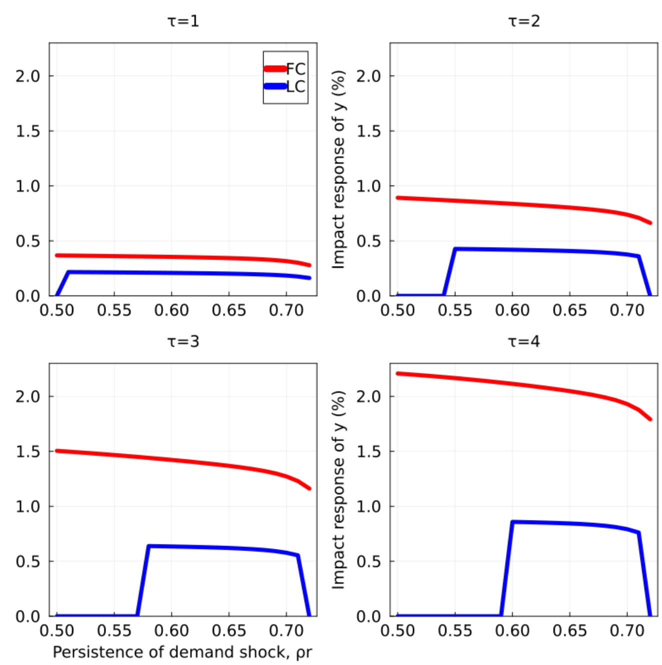

Figure 1.5 – Forward guidance and the persistence of the natural rate of interest

6. We perform a similar exercise for the standard deviation of the demand shock, and we report the results in appendix 1.D.

We see that there is a non-monotonic relation between the effect of forward guidance and the persistence of the demand shock for some values of τ when the central bank cannot commit, while the effect of forward guidance is smoothly decreasing in the persistence of the demand shock in a model with full commitment.

Intuitively, a higher persistence of the demand shock implies that the economy is less likely to exit the ZLB in the following period which in turn means that the economy is less likely to transition to the promise keeping state. This explains why the effect of forward guidance is decreasing in the persistence of the demand shock with full commitment.

With limited commitment, there is no longer a monotonic relationship between the persistence of the demand shock and the effect of forward guidance. For a very low persistence of the shock, we find that a promise of keeping the interest rate at 0 for one quarter after the economy exits the ZLB has a zero impact on the output gap at the ZLB. This happens as a low persistence of the demand shock renders the ZLB less severe due to the fact that the economy is likely to exit in the following period, causing expectations about future output gap and inflation to be rather expansionary. Due to the fact that the ZLB is not very severe, the value of having a high reputation, allowing the central banker to stabilize economic outcomes at the ZLB, is modest. Hence, the central banker prefers stabilizing economic outcomes once out of the ZLB by deviating to discretion instead of building reputation.

We see that for a persistence of around 0.51, the effect of forward guidance on output gap jumps from 0 % to 0.25 % when $\tau = 1$. The reason is that as the persistence increases, the ZLB gets more severe implying that the central banker faces stronger incentives to keep her promise to gain reputation. After this jump, the effect of forward guidance then smoothly declines in the persistence of the demand shock for the same reasons as under full commitment: the higher the persistence, the less likely that the economy will transition to the promise keeping state in the following period. We see that eventually the effect of forward guidance drops to zero. This is due to the fact that for a too high persistence, the promise keeping state gets too expansionary implying that the central banker will have to deviate from her promise as it is too costly to keep it.

We find that the higher the value of τ, the smaller the set of values of the persistence such that forward guidance can be sustained as an equilibrium strategy. The intuition is that longer promises are associated to larger expansions implying that the central banker needs stronger incentives to keep the promise, i.e. a very bad ZLB while at the same the persistence cannot be too high as the promise keeping state would then be too expansionary.

1.5.3 Welfare loss

The presence of a ZLB causes a welfare loss as the central bank is unable to stabilize inflation and output gap through adjustments to the nominal interest rate. Instead, the central bank is forced to rely on unconventional policy measures such as forward guidance to stabilize the economy G. Eggertsson and Woodford (2003).

We compute the loss associated to the ZLB by applying the discounted loss function of the forward-looking central banker which approximates household welfare. Following Galí (2015), we can interpret this loss as the fraction of steady state consumption up to an additive constant that the household is willing to give up to avoid fluctuations in the economy. We find that the loss is 0.003 % of steady state consumption if the central bank cannot rely on forward guidance. It is not surprising that we do not find important quantitative welfare losses of the ZLB as we follow the approach by Lucas (1987), and it is well known that the estimated welfare losses of business cycle fluctuations are very small.

We then report the relative loss at the ZLB for central bank that can conduct forward guidance. We compare a central bank that has only limited commitment to a central bank of full commitment. The relative losses are reported in Table 1.2.

Scenario	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$
Loss with FC	**95.6 %**	95.8 %	105.4 %	132.8 %	191.9 %
Loss with LC	94.4 %	88.8 %	83.8 %	80.1 %	**78.7 %**

Table 1.2 – Welfare losses relative to ZLB with no FG. Initial belief distribution is $\mu_0(\beta) = 0.6$.

In general, we find that the relative losses are smaller once forward guidance is available to the central bank. We find that a central bank which is forced to commit to a zero interest rate should optimally choose a length of $\tau = 1$. On the contrary, a central bank that does not have access to a commitment technology should optimally choose a length of $\tau = 5$ corresponding to five quarters.

We therefore find that not having access to a commitment technology causes the central bank to give a longer promise. The intuition is two-fold. First, as agents realize that the central bank may deviate from its promise, the expansionary effect of a promise at the ZLB is reduced under limited commitment. Second, as the central bank has not committed, it can choose to deviate from the given promise in a state where it is very costly to keep it.

1.6 Relation to Discounted Euler Equation

McKay, Nakamura, and Steinsson (2016a) emphasizes that the highly forward looking nature of the Euler equation is at the core of the forward guidance puzzle. A large literature solving the this puzzle can be summarized by the so-called "discounted Euler equation" which takes the following form:

$$y_t = \Omega\, \mathbb{E}_t[y_{t+1}] - \frac{\chi}{\eta}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)$$

and can be iterated forward to obtain:

$$y_t = -\frac{\chi}{\eta} \sum_{j=0}^{\infty} \Omega^j \, \mathbb{E}_t \left[i_{t+j} - \pi_{t+j+1} - r^n_{t+j} \right]$$

In the standard New Keynesian model $\Omega = \chi = 1$. The condition $\Omega < 1$ dampens the effect of future expected interest rates on current economic outcomes which explains why forward guidance about interest rates in the distant future is mitigated with a discounted Euler equation compared to the standard model.

However, we find it paradoxical that agents make consumption-savings decisions looking infinitely into the future while failing to understand the incentives of the central bank to keep a time-inconsistent promise. The ability of agents to understand the incentives of the central bank is what we add to the standard model.

The difference between our limited commitment-reputation model and this literature is that we do not obtain attenuation of forward guidance through discounting of the Euler equation, i.e. we have $\Omega = 1$. Instead, we obtain attenuation of forward guidance by assuming that agents understand the incentives of the central bank to keep or renege on a promise. Hence, our attenuation works through the lack of effect that a promise has on expectations regarding the future interest rate path.

In Table 1.3, we compare the output gap and inflation that pertain at the ZLB in the standard model, the discounted Euler equation model, and our model with limited commitment and reputation.

Model		Output gap	Inflation
Standard NK	$(\Omega = 1.00, \chi = 1.00)$	-6.49%	-0.33%
Discounted Euler	$(\Omega = 0.97, \chi = 0.75)$	-4.48%	-0.18%
LC and reputation	$(\Omega = 1.00, \chi = 1.00)$	-6.49%	-0.33%

Table 1.3 – Recession at the ZLB with no forward guidance

A discounted Euler equation attenuates the fall in output and inflation at the ZLB compared to the standard model as stressed by McKay, Nakamura, and Steinsson (2016a). Instead, our model with limited commitment and reputation does not impact economic outcomes at the ZLB in the absence of forward guidance compared to the standard NK model. This is due to the fact that the discounted Euler literature obtains mitigation of a ZLB by reducing the front-loading effect of future drops in inflation and output gap.

We show how the economy responds to a promise of keeping the interest rate at 0 for τ periods in Figure 1.6 in the three different models.

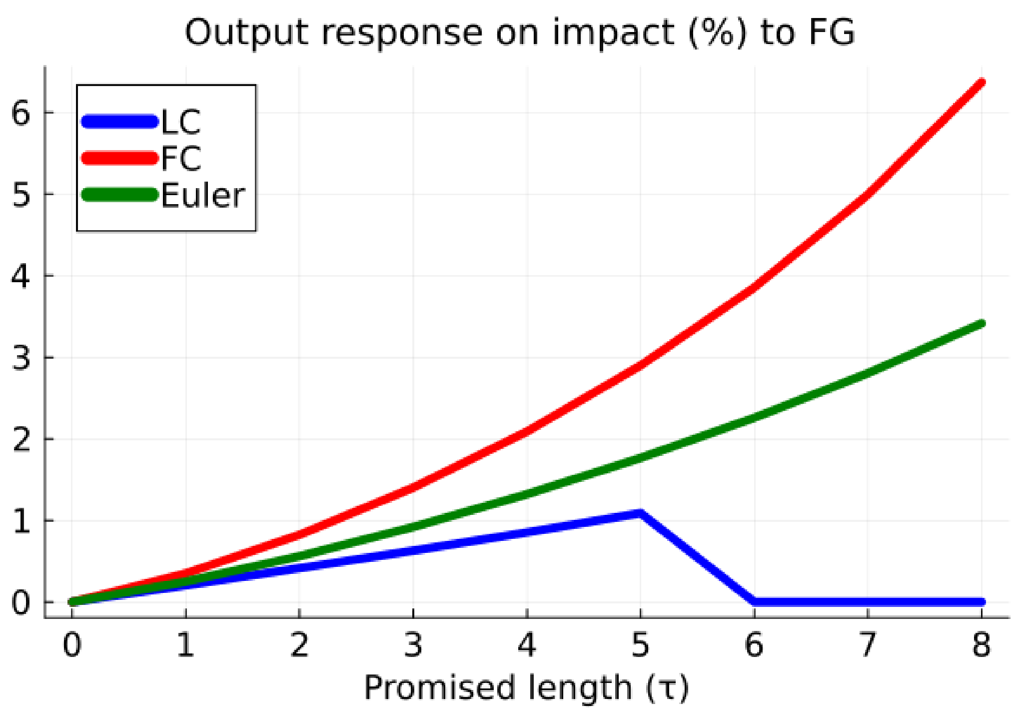

Figure 1.6 – Discounted Euler equation model

Both a model with a discounted Euler equation and our limited commitment-reputation model provide attenuation of the effect of forward guidance, but an important difference between the two models is that with the former, the marginal effect of giving a more generous promise is always weakly positive while in the latter we obtain that the marginal effect may be negative if the promise induces agents to lose faith in the ability of the central bank to keep its promise. Hence, central bank communication may be more expansionary by giving a humble promise that can be kept compared to a more expansionary one, which agents expect central banks to renege on.

While an empirical investigation of forward guidance to discriminate between lack of commitment or behavioral agents is beyond the scope of this paper, we notice that our paper is consistent with the notion that central banks have only limited influence on expected interest rates far out in the future. This is consistent with the results in Campbell et al. (2019) that the FED has limited ability to affect expectations at long horizons.

1.7 Extension: Forward Guidance and Quantitative Easing

Bhattarai, G. B. Eggertsson, and Gafarov (2023) show that quantitative easing (an increase in short run government debt) stimulates the economy by inducing the central bank to choose a lower interest rate in the future. They informally discuss how quantitative easing therefore can serve as a commitment device to forward guidance and hence work to strengthen the effect of it. We now extend our quantitative model to study the interaction between forward guidance and government debt in our limited commitment-reputation model.

We assume that there is a level of government spending at the ZLB that is pure waste in the sense that it does not affect aggregate demand nor consumer welfare. Whatever is spent at the ZLB is financed by short term debt that has to be paid back in the following period.

Bhattarai, G. B. Eggertsson, and Gafarov (2023) microfounds a loss function that is increasing in the level of taxation as tax payments are associated to a decrease in welfare due to their distortionary nature. We then follow their approach and assume a loss function of the following form:

$$L\big(\{y_t\}_{t=0}^{\infty},\{\pi_t\}_{t=0}^{\infty},\{T_t\}_{t=0}^{\infty}\big)=\sum_{t=0}^{\infty}\delta^t\Big[\ \lambda\cdot(y_t-\bar{y})^2+\pi_t^2+\lambda_T\cdot T_t^2\Big]$$

The term T denotes taxes that are raised to repay government debt raised in the previous period. In particular, $T_t=(1+i_{t-1})\cdot G_{t-1}$. The model is otherwise identical to the one described in Section 1.3. We choose the same calibration as in the baseline model, and we put the weight on taxes $\lambda_T=0.003$.

In Table 1.4 we report the interest rate chosen by the central bank at the median realization of the demand shock in the model with and without government debt but with no forward guidance. We find that having higher government debt reduces the interest rate chosen by the central bank when they conduct discretionary policy. This result replicates the finding of Bhattarai, G. B. Eggertsson, and Gafarov (2023).

Debt level		i_t
$G_t=0$	(baseline model)	0.66 %
$G_t=0.03$	(model with debt)	0.35 %
$G_t=0$	(model with debt)	0.67 %

Table 1.4 – Interest rates under discretionary policy with no forward guidance.

We find that the interest rate in discretion is 0.66 % in the baseline model where government debt is always zero. We then compute the interest rate in a model where the government issues debt at the ZLB. The interest rate is lower in discretion (0.35%) when the government has a positive level of debt. The intuition is that the central bank tries to minimize the tax payments of the private agents while still taking into account how the interest rate influences output and inflation. As a consequence of this, the zero lower bound is less severe in such a model which causes the interest rate to be higher (0.67%) in discretionary periods with no debt as expectations about future outcomes are more expansionary.

We then study the consequences of forward guidance in the model with government debt. We display in Figure 1.7 the increase in output at the ZLB of promising an interest rate of zero in the period where the economy exists the ZLB.

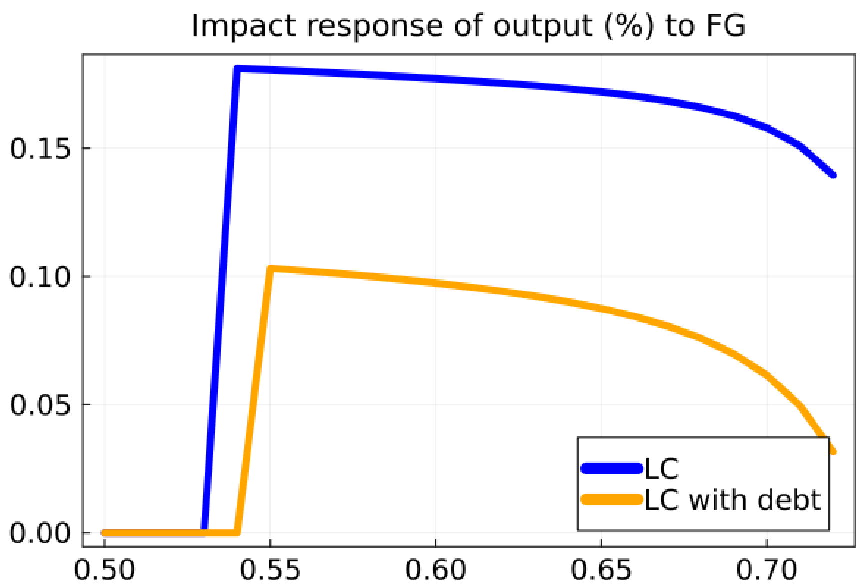

Figure 1.7 – Forward guidance with government debt

From the graph, it follows that the effect of forward guidance is muted once there is government debt. The reason for this is that the presence of government debt reduces the optimal interest rate of the central bank in discretion. This renders a promise of low interest rates in the future superfluous as private agents already expect the central bank to keep rates low. We therefore obtain that short term government debt and forward guidance are substitute policies whereas Bhattarai, G. B. Eggertsson, and Gafarov (2023) argue that they are complements.

1.8 Conclusion

We showed how a standard New Keynesian model augmented with limited commitment and type-based reputation bounds the effect of forward guidance. An important implication of the model is that the marginal effect of the promised horizon becomes negative when the promise is judged unsustainable by private agents. We demonstrated a strong interaction between the persistence of the ZLB, and the sustainability and effect of forward guidance. We conjecture that our model could be useful for studying other time inconsistent policies.

References

Acharya, Sushant and Keshav Dogra (2020). « Understanding HANK: Insights From a PRANK ». In: *Econometrica* 88.3, pp. 1113–1158.

Angeletos, George-Marios and Chen Lian (2018). « Forward Guidance without Common Knowledge ». In: *American Economic Review* 108.9, pp. 2477–2512.

Barro, Robert J. and David B. Gordon (1983). « Rules, discretion and reputation in a model of monetary policy ». In: *Journal of Monetary Economics* 12.1, pp. 101–121. ISSN: 0304-3932.

Barthélemy, Jean and Eric Mengus (2018). « The signaling effect of raising inflation ». In: *Journal of Economic Theory* 178, pp. 488–516.

Bernanke, Ben S. (2020). « The New Tools of Monetary Policy ». In: *American Economic Review* 110.4, pp. 943–83.

Bhattarai, Saroj, Gauti B. Eggertsson, and Bulat Gafarov (2023). « Time Consistency and Duration of Government Debt: A Model of Quantitative Easing ». In: *Quarterly Journal of Economics* Forthcoming.

Bilbiie, Florin O. (2019). « Optimal Forward Guidance ». In: *American Economic Journal: Macroeconomics* 11.4, pp. 310–45.

Campbell, Jeffrey R. et al. (2019). « The limits of forward guidance ». In: *Journal of Monetary Economics* 108, pp. 118–134. ISSN: 0304-3932.

Chari, V. V. and Patrick J. Kehoe (1990). « Sustainable Plans ». In: *Journal of Political Economy* 98.4, pp. 783–802.

Debortoli, Davide and Aeimit Lakdawala (2016). « How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-optimizations ». In: *American Economic Journal: Macroeconomics* 8.3, pp. 42–76.

Debortoli, Davide, Junior Maih, and Ricardo Nunes (2014). « LOOSE COMMITMENT IN MEDIUM-SCALE MACROECONOMIC MODELS: THEORY AND APPLICATIONS ». In: *Macroeconomic Dynamics* 18.1, pp. 175–198.

Debortoli, Davide and Ricardo Nunes (2010). « Fiscal policy under loose commitment ». In: *Journal of Economic Theory* 145.3, pp. 1005–1032.

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson (2012). « The Forward Guidance Puzzle ». In: *FRB of New York Staff Report*, pp. 3133–58.

Dovis, Alessandro and Rishabh Kirpalani (2020). « Fiscal Rules, Bailouts, and Reputation in Federal Governments ». In: *American Economic Review* 110.3, pp. 860–88.

Eggertsson, Gauti (2006). « The Deflation Bias and Committing to Being Irresponsible ». In: *Journal of Money, Credit and Banking* 38, pp. 283–321.

Eggertsson, Gauti and Michael Woodford (2003). « The Zero Bound on Interest Rates and Optimal Monetary Policy ». In: *Brookings Papers on Economic Activity* 34.1, pp. 139–235.

Farhi, Emmanuel and Iván Werning (2019). « Monetary Policy, Bounded Rationality, and Incomplete Markets ». In: *American Economic Review* 109.11, pp. 3887–3928.

Gabaix, Xavier (2020). « A Behavioral New Keynesian Model ». In: *American Economic Review* 110.8, pp. 2271–2327.

Galí, Jordi (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition*. Economics Books 10495. Princeton University Press.

García-Schmidt, Mariana and Michael Woodford (2019). « Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis ». In: *American Economic Review* 109.1, pp. 86–120.

King, Robert G. and Yang K. Lu (2022). *Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation*. NBER Working Papers 30763. National Bureau of Economic Research, Inc.

Kreps, David M and Robert Wilson (1982). « Reputation and imperfect information ». In: *Journal of Economic Theory* 27.2, pp. 253–279.

Kydland, Finn E. and Edward C. Prescott (1977). « Rules Rather than Discretion: The Inconsistency of Optimal Plans ». In: *Journal of Political Economy* 85.3, pp. 473–491.

Loisel, Olivier (2008). « Central bank reputation in a forward-looking model ». In: *Journal of Economic Dynamics and Control* 32.11, pp. 3718–3742.

Lu, Yang K., Robert G. King, and Ernesto Pasten (2016). « Optimal reputation building in the New Keynesian model ». In: *Journal of Monetary Economics* 84, pp. 233–249. ISSN: 0304-3932.

Lucas, Robert E. (1987). *Models of business cycles.* Yrjö Jahnsson lectures. Oxford: Basil Blackwell.

Maskin, Eric and Jean Tirole (2001). « Markov Perfect Equilibrium I. Observable Actions ». In: *Journal of Economic Theory* 100, pp. 191–219.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016a). « The Discounted Euler Equation: A Note ». In: *NBER Working paper.*

— (2016b). « The Power of Forward Guidance Revisited ». In: *American Economic Review* 106.10, pp. 3133–58.

Milgrom, Paul and John Roberts (1982). « Predation, reputation, and entry deterrence ». In: *Journal of Economic Theory* 27.2, pp. 280–312.

Nakata, Taisuke (2018). « Reputation and liquidity traps ». In: *Review of Economic Dynamics* 28, pp. 252–268. ISSN: 1094-2025.

Nakata, Taisuke et al. (2019). « Attenuating the forward guidance puzzle: Implications for optimal monetary policy ». In: *Journal of Economic Dynamics and Control* 105, pp. 90–106.

Stokey, Nancy L. (1991). « Credible public policy ». In: *Journal of Economic Dynamics and Control* 15.4, pp. 627–656.

Tauchen, George (1986). « Finite state markov-chain approximations to univariate and vector autoregressions ». In: *Economics Letters* 20.2, pp. 177–181.

Walsh, Carl E. (2018). « Simple Sustainable Forward Guidance at the ELB ». In: *Working paper.*

Werning, Iván (2015). *Incomplete Markets and Aggregate Demand.* Working Paper 21448. National Bureau of Economic Research.

Woodford, Michael (2003). *Interest and prices.* Princeton Univ. Press.

Appendices

1.A Interest Rate under Optimal Discretion

When the central bank is not honoring a promise[7], the interest rate is set to minimize the instantaneous loss function by taking expectations as given. We impose the constraint on the minimization problem that the interest rate cannot be negative. The interest rate chosen by the central banker will therefore be

$$i_t = \arg\min \ell(y_t, \pi_t)$$

subject to

$$\begin{aligned}
&\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa y_t \\
&y_t = \mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \\
&i_t \geq 0
\end{aligned}$$

Using the IS curve and the Phillips curve, we can rewrite the loss function of the central bank in terms of the nominal interest rate rather than inflation and output gap:

$$\ell(y_t, \pi_t) = \ell\left(\mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)\,,\, \beta\mathbb{E}_t[\pi_{t+1}] + \kappa\mathbb{E}_t[y_{t+1}] - \kappa\frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)\right)$$

We denote by $\zeta_t \leq 0$ the multiplier associated to the constraint $i_t \geq 0$. The Lagrangian

7. Either because no promise has been given or because the central bank decides to break it.

associated to the minimization problem of the central bank writes:

$$\mathcal{L}(i_t, \zeta_t) = \left(\beta \mathbb{E}_t[\pi_{t+1}] + \kappa \mathbb{E}_t[y_{t+1}] - \kappa \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)\right)^2$$
$$+ \lambda \left(\mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) - \bar{y}\right)^2$$
$$+ \zeta_t i_t$$

The first order condition is:

$$2\frac{\kappa}{\rho} \underbrace{\left(\beta \mathbb{E}_t[\pi_{t+1}] + \kappa \mathbb{E}_t[y_{t+1}] - \kappa \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)\right)}_{=\pi_t} + 2\frac{\lambda}{\rho} \underbrace{\left(\mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)\right)}_{=(y_t - \bar{y})} = \zeta_t$$

$$\iff 2\frac{\kappa}{\rho}\pi_t + 2\frac{\lambda}{\rho}(y_t - \bar{y}) = \zeta_t$$

So that substituting back π_t and y_t we obtain

$$\kappa\pi_t + \lambda\left(y_t - \bar{y}\right) = \frac{\rho}{2}\,\zeta_t \leq 0$$

Whenever the ZLB is binding, the multiplier ζ_t will generically be strictly negative. On the contrary, whenever the ZLB is slack, the multiplier ζ_t is exactly zero, and the FOC writes:

$$\lambda\left(y_t - \bar{y}\right) + \kappa\pi_t = 0$$

The second order condition is:

$$2\frac{\kappa^2}{\rho^2} + 2\frac{\lambda}{\rho^2} > 0$$

So that the FOC characterizes the unique global minimum.

1.B Equilibrium Description

We describe here in detail how the equilibrium is computed. For fixed strategies σ and beliefs μ, output gap, inflation and interest rates must constitute a solution to the New Keynesian model. The key differences are that interest rates are set to zero whenever a promise is kept, and that the expectations about future outcomes depend on the private sector assessment about the type of the central bank and how they behave.

1.B.1 At the ZLB

Whenever the economy is at the ZLB at time t, the nominal interest rate is set to zero, $i_t^Z = 0$. An Odyssean promise is given, and therefore the private sector forms an assessment about how likely it is that the promise will be kept. We have to consider three possibilities:

1. If the ZLB is binding again in the next period, then the promise is renewed. This corresponds to the event $\{r_{t+1}^n \in Z_{\sigma,\mu}\}$.
2. If the ZLB is not binding at time $t+1$, then each type δ of the central bank will keep the promise with probability $\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)$. As the type of the central bank is private information, the private sector forms an expectation with respect to δ using the probability measure μ_t. Thus conditional on exiting the ZLB, i.e. $r_{t+1}^n \notin Z_{\sigma,\mu}$, the expected probability of transitioning to the first promise-keeping period is $\mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)]$.
3. Finally, if the ZLB is not binding at time $t+1$, the conditional probability of transitioning to normal times is $\mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)]$.

Output gap, inflation and interest rates therefore satisfy the following system:

$$
\begin{cases}
\pi_t^Z = \beta\, \mathbb{E}_{r_{t+1}^n} \begin{bmatrix} \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ \quad + \pi_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \\ \quad + \pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \end{bmatrix} + \kappa y_t^Z \\
\\
y_t^Z = \mathbb{E}_{r_{t+1}^n} \begin{bmatrix} y_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ \quad + y_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \\ \quad + y_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \end{bmatrix} \\
\qquad - \frac{1}{\rho} \left(0 - \mathbb{E}_{r_{t+1}^n} \begin{bmatrix} \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ \quad + \pi_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \\ \quad + \pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \end{bmatrix} - r_t^n \right) \\
\\
i_t^Z = 0
\end{cases}
$$

For the ease of exposition we have suppressed the dependence on the current shock realization. A similar omission applies to the remaining cases.

1.B.2 Promise-keeping period with $P_t < \tau$

If the economy is away from the ZLB and a promise is to be kept at time t, the interest rate is set to zero, $i_t^P = 0$. In the next period, the ZLB could be binding, so a promise is renewed and the process starts over. If the economy stays away from the ZLB for one more period, a transition to the next period of promise keeping occurs with conditional probability $\mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)]$. The promise is broken in next period with conditional probability $\mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)]$.

Output gap, inflation and interest rates therefore satisfy the following system:

$$
\begin{cases}
\pi_t^P = \beta\, \mathbb{E}_{r^n_{t+1}} \begin{bmatrix} \pi^Z_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\} \\ \qquad + \pi^P_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}^{\mu_t}_{\delta}[\sigma(\delta, r^n_{t+1}, \mu_t, P_t+1)] \\ \qquad + \pi^D_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}^{\mu_t}_{\delta}[1-\sigma(\delta, r^n_{t+1}, \mu_t, P_t+1)] \end{bmatrix} + \kappa y_t^P \\
\\
y_t^P = \mathbb{E}_{r^n_{t+1}} \begin{bmatrix} y^Z_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\} \\ \qquad + y^P_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}^{\mu_t}_{\delta}[\sigma(\delta, r^n_{t+1}, \mu_t, P_t+1)] \\ \qquad + y^D_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}^{\mu_t}_{\delta}[1-\sigma(\delta, r^n_{t+1}, \mu_t, P_t+1)] \end{bmatrix} \\
\qquad -\frac{1}{\rho} \left(0 - \mathbb{E}_{r^n_{t+1}} \begin{bmatrix} \pi^Z_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\} \\ \qquad + \pi^P_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}^{\mu_t}_{\delta}[\sigma(\delta, r^n_{t+1}, \mu_t, P_t+1)] \\ \qquad + \pi^D_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}^{\mu_t}_{\delta}[1-\sigma(\delta, r^n_{t+1}, \mu_t, P_t+1)] \end{bmatrix} - r^n_t \right) \\
\\
i_t^P = 0
\end{cases}
$$

1.B.3 Last promise-keeping period $P_t = \tau$

If the economy is in the last promise-keeping period the interest rate is $i_t^P = 0$. There are only two possible cases in the following period: either the ZLB binds, or the ZLB does not bind and the central bank conducts discretionary policy. Output gap, inflation and interest rates therefore satisfy the following system:

$$
\begin{cases}
\pi_t^P = \beta\, \mathbb{E}_{r^n_{t+1}} \Big[\pi^D_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} + \pi^Z_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\} \Big] + \kappa y_t^P \\
\\
y_t^P = \mathbb{E}_{r^n_{t+1}} \Big[y^D_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} + y^Z_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\} \Big] \\
\qquad -\frac{1}{\rho} \Big(0 - \mathbb{E}_{r^n_{t+1}} \Big[\pi^D_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \notin Z_{\sigma,\mu}\} + \pi^Z_{t+1} \cdot \mathbb{1}\{r^n_{t+1} \in Z_{\sigma,\mu}\} \Big] - r^n_t \Big) \\
\\
i_t^P = 0
\end{cases}
$$

1.B.4 Normal times

Away from the ZLB and with no promise to be kept, the central bank conducts discretionary policy. Thus, the interest rate is set to satisfy the derived FOC:

$$-\kappa\, \pi_t^D = \lambda\, (y_t^D - \bar{y})$$

In the following period, the economy may either hit the ZLB or stay in discretion. Output gap, inflation and interest rates therefore satisfy the following system:

$$\begin{cases} \pi_t^D = \beta\, \mathbb{E}_{r_{t+1}^n}\Big[\pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z(\sigma)\} + \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z(\sigma)\}\Big] + \kappa y_t^D \\ \\ y_t^D = \ \mathbb{E}_{r_{t+1}^n}\Big[y_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z(\sigma)\} + y_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z(\sigma)\}\Big] \\ \qquad -\dfrac{1}{\rho}\Big(i_t^D - \mathbb{E}_{r_{t+1}^n}\Big[\pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z(\sigma)\} + \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z(\sigma)\}\Big] - r_t^n\Big) \\ \\ -\kappa\, \pi_t^D = \lambda\, (y_t^D - \bar{y}) \end{cases}$$

1.B.5 Decisions of the central bank

Once we have computed output gap, inflation and interest rates for fixed (σ, μ) we need to verify that the strategies σ are optimal.

When the economy is at the ZLB, the strategies σ are optimal if they solve the following Bellman equation:

$$\begin{aligned} L_t^Z(\delta, r_t^n, \mu_t, 0) = \ell\Big(y_t^Z(r_t^n, \mu_t, 0), \pi_t^Z(r_t^n, \mu_t, 0)\Big) \\ + \delta\, \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{c} \min\limits_{x \in \{0,1\}} \left\{ \begin{array}{l} x\, L_{t+1}^P(\delta, r_{t+1^n}, \mu_{t+1}, 1) \\ \quad + (1-x)\, L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \end{array} \right\} \cdot \mathbb{1}\{r_{t+1} \notin Z_{\sigma,\mu}\} \\ + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \end{array} \right] \end{aligned}$$

That is, the optimal action prescribes whether to transit to either discretion (break the promise) or promise-keeping. If the economy is in a promise keeping period with $P_t < \tau$, optimal strategies solve:

$$
\begin{aligned}
L_t^P(\delta, r_t^n, \mu_t, P_t) = \ell\Big(y_t^P(r_t^n, \mu_t, P_t), \pi_t^P(r_t^n, \mu_t, P_t)\Big) \\
+ \delta\, \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{c} \min\limits_{x \in \{0,1\}} \left\{ \begin{array}{l} x\, L_{t+1}^P(\delta, r_{t+1}^n, \mu_{t+1}, P_t + 1) \\ \quad + (1-x)\, L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \end{array} \right\} \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \\ + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \end{array} \right]
\end{aligned}
$$

When the economy is in the last promise-keeping period, no decision is taken taken with respect to the transition, so the Bellman equation writes:

$$
L_t^P(\delta, r_t^n, \mu_t, \tau) = \ell\Big(y_t^P(r_t^n, \mu_t, \tau), \pi_t^P(r_t^n, \mu_t, \tau)\Big) + \delta \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \\ \qquad + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \end{array} \right]
$$

Finally, if the economy is in discretion, no strategic decision is made, so the Bellman equation is:

$$
L_t^D(\delta, r_t^n, \mu_t, 0) = \ell\Big(y_t^D(r_t^n, \mu_t, 0), \pi_t^D(r_t^n, \mu_t, 0)\Big) + \delta \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \\ \qquad + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \end{array} \right]
$$

1.C Proofs

Proof of Proposition 1.4.1. The myopic central banker only compares the static loss $\lambda y_t^2 + \pi_t^2$. If she breaks her promise, output and inflation are given by:

$$\pi_D = \beta \pi_D^e + \kappa y_Z$$
$$0 = \lambda y_D + \kappa \pi_D$$

Where (π_D^e, y_D^e) are inflation and output expectations when the central banker breaks the promise (they are fixed once the promise is broken). The IS curve is irrelevant, because the shock is absorbed by the nominal interest rate (and this can be done as we assume the ZLB is not binding). Solving for CB's FOC and Phillips Curve:

$$\pi_D = \pi_D^e \frac{\beta\lambda}{\lambda + \kappa^2}$$
$$y_D = -\pi_D^e \frac{\beta\kappa}{\lambda + \kappa^2}$$

The static loss is:

$$\ell_D = \lambda y_D^2 + \pi_D^2 = \frac{\lambda\beta^2}{\lambda + \kappa^2} \left(\pi_D^e\right)^2$$

The loss decreases to zero as $\lambda \to 0$. The nominal interest rate should be positive:

$$i_D = r + \pi_D^e + \rho y_D^e + \frac{\rho\beta\kappa}{\lambda + \kappa^2} \pi_D^e > 0$$

If she keeps her promise, output and inflation are given by:

$$y_Z = y_Z^e - \frac{1}{\rho}(0 - r - \pi_Z^e)$$
$$\pi_Z = \beta \pi_Z^e + \kappa y_Z$$

Where (π_Z^e, y_Z^e) are inflation and output expectations when the central banker keeps the promise. Solving:

$$y_Z = y_Z^e + \frac{\pi_Z^e}{\rho} + \frac{r}{\rho}$$
$$\pi_Z = \beta \pi_Z^e + \kappa y_Z^e + \frac{\kappa}{\rho} \pi_Z^e + \frac{\kappa}{\rho} r$$

So the loss is:

$$\ell_Z = \lambda \left(y_Z^e + \frac{\pi_Z^e}{\rho} + \frac{r}{\rho} \right)^2 + \left(\beta \pi_Z^e + \kappa y_Z^e + \frac{\kappa}{\rho} \pi_Z^e + \frac{\kappa}{\rho} r \right)^2$$

This loss does not vanish for $\lambda \to 0$:

$$\lim_{\lambda \to 0} \ell_Z = \left(\beta \pi_Z^e + \kappa y_Z^e + \frac{\kappa}{\rho} \pi_Z^e + \frac{\kappa}{\rho} r \right)^2$$

which is generically positive.

Finally, the map $\ell_D(\lambda) - \ell_Z(\lambda)$ is continuous in λ, so $\ell_D(0) < \ell_Z(0)$ guarantees that there exists some $\lambda^* > 0$ such that $\ell_D(\lambda) < \ell_Z(\lambda)$ for all $\lambda \in [0, \lambda^*)$. □

Proof of Proposition 1.4.2.
A central banker with type δ will not keep her promise if

$$\ell(y_t^D, \pi_t^D) + \delta \mathbb{E}_t^{\chi} \left[L\left(\{y_i\}_{i=t+1}^{\infty}, \{\pi_i\}_{i=t+1}^{\infty} \right) \right] < \ell(y_t, \pi_t) + \delta \mathbb{E}_t^{\tilde{\chi}} \left[L\left(\{y_i\}_{i=t+1}^{\infty}, \{\pi_i\}_{i=t+1}^{\infty} \right) \right]$$

where χ denotes the probability distribution over paths of inflation and output gap if the promise is not kept, while $\tilde{\chi}$ denotes the probability distribution when the promise is kept.

It is easy to verify that the left-hand side is bounded while the right-hand side grows unboundedly in r_t^n for $i_t = 0$. Hence, there is a value of r_t^n such that even the patient central banker will have to deviate from her promise. □

Proof of Proposition 1.4.3.
Given the pooling strategies and the off-path beliefs, the private sector assigns probability zero to promises being kept. Thus, promise-keeping induces no change in expected output gap and inflation, implying that there are no gains from deviating. □

To prove Proposition 1.4.2, we use the following lemma:

Lemma 1.C.1. *The loss is bounded if the central bank conducts optimal discretionary policy.*

Proof of Lemma 1.C.1.
Losses are bounded for negative shock realizations, as they are assumed to be bounded. Moreover, the central bank can always set $i_t = r_t^n$ whenever the natural rate of interest is positive.

This bounds from above the values of output gap and inflation for fixed expectations. Furthermore, this policy being infinitely repeated whenever possible bounds expectations from above as well. The instantaneous loss is thus:

$$\ell\left(\mathbb{E}_t[y_{t+1}] + \frac{1}{\rho}\mathbb{E}_t[\pi_{t+1}], \left(\beta + \frac{\kappa}{\rho}\right)\mathbb{E}_t[\pi_{t+1}] + \kappa\mathbb{E}_t[y_{t+1}]\right)$$

Which is bounded as all the terms are bounded. □

Proof of Proposition 1.4.4.

A benevolent central banker ($\delta = \beta$) will not keep her promise if

$$\ell(y_t^D, \pi_t^D) + \beta\mathbb{E}_t^{\chi}\left[L\big(\{y_i\}_{i=t+1}^{\infty}, \{\pi_i\}_{i=t+1}^{\infty}\big)\right] < \ell(y_t^{PK}, \pi_t^{PK}) + \beta\mathbb{E}_t^{\tilde{\chi}}\left[L\big(\{y_i\}_{i=t+1}^{\infty}, \{\pi_i\}_{i=t+1}^{\infty}\big)\right]$$

where χ denotes the probability distribution over paths of inflation and output gap if the promise is not kept, while $\tilde{\chi}$ denotes the probability distribution when the promise is kept. We notice that the LHS is independent of the promised length τ. As the RHS increases without bounds in the promised length of zero interest rates, a too generous promised cannot be sustained. □

Proof of Corollary 1.4.1.

Follows from Proposition 1.4.4. □

1.D Forward guidance and the variance of the demand shock

We vary the standard deviation (ς) of the innovation ξ_t to assess its effect on credibility and power of forward guidance.

In Figure 1.8 we show the effect on the output gap at the ZLB of promising a zero interest rate for a time horizon of $\tau = 1$ periods as a function of the standard deviation of the innovation to the demand shock. We also show the effect of forward guidance with full commitment in red for comparison.

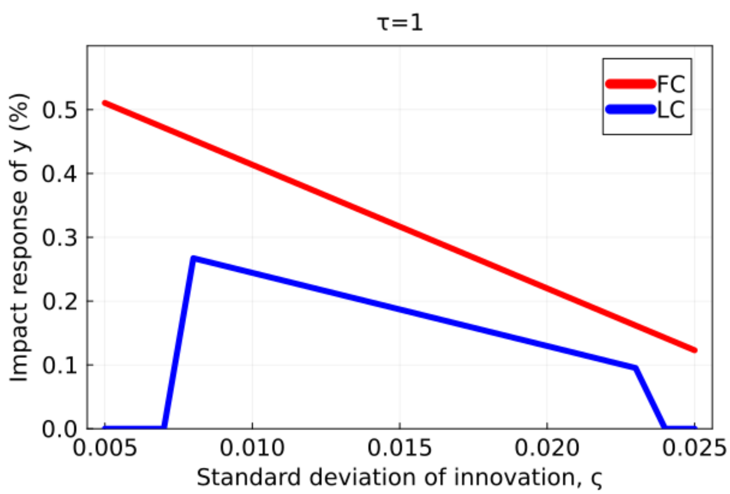

Figure 1.8 – Forward guidance and the standard deviation of the innovation to the demand shock

We see that there is a range of values for the standard deviation ς for which FG can be sustained in equilibrium. In particular, we notice that for a standard deviation of less than 0.007, forward guidance cannot be sustained in equilibrium, as well as for a standard deviation of more than 0.023.

Conditional on being sustainable, we find that the higher the standard deviation, the smaller the effect of FG. This occurs as we consider promises of keeping the rate at zero, and interest rates are already low for high levels of the standard deviation, as can be seen from Figure 1.9. Thus, the promise is less expansionary by construction.

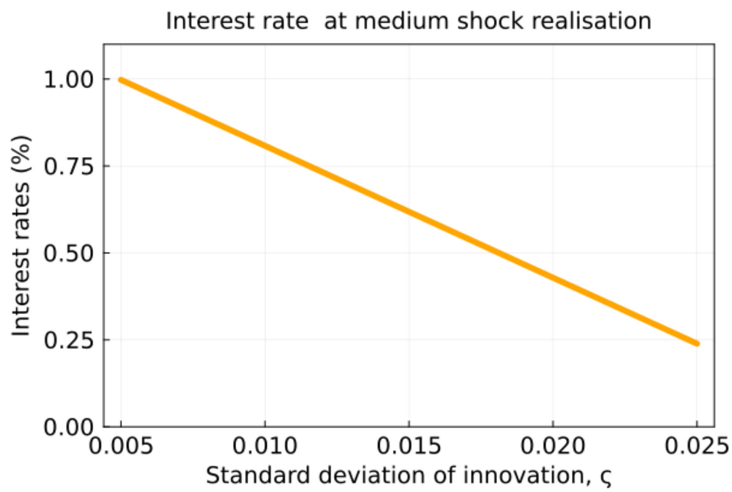

Figure 1.9 – Nominal interest rates with medium realizations of the demand shock.

Chapter 2

The Distributional Consequences of Property Taxes

Chenchen Huang[1] & Emil Mortensen[2]

Abstract

We document three stylized facts: 1) the expenditure share of housing is decreasing in income, 2) the portfolio share of housing is inversely u-shaped in wealth, and 3) conditional on home ownership the portfolio share of housing is decreasing in wealth. Motivated by these stylized facts we propose a model where households have non-homothetic preferences over housing and non-durable consumption. We use our model to study the distributional consequences of a property value tax. We find that a higher property value tax increases wealth inequality while reducing consumption inequality. While this may appear puzzling at first, it is due to the fact higher property taxes discourage future home ownership for wealth-poor households, and they respond by reducing savings and increasing consumption of non-durables.

Keywords: Inequality, heterogeneous agents, property taxation

JEL classification: E21, H20

1. Toulouse School of Economics. Email: chenchen.huang@tse-fr.eu

2. Toulouse School of Economics. Email: emitm@outlook.dk. We are grateful to Fabrice Collard for providing useful feedback on earlier versions of this paper. We also thank Martial Dupaigne, Eugenia Gonzalez-Aguado, and Christian Hellwig for useful comments and discussion.

2.1 Introduction

Housing occupies a unique position in household portfolios: it serves both as a consumption good and a store of wealth. Poor households tend to rent with a homeownership rate of only around 15 % for the three lowest wealth deciles, but housing is the most prominent asset in the portfolio of the middle class having a portfolio share of more than 50 %, while richer households tend to hold other forms of wealth with a portfolio share of housing less than 50 %. In addition to being an important part of households' portfolio, housing wealth is also a prominent source of income for governments, and in particular local governments, given that property is immobile and difficult to avoid Norregaard (2013). In the United States, whilst over 10% of overall tax revenue is generated from property tax revenues,[3] property taxes account for 72% of local government tax revenues.[4]

Property taxes are often considered a progressive form of taxation based on two premises. First, renters are less likely to be affected than homeowners owing to the potential for imperfect pass-through. Second, the home-ownership rate increases with wealth. In this paper, we revisit the implications of a property value tax on the distribution of wealth and household welfare. In particular, we ask the following question: are property taxes effective at redistribution? In the context of rising wealth inequality and growing concerns about housing affordability, understanding the distributional and welfare effects of property taxation is critical for designing an equitable and efficient property taxation system.

In this paper, we first want to characterize the role of housing in the wealth distribution in order to capture how changes in property taxes affect households at different wealth deciles. We investigate the distribution of housing in total wealth in the US using micro data from the Survey of Consumer Finances (SCF) and the American Housing Survey

3. https://taxfoundation.org/data/all/federal/us-tax-revenue-by-tax-type/

4. https://taxfoundation.org/research/all/state/local-tax-revenue-local-sales-taxes-local-income-taxes/

(AHS). The SCF provides us with detailed information on wealth, including housing wealth, while it is less informative on housing costs. Instead, we rely on the AHS for information on housing costs. We document three empirical patterns: (1) the share of income spent on housing declines with income; (2) the share of housing in total assets is inversely U-shaped with respect to wealth, and (3) conditional on homeownership the portfolio share of housing is decreasing in wealth.

Motivated by these patterns, we develop a heterogeneous agents model with idiosyncratic income risk such that our model features a distribution over net worth, which will allow us to study the welfare implications for agents located at different wealth deciles. We model housing as a discrete choice over different housing sizes, where there are significant transaction costs associated with trading houses, and to purchase a house, agents face a down payment requirement. Agents are allowed to rent if they choose not to own housing. A key feature of our model in contrast to much of the existing literature is that we assume preferences over housing and non-durable consumption are non-homothetic. In particular, we let housing be a necessity good, i.e. the income elasticity of housing is less than one, to align with our empirical findings. Age is another important variable that can determine whether to own or rent. To capture this, we model the life-cycle in our model, but we choose to do this by modeling aging stochastically which allows us to reduce the dimensionality of the state space while still allowing the expected income to grow over time. The model features a government that levies a proportional property value tax of one percent on owner-occupied housing. We then shock this tax level to study how it affects households across the wealth distribution. We choose to calibrate the model to match the aforementioned empirical facts on housing and the wealth distribution as well as possible.

We then use our model to study the effect of increasing property taxes. We increase it from 1 percent in the baseline model to 2 and 3 percent, respectively. We find that increasing the property tax rate leads to higher wealth inequality, while consumption of non-durables becomes more equitable. While it may appear puzzling at first that property taxes have a different sign effect on wealth and consumption inequality, it is due to the fact that many households in the middle of the wealth distribution choose to become renters rather than home owners in response to increasing taxes. This, in turn, reduces the incentive to save for homeownership for wealth poor households. Rather than saving for homeownership, they choose to increase the consumption level of non-durable goods.

We further study how the effect of a property tax increase depends on our preference assumption. We find that when preferences are homothetic, the sign of the effect on consumption inequality is reversed such that higher property taxes increase consumption inequality, but on the other hand, the effect on wealth inequality is muted. This indicates that the role of homothetic preferences is crucial in determining the distributional impacts of tax changes.

We also investigate the role of tax pass-through onto rents. If increases in property taxes raise rents, the budgets of low-income renters are affected as well as those of homeowners. For all of the results so far, we have assumed that there is perfect pass-through of property taxes on rents. To check the importance of this assumption, we fix rents to the level in the baseline scenario, and thereby assume that all of the tax incidence falls on landlords rather than renters. This assumption turns out to matter in important quantitative ways. The drop in the homeownership rate when the property tax is increased from 1 to 3 percent is 7 percentage points higher when rents are fixed. This is due to a substitution effect whereby renting becomes cheaper than owning. While increasing property taxes still increases wealth inequality when rents are fixed, this change also reduces consumption inequality significantly more compared to when tax increases are passed through to renters.

The paper is organized as follows. Section 2.2 relates our paper to the existing literature. Section 2.3 introduces the data that we use, and some empirical evidence. Section 2.4 introduces the model. Section 2.5 presents our calibration and the solution method we apply to solve the model. Section 2.6 discusses our results. Section 2.7 concludes.

2.2 Related Literature

This paper contributes to three different strands of literature: 1) a vast and rapidly growing literature in macroeconomics that explicitly models housing, 2) the optimal portfolio allocation of households with respect to housing, and 3) we relate to the literature on optimal taxation.

Most papers in the early macroeconomics literature tended to not explicitly include housing. In fact, housing was barely mentioned in the first *Handbook of Macroeconomics* of 1999 according to Piazzesi and Schneider (2016). This has drastically changed, and there is now a large and growing literature on this subject. For an excellent review, see Piazzesi and Schneider (2016). Papers in macro and housing differ in terms of how they model i) the housing market, ii) preferences, and iii) the life cycle.

Regarding i) the housing market, many papers now build on Aiyagari (1994) by augmenting this type of model with a discrete choice over housing. Some papers assume no rental markets such as Favilukis, Ludvigson, and Van Nieuwerburgh (2017), while other papers assume the presence of both rental and owner-occupied housing such as Kaplan, Mitman, and Violante (2020), Sommer, Sullivan, and Verbrugge (2013), Bontemps, Cherbonnier, and Magnac (2023), and Sommer and Sullivan (2018). In terms of modeling the housing market, our paper is most closely related to Kaplan, Mitman, and Violante (2020) as we follow them in assuming that equilibrium rents are determined by a user-cost formula derived from a no-arbitrage condition. While Kaplan, Mitman, and Violante (2020) allow for renting units of multiple sizes, we only allow for renting a small house.

Related to ii) preferences, our paper is close to Cioffi (2021), who documents that the expenditure share of housing is decreasing in income, and then studies the role of non-homothetic preferences in shaping wealth inequality. Our question of interest is different, and we further contribute to the paper by showing how housing's share of total assets is inversely U-shaped in wealth. These findings are central to our argument that taxing housing wealth may unintentionally worsen inequality. The vast majority of the other papers in the housing literature assume homothetic preferences.

In terms of iii) the life cycle, we choose to model age stochastically following Sommer and Sullivan (2018), while there is a large literature that incorporates a proper life-cycle profile by augmenting the state space with aging. Some prominent papers are Iacoviello and Pavan (2013), which study a life-cycle model of housing to better understand the cyclical properties of mortgage debt, and Halket and Vasudev (2014), which develops a life-cycle model to study the importance of financial constraints for homeownership rates.

We also contribute to the literature that has the portfolio allocation of households at the heart of the problem they study. One important paper in this literature is Flavin and Yamashita (2002), which stresses how housing is the single most important asset for most home owners. Another example is Bacher (2023), which shows the importance of marital status for housing choices. Compared to her paper, we show that once wealth and income are conditioned on, the impact of marital status is significantly reduced. We follow her assumption of a fully elastic housing supply curve, which simplifies the computation of the equilibrium as the housing price is unresponsive to demand conditions. Fischer and Khorunzhina (2019) study the optimal choice of housing in the context of divorce risk. In our paper, we do not allow for such family shocks, which is similar to Halket and Vasudev (2014) that studies the importance of financial constraints to homeownership over the life cycle. While family shocks may be important for understanding the behavior of households' housing decisions over the life cycle, as documented by Fisher and Gervais (2011), we consider wealth dynamics alone to be able to capture the distributional consequences of property taxes.

Our work connects to the literature on optimal taxation, particularly taxation of capital and wealth. Classical results such as Chamley (1986) and Judd (1985) suggest that capital taxation is inefficient in the long run. However, more recent work by Straub and Werning (2020) revisits these conclusions. Another important insight is the uniform taxation principle by Atkinson and Stiglitz (1976) that states that access to fully non-linear income taxes would be sufficient, and we therefore do not need to rely on commodity taxes, such as taxation of housing services. Other papers have studied the importance of non-homothetic preferences for optimal taxation such as Christiansen (1984).

Finally, the paper which is most closely related to ours is Sommer and Sullivan (2018), as we both consider a model with stochastic aging following Heathcote (2005), and we study the impact of taxation. Yet, there are some key differences. Sommer and Sullivan (2018) model preferences as being homothetic over housing and non-durable consumption, while we assume that preferences are non-homothetic, which is a key distinction from their paper. Further, we model the rent price as being determined by a no-arbitrage condition while they explicitly solve for an equilibrium rent price under the assumption that all house sizes can either be owned or rented. Another important difference is that we do not allow households to downsize their housing consumption without selling their home, while Sommer and Sullivan (2018) allow for this through renting out parts of their housing stock. We further introduce a taste shock to housing which is not present in Sommer and Sullivan (2018), that allows us to capture idiosyncratic variation in households' housing choices and therefore better match the observed heterogeneity in ownership rates. In our work, we consider the role of property taxes in general, whilst Sommer and Sullivan (2018) focuses on the role of the mortgage interest deduction. Sommer and Sullivan (2018) find that removing the mortgage interest tax deduction for owner-occupiers leads to an increase in homeownership, which is generated as in their model, house prices fall in equilibrium leading to substitution towards homeownership and away from the rental sector. By comparison, our model considers an increase in property taxes in general, rather than an increase only on owner-occupiers.

2.3 Empirical Evidence

In this section, we first introduce the data that we use for this paper and then move on to present some facts regarding the role of housing across the wealth distribution.

2.3.1 Data

To study the impact that property taxes have on inequality, we need to properly capture the importance of housing in the portfolio over the wealth distribution. For instance, it is important to correctly capture the fact that people in the lower end of the distribution tend to rent, and therefore are only affected by property taxes insofar as they influence rents. Further, we must look at the expenditure share of housing over the income distribution as this will be important for understanding how increased property taxes will affect consumption inequality.

In order to look into these facts, we need rich micro data on households' portfolio and wealth, and we need data on housing costs. For wealth data, we use the SCF as this specifically makes an effort to include wealthy individuals that are often under-represented in survey data. Dettling et al. (2015) find that SCF estimates align closely with macroeconomic benchmarks, providing reliable measures of household wealth and income aggregates. One drawback of the SCF is that it does not contain detailed data on housing costs. We overcome this deficiency by using the American Housing Survey. It contains rich data on housing costs and income while lacking wealth data. For the purpose of this project, we thus choose to work with both data sources as they allow us to capture the portfolio share of housing over the wealth distribution as well as the expenditure share of housing over the income distribution. Below, we describe briefly the data sources as well as the samples used in this paper.

The Survey of Consumer Finances

The Survey of Consumer Finances (SCF) is a cross-sectional dataset that collects data on the financial status, income, wealth, and demographic characteristics of U.S. households. It includes information on a wide range of financial topics such as assets (e.g., housing, retirement savings, and directly held stocks etc), liabilities (e.g., mortgage debt, credit card debt etc.), and income. The survey is conducted by the Federal Reserve Board every three years. We work with the most recent survey from 2022 in this paper.

Sample restrictions: We focus on households with heads that are in the age-interval 25-65, and have positive labor income. We choose to not include retirees as their housing decisions may be primarily driven by bequest motives as well as widowhood, which we do not model to keep the model simple. We further only consider households that have positive net worth as we follow Aiyagari (1994) in assuming a borrowing constraint for the safe asset which we set to zero implying that no households have negative net worth in our model.

The sample restrictions are inconsequential for the main results. In appendix 2.A, we reproduce the same graphs using respectively our chosen sample and the full sample of the SCF. As is clear from the appendix, the results are very similar.
For the purpose of this project, we define income as the sum of labor income received by all household members to have a direct mapping from productivity in the model which determines labor income to the income measured in the data. For net worth, we follow the definition of the SCF [5].

In Table 2.1 below we present descriptive statistics for the final sample that we use.

5. For an overview of the included variables, the reader is referred to the flow chart provided by the SCF which we report in appendix A.

	Mean	SD	Min	Max	N
age	46.41	11.62	25.00	65.00	14,784
homeowner	0.69	0.46	0.00	1.00	14,784
income	162,734.66	657,161.03	0.00	272,649,198.40	14,784
networth	959,565.66	5982839.50	0.00	1529792000.00	14,784
houses	330,762.53	542,098.31	0.00	63,160,000.00	14,784
asset	1118347.34	6094618.62	0.00	1542822000.00	14,784
portfolio share of housing	0.44	0.35	0.00	1.00	14,700

Table 2.1 – Summary Statistics SCF

One interesting feature of our summary statistics is the importance of housing. We notice that almost 70 percent of all American households are home owners. Further, the value of housing is more than half of net worth, and not far from half of total assets making housing a crucial asset in the portfolio of most households. Property taxes must therefore be expected to have important consequences for the wealth distribution. We further find that, on average, the stock of housing is around twice the yearly income of households. This implies that the monthly cost of housing, as determined for instance by property taxes, are potentially important components of the household budget and can therefore have substantial effects on household consumption.

The American Housing Survey

The American Housing Survey (AHS), sponsored by the Department of Housing and Urban Development (HUD), is conducted by the U.S. Census Bureau. It provides information on the size, composition, and quality of US housing. The AHS is a longitudinal housing unit survey conducted biennially in odd-numbered years. We will be using the survey wave in 2023 for the purposes of this paper [6]. It provides information on the costs of financing and maintaining homes, and the characteristics of people who live in

6. Ideally, we would have liked to choose the same year for the AHS and SCF, but as the AHS are only available biannually and the SCF tri-annually, we would have needed to choose 2019. Rather than going further back in time, we opted for the more recent data.

these homes, which is especially useful for this project.

Sample restrictions: to be consistent with the Survey of Consumer Finances, we focus on households with heads that are in the age-interval 25-65. We further restrict the sample to housing units that are currently occupied.

For housing costs we simply follow the definition used by the American Housing Survey. That implies that housing costs are defined as the sum of the following variables: mortgage expenses, rent, utilities, property taxes, insurance, homeowners association fees, lot rent, and routine maintenance costs. For income, we use the sum of all types of pre-tax income for all household members age 16 and over.

In Table 2.2 below, we report some summary statistics for the final sample that we use.

	Mean	SD	Min	Max	N
Age	46.42	11.52	25.00	65.00	29,433
Housing expenditures	23,489.01	18,930.58	0.00	776,652.00	29,433
Housing expenditures, owner	26,826.08	21,320.32	0.00	776,652.00	17,194
Housing expenditures, renter	19,233.03	13,545.45	120.00	176,448.00	11,916
Income	133,779.49	164,505.31	3,000.00	6,445,000.00	29,433
Housing costs to income	0.26	0.18	0.00	1.00	29,433
Housing costs to income, owner	0.22	0.16	0.00	1.00	17,194
Housing costs to income, renter	0.33	0.19	0.00	1.00	11,916

Table 2.2 – Summary Statistics AHS

We see from Table 2.2 that housing indeed is important for the household budget. About 25 percent of income is spend on housing expenditures on average. As we shall later see, this varies a lot across the income distribution.

2.3.2 Housing and the Wealth Distribution

In this section we present motivating evidence for the model. First, we find that housing expenditure as a share of income is decreasing in income. Second, we find that the home ownership rate is increasing in wealth. Third, we find that there is an inverse u-shaped relationship for housing as a share of assets, as wealth increases.

Expenditure share of Housing decreases w/ income. Using the AHS, we calculate the share of housing expenditures to income, and we present the average share for each income decile in Figure 2.1.

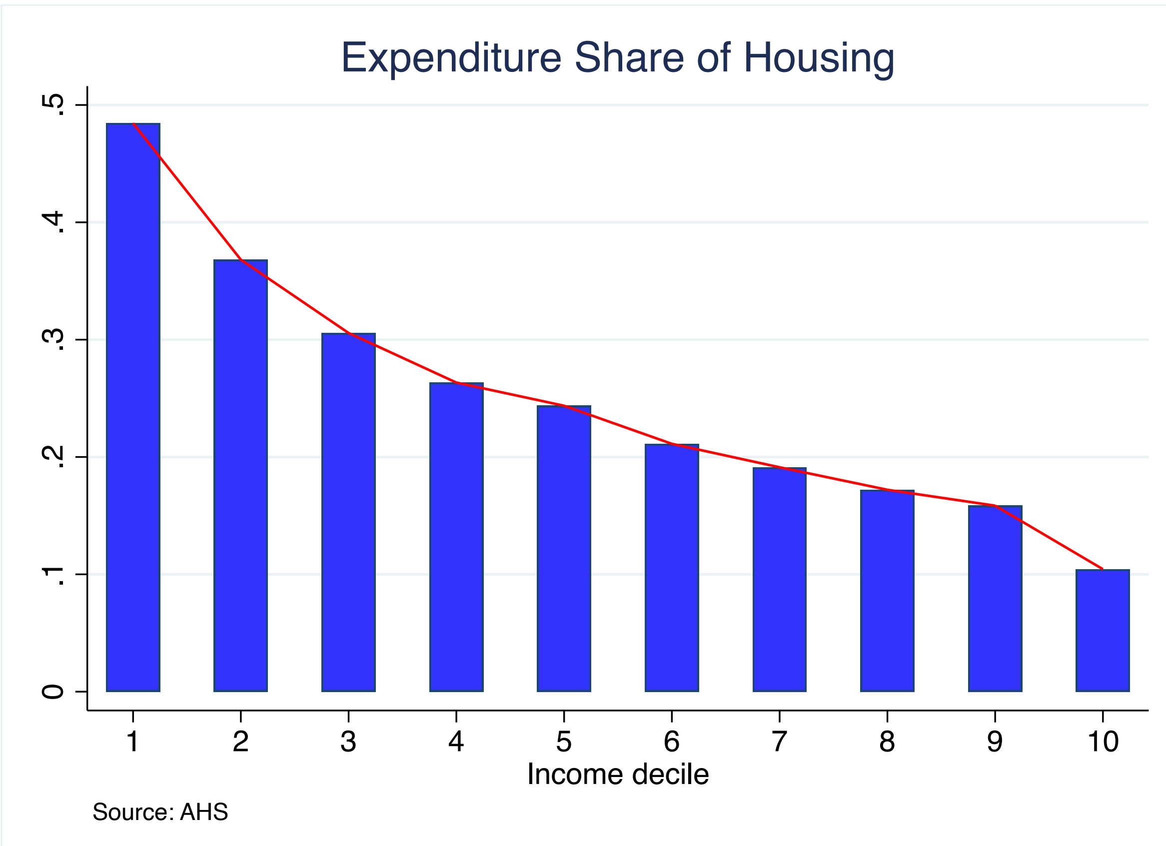

Figure 2.1 – Expenditure share

As we find that the expenditure share is decreasing in income we will follow Cioffi (2021), and model housing as a necessity good by using a utility function that allows for preferences to be non-homothetic over housing and non-durable consumption.

Home ownership is S-shaped in Wealth. Using the SCF, we calculate the average home-ownership rate for each wealth decile. We show these results in Figure 2.2. The home ownership rate stays close to 0 for lower wealth deciles, and then sharply rises around wealth decile 3. For households in the 5th wealth decile and above the home ownership rate stays at about 0.9. Hence, looking at the extensive margin we find that higher wealth is associated with an increase in the propensity to own a house rather than to rent.

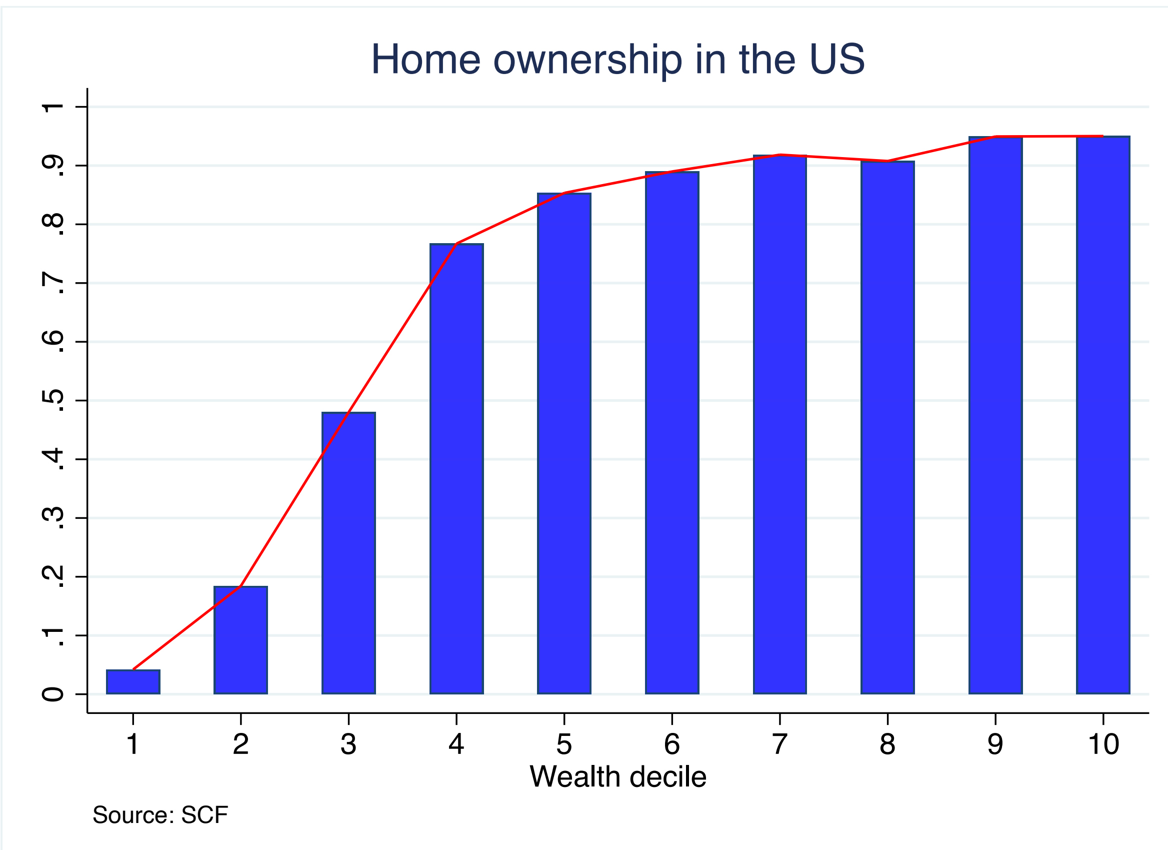

Figure 2.2 – Home ownership rate

As individuals' wealth decile may be confounded with other factors that affect home ownership, we regress the home ownership rate on wealth decile while controlling for other potentially important factors. Specifically, we estimate equation 2.1 by OLS, where we include controls for marital status, age of the head, income decile, race, and education. On account of the non-linear relationship between wealth decile and home ownership that is evident from Figure 2.2, we choose to include a dummy variable for wealth decile rather than estimating a linear regression. As we include a constant, we omit wealth decile 1.

$$HO_i = \beta_0 + \sum_{j=2}^{10} \alpha_j \cdot \text{wealth decile}_{ij} + X_i\beta + \epsilon_i \tag{2.1}$$

We show the full results with the estimated coefficients and standard errors for all of the variables in Table 2.9 in Appendix 2.C. In Figure 2.3, we plot the estimated coefficients for the different wealth deciles.

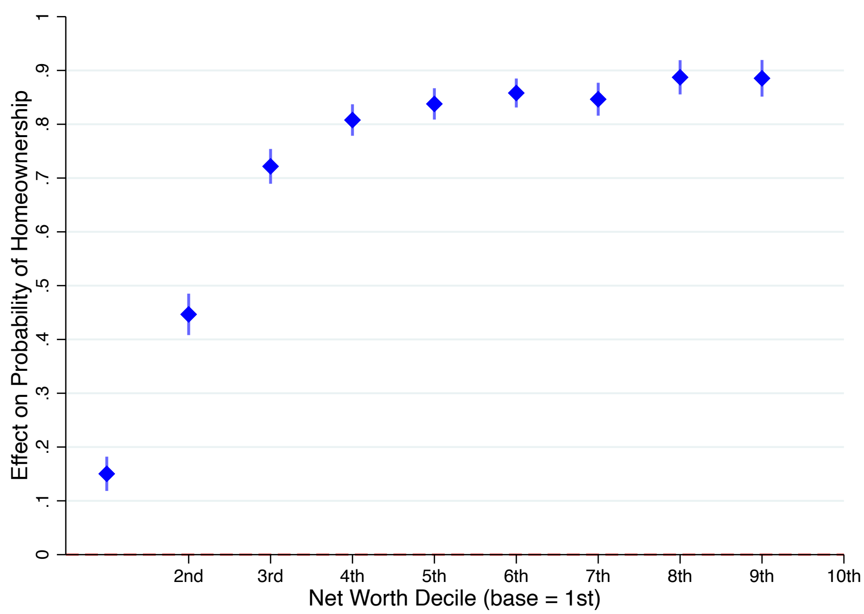

Figure 2.3 – Effect of Wealth Decile on Home Ownership

We find that the wealth decile strongly correlates with homeownership rate even after adding control variables.

We notice that the S-shape is preserved. There is a large increase in the probability of home ownership for lower wealth deciles, while the effect stays fairly constant around wealth decile 5, though slightly increasing.

Housing is inversely u-shaped. Also using the SCF, we compute housing wealth as a share of households' assets. We find that there is an inverse u-shaped relationship between the portfolio share of housing and wealth. Hence, while on the extensive margin there is a positive correlation between homeownership and wealth, we find that the intensive margin shows a different pattern. We present the average housing wealth share of assets for each wealth decile in Figure 2.4, which shows this inverse u-shape relationship.

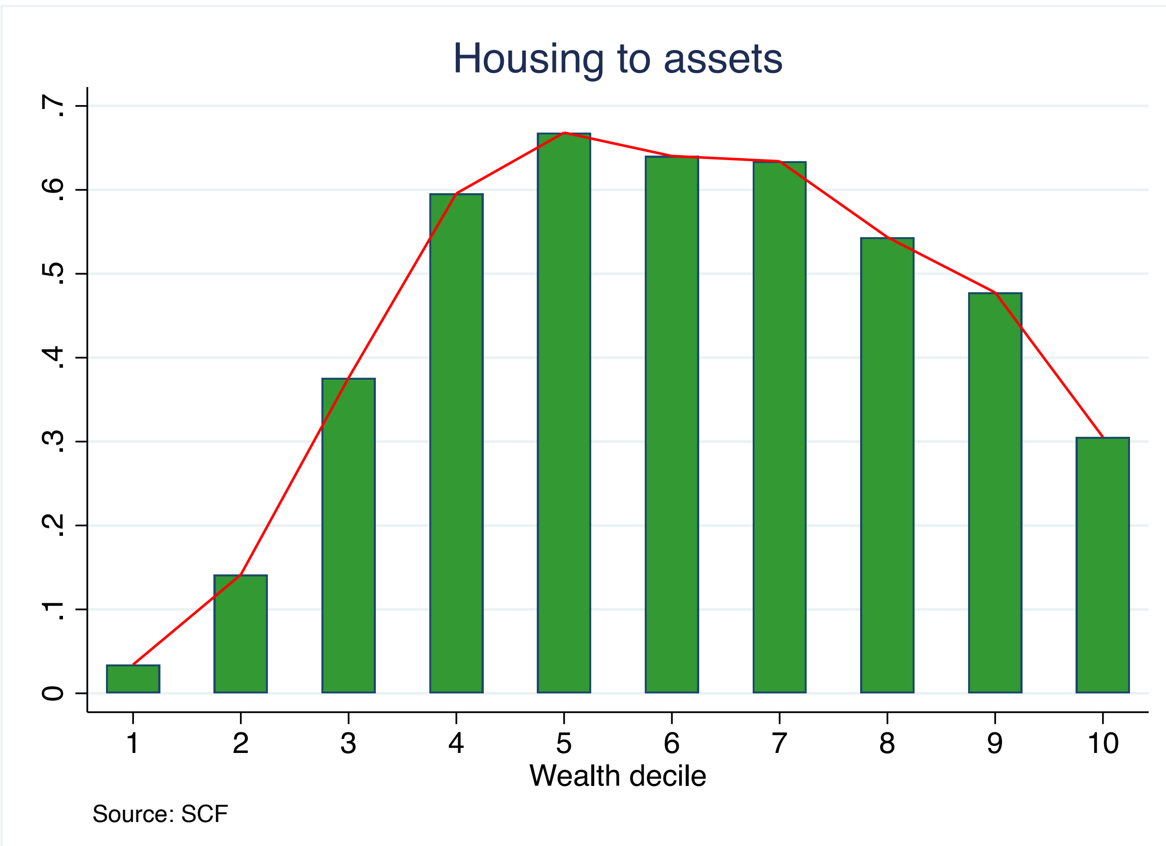

Figure 2.4 – Portfolio allocation

The reason for this inverse u-shaped relation is that despite the fact that the home ownership rate is increasing in wealth decile, households invest proportionally less in housing relative to other assets conditional on home ownership as illustrated in Figure 2.5 where we compute the average housing wealth as a share of households' net worth, conditional on home ownership, for each wealth decile.

Figure 2.5 – Portfolio allocation

We further estimate the following linear regression model of the portfolio share of housing (housing to assets = HtA) for the subsample of homeowners

$$HtA = \alpha \cdot \text{wealth decile}_i + X_i\beta + \epsilon_i \tag{2.2}$$

We report the results in Table 2.10 in Appendix B where we confirm the decreasing relationship between the portfolio share of housing and the wealth decile of individuals.

Taking stocks: we have documented three important facts that are relevant for assessing the distributional impact of property taxes on inequality. First, we have shown how the expenditure share of housing is decreasing over the income distribution. Second, we have found that while home ownership is increasing in wealth share, the portfolio share of housing is inversely U-shaped. Third, we find that this is driven by the fact that conditional on home ownership, higher wealth individuals allocate a smaller share of their portfolio to housing. We therefore aim to lay out a model that is able to replicate these key facts, and then use the model to study the distributional consequences of property taxes.

2.4 Model

In this section, we lay out the model that we will use for studying the distributional effects of property taxation. We study an Aiyagari-style model augmented with a housing sector, where households have non-homothetic preferences over housing and non-durable consumption. This framework allows us to match the empirical pattern of a declining share of housing expenditures as income increases. To keep the model simple, we model mortgage debt as an interest-only loan. Households are not permitted to act as landlords. Instead, a separate rental sector supplies fully elastically small apartment units, each smaller than the minimum size of an owner-occupied house.

2.4.1 Demographics and Income

We suppose that households are born as renters with zero wealth and housing. Upon death all remaining assets are taxed at 100 %, and the proceeds are used to finance wasteful government spending that does not affect the utility of households.

We follow Sommer and Sullivan (2018) in that we consider an overlapping generations economy with stochastic aging. It allows the expected household income to rise over time, without the need to keep track of the household's age as a state variable. Households transition between discrete labor income levels over time due to two mutually exclusive stochastic events: (i) aging and (ii) productivity shocks:

$$\Pi = \begin{bmatrix} 0 & \chi_1 & 0 & 0 \\ 0 & \ddots & \ddots & 0 \\ 0 & 0 & 0 & \chi_{J-1} \\ \chi_J & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1-\chi_1 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1-\chi_J \end{bmatrix} P \tag{2.3}$$

The first part of the Π matrix denotes the probability of aging which is associated with an increase in income except for the people with the highest productivity that dies, and are then reborn as poor renters with no savings. In this equation, the aging shock is given by $\chi_i = 1/(p_i L)$ where p_i is the fraction of the population with productivity w, and L is equal to the expected lifetime. The second part of the Π matrix is the probability that the worker transitions to a different income level due to an idiosyncratic income shock just as in baseline Aiyagari (1994) model.

2.4.2 Households

One important take-away from the empirical section is that housing is a necessity good. To capture this in the model, we follow Cioffi (2021) and introduce a non-homothetic utility function over housing and non-durable consumption. In particular, we suppose that the preferences take the following form:

$$U(c,s) = \frac{u(c,s)^{1-\sigma}}{1-\sigma}$$
$$u(c,s) = \left(\alpha \frac{1-\epsilon_h^{-1}}{1-\epsilon_c^{-1}} c^{\frac{\epsilon_c - 1}{\epsilon_c}} + (1-\alpha) s^{\frac{\epsilon_h - 1}{\epsilon_h}}\right)^{\frac{\epsilon_h}{\epsilon_h - 1}}$$

where s is shelter service, and c is non-durable consumption. α captures the relative preference for housing services relative to non-durable consumption, and σ is the coefficient of relative risk aversion. ϵ_c, ϵ_h govern the income elasticity of consumption and housing. When $\epsilon_c, \epsilon_h \to 0$, $u(c,s)$ forms a Cobb-Douglas specification as used in

Sommer and Sullivan (2018):

$$U(c,s) = \frac{(c^{\alpha} s^{1-\alpha})^{1-\sigma}}{1-\sigma}$$

When $\epsilon_h = \epsilon_c$, the utility function reduces to a constant elasticity of substitution (CES) specification. Motivated by our empirical findings in Section 2.3, we observe that as income increases, housing expenditures as a share of income decrease. This motivates us to consider specifications of utility where $\epsilon_h \neq \epsilon_c$, where as income increases, housing expenditures increase at a slower rate relative to non-housing consumption. Housing is therefore more of a necessity good relative to non-housing consumption.

Households can choose to rent or own a house to enjoy shelter services, which are differentiated by size. Households can either choose to own a house of size $h \in \{h_1, ..., h_K\}$, or choose to rent a small apartment of size $\underline{h} < h_1$. Each unit of housing provides an equivalent unit of shelter services s. This implies that the apartments available for renting is smaller than the ones available for purchase as found by Halket, Nesheim, and Oswald (2020). We follow Kaplan, Mitman, and Violante (2020) and assume that markets for rental and owner-occupied housing are both frictionless and competitive, meaning that buying and selling does not take time and the law of one price holds.

Costs associated with owner-occupied housing are determined as follows.When buying a house there is a per unit price of housing denoted by q. We suppose that the household gets an interest only loan to finance a fraction $(1-\theta)$ of the house purchase with a mortgage rate with a wedge of ι over the risk free rate. The household must put a down payment of $\theta q h'$, which creates an incentive to save for households that want to either become homeowners or up-size. Owner-occupied houses carry a per-period maintenance and tax cost of $\delta q h'$ and $t q h'$. Maintenance fully offsets physical depreciation of the dwelling.

When a household sells its home, it incurs a transaction cost $\tau^s q h$ that is linear in the house value. Likewise, when a household buys a home, it incurs a transaction cost of $\tau^b qh'$. As such, the transaction costs can be described as follows:

$$\Gamma(h, h') = \mathbb{1}\{h \neq h'\}(\tau^b qh' + \tau^s qh)$$

Households that choose to rent rather than own a home simply pay a per-unit price of renting equal to ρ.

Summarizing, the household description we have that households enter each period with an amount of housing $h \in \{0, h_1, ..., h_K\}$, assets a, and observe their idiosyncratic wage shock w. Households also observe the extreme-value type-1 preference shock $\epsilon(h')$ for housing types in the next period, which is unobserved by the econometrician. Given prices q, ρ, they solve the following dynamic problem:

$$V(a, h, w) = \max_{h', a'} U(c, s) + \epsilon(h') + \beta E[V(a', h', w')]$$

subject to:

$$c + a' + \theta q(h' - h) + \Gamma(h, h') + \rho(s - h') \leq w + (1 + r)a - (t + \delta + (1 - \theta)(r + \iota))qh' \tag{2.4}$$

$$a' > 0 \tag{2.5}$$

$$s = \begin{cases} h' & \text{if } h' > 0 \\ \underline{h} & \text{if } h' = 0 \end{cases} \tag{2.6}$$

We assume that housing h' in each period provides an equivalent unit of housing services $s = h'$ if the household owns a positive amount of housing, i.e. are homeowners. If households do not own any housing, they pay a rental cost ρ per unit of shelter service that they rent. Households also pay a transaction cost Γ to adjust their housing stock, a maintenance cost δ, and pay property taxes of t on the current value of their choice of property h'.

The household faces the following Euler equation for consumption:

$$U'(c,s) = \beta(1+r)E_{w,\epsilon}\left[\frac{\partial V(a',h',w')}{\partial a'}\right]$$

At first glance, this Euler equation seems standard as it states that the marginal loss of savings (reduced current consumption), must be equal to the marginal gain of savings (higher future consumption.) One important thing to notice, however, is how the expectation is not only over income but also the future discrete choice of housing which the household does not fully now as it cannot anticipate the realization of future taste shocks. The household will use the conditional choice probability to evaluate the expectation of the value function over future housing choices. The conditional choice probability can be solved in closed form as:

$$\mathbb{P}(h'|a,h,w) = \frac{\exp\{E_w[V(a',h',w'|h''=j)]/\sigma_\epsilon\}}{\sum_k \exp\{E_w[V(a',h',w'|h''=k)]/\sigma_\epsilon\}}$$

It is also possible to solve for the expectation of the value function over future housing choices as

$$E_{w,\epsilon}[V(a,h,w)] = \sigma_\epsilon \log(\sum_j \exp\{\frac{E_w[V(a',h',w'|h''=j)]}{\sigma_\epsilon}\})$$

Hence, the household takes into account the future uncertainty related to their housing choices. Regarding the optimal housing choice, the household simply chooses the housing option that yields the highest utility after observing their taste shock. The higher the levels of savings, the more likely the household is to enter homeownership or up-size. This, in turn, provides an incentive for savings in our model in addition to the standard consumption smoothing motive.

2.4.3 Housing Supply

2.4.4 Construction Sector

We suppose that there is a competitive construction sector in the economy using a constant returns to scale production technology. They construct housing using a linear production function employing only materials that can be bought at a unit price of ϕ such that the profit maximization problem of the construction firm reads

$$\bar{\Pi} = \max_{M}\{M - \phi M\}$$

where we let M denote the amount of material used for the production of housing. The convenience of this assumption is that we study the effect of property prices in an environment where house prices are fixed.

Rental Sector

There is a competitive rental sector in the economy that can buy housing at the unit price of q, and rent it out to households. The rental rate of a unit of housing is denoted by ρ. The rental sector has deep pockets so they do not rely on mortgage financing. The rental companies are responsible for maintaining the homes, just like home owners, so they face a maintenance cost of δ times the value of the house each period. The rental company further pays a property value tax of the same size as home owners. Finally, we suppose that there are some operating costs associated with renting which we will denote by ϕ. Due to no arbitrage the return to renting out an apartment must be the same as the return that could be obtained by investing in the risk-free asset. Hence, the rental price must satisfy the following user-cost formula:

$$\rho = (\psi + \delta + t + r)q$$

Notice, we consider a stationary environment where house prices are constant so there is no appreciation of housing prices over time. The user-cost formula can be extended to such an environment as done by Kaplan, Mitman, and Violante (2020).

2.4.5 The Government

The model features a government that levies a value tax on housing. The proceeds from this tax are used to finance wasteful government spending in a way that does not influence household welfare. We make this simplifying assumption as we are interested in studying only the distributional consequences of a property value tax.

The government must satisfy the following budget constraint:

$$\int tqh'(x)d\lambda(x) = G$$

where λ is the distribution over the states that we denote by x.

2.4.6 Stationary Equilibrium

An equilibrium is defined by decision rules $c(x)$, $a'(x)$, and $h'(x)$, prices q^* and ρ^*, and a stationary distribution λ such that:

1. $c(x)$, $a'(x)$, and $h'(x)$ are optimal decision rules for the household's problem, given prices q^* and ρ^*.
2. The rental price follows the following user-cost formula

$$\rho = (\psi + \delta + t + r)q$$

3. The measure λ is stationary, i.e.,

$$\lambda(B) = \int_{\mathcal{S}} P(x, B)\, d\lambda \quad \text{for any Borel set } B.$$

When solving for the stationary equilibrium, we exploit that the expectation of the value function can be solved in closed form following Iskhakov et al. (2017). We then solve the model by use of value function iteration.

2.5 Calibration and Solution Method

In this section, we describe the calibration of the model. WE fix some parameters externally, while others are calibrated to improve the fit of the model.

Externally Fixed Parameters. The model period is one year. We set $\beta = 0.985$. For the utility function, we set the consumption share of housing to 0.685. ϵ_c and ϵ_h are going to be important parameters as they determine the degree to which housing is a necessity good. We choose to follow Cioffi (2021), and we set ϵ_c to 0.91, and ϵ_h to 0.75.

For the income process, we choose the expected life time L to be 50 following Sommer and Sullivan (2018). We further follow the tradition in the quantitative macroeconomics literature, and assume that the labor income process is governed by an AR(1) process

$$\ln(y_t) = \rho_y \ln y_{t-1} + \epsilon_t, \quad \epsilon_t \sim \text{LogNormal}(0, \epsilon_y^2)$$

Following Sommer and Sullivan (2018), the persistent labor productivity process is parameterized by $\rho = 0.9$.

For housing costs, we follow Kaplan, Mitman, and Violante (2020), and we set the property value tax to 1 percent, and the housing maintenance cost to 1.5 percent.

We allow for asymmetric transactions costs when selling or buying housing, and we set $\tau^s = 0.07$ following Sommer and Sullivan (2018) while we calibrate the transaction cost of buying.

We report all of the externally fixed parameters in Table 2.3 as well as their source.

Parameter	Value	Description	Source
β	0.985	Discount factor	Sommer and Sullivan (2018)
α	0.685	Consumption share	Sommer and Sullivan (2018)
ϵ_c	0.91	elasticity of substitution for consumption	Cioffi (2021)
ϵ_h	0.75	elasticity of substitution for housing	Cioffi (2021)
L	50	Expected life time	Sommer and Sullivan (2018)
ρ_y	0.9	Auto correlation of wage shocks	Sommer and Sullivan (2018)
t	0.01	Property value tax	Kaplan, Mitman, and Violante (2020)
δ	0.015	Housing maintenance	Kaplan, Mitman, and Violante (2020)
ι	0.33	mortgage rate wedge	Kaplan, Mitman, and Violante (2020)
τ^b	0.07	proportional transaction cost of selling	Sommer and Sullivan (2018)

Table 2.3 – Exogenous parameters

After having fixed these parameters, we are left with the following parameters:

$$R, \tau^b, \sigma, \phi, \psi, \theta, \sigma_\epsilon, \text{ and } \sigma_w$$

that we calibrate in order to match the data as well as possible. As our primary concern is the impact that property taxes have on inequality, we want to match well the home ownership rate over the wealth distribution, the portfolio share of housing over the wealth distribution, and the expenditure share of housing over the income distribution. Hence, we choose to calibrate the vector of parameters

$$\Theta = \{R, \tau^b, \sigma, \phi, \psi, \theta, \sigma_\epsilon, \sigma_w\}$$

to minimize the distance between the moments computed from the data and those generated by the model by minimizing the following loss function.

$$\text{loss} = \sum_{i=1}^{3}\big((\text{data}_{\text{hti,i}} - \text{model}_{\text{hti,i}})^2\big) + \sum_{i=1}^{10}\big((\text{data}_{\text{ho,i}} - \text{model}_{\text{ho,i}})^2\big) + \sum_{i=1}^{10}\big((\text{data}_{\text{htw,i}} - \text{model}_{\text{htw,i}})^2\big)$$

where the first term is the squared different in housing expenditures for each income tercile, the second term is the squared different in homeownership rate for each wealth decile, and the third term is the squared difference in the portfolio share of of housing for each wealth decile.

We minimize the loss function by defining a grid over the parameter values, and then we search over this grid to find the combination that minimizes the loss function.

We report the calibrated parameters in Table 2.4 below:

R	1.025	Gross interest rate
τ^s	0.07	proportional transaction cost of buying
σ	0.5	Risk aversion
ϕ	1.7	Marginal cost of construction
ψ	0.02	Operating cost of renting
θ	0.3	Down payment requirement
σ_ϵ	0.025	Scale parameter of taste shock
σ_w	0.22	Std. dev. of wage shocks

Table 2.4 – Calibrated parameters

Overall the calibrated parameters appear reasonable. We calibrate the return on the safe asset to be at 2.5 %, which is not far from the average US fed funds rate over the past 10 years. We get a transaction cost of buying of 7 percent, which may be somewhat high but not unreasonable. The calibrated sigma of 0.5 is a bit far from the one used by Cioffi (2021) as he uses two. The operating cost of renting is a bit higher than the 0.008 used by Kaplan, Mitman, and Violante (2020) but it does not appear unreasonably large. The down payment requirement of 0.3 is close to the original Loan-to-Value according to FRED [7]. The scale parameter of the taste shock is small enough that it's not only randomness that determines housing choices, and finally, the standard deviation of wage shocks of 0.22 is close to the 0.2 frequently used in the literature.

7. https://fred.stlouisfed.org/series/RCMFLOLTVPCT50

To see how well the model matches the data based on the calibration, we show the home ownership rate and portfolio share of housing over the wealth distribution in the model and the data as well as the expenditure share of housing in both the model and the data in the figure below.

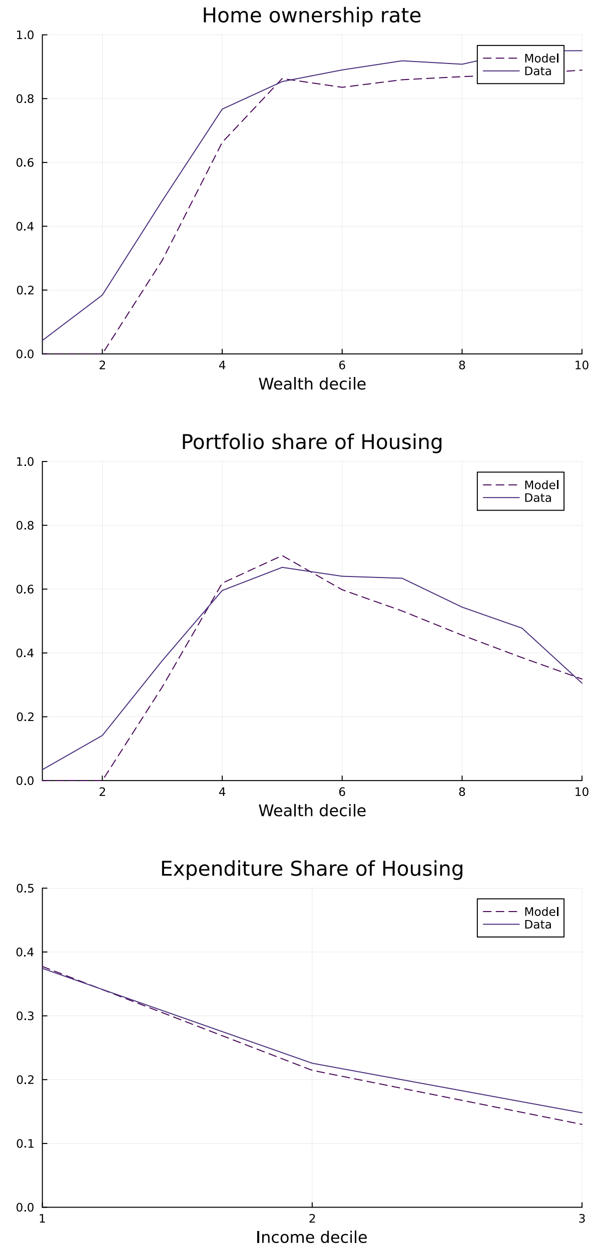

Figure 2.6 – Model vs Data

All in all, the model captures decently the stylized facts that we laid out as motivating evidence. We are now ready to consider the distributional consequences of counterfactual levels of property taxes.

2.6 Results

In this section, we show the consequences of increasing the property value tax from the initial 1 percent to either 2 or 3 percent, with the purpose of studying the distributional consequences and equilibrium impact. To show the importance of using a model with non-homothetic preferences, we compare our results to a model where we use homothetic preferences instead. In our model, we have assumed that house prices are fixed, and that changes in taxes are fully passed through to renters. We also check the sensitivity of our results to this assumption by comparing our model to one where tax changes are not passed through to renters.

2.6.1 The effects of increasing property taxes

We report the main economic consequences in Table 2.5. We report the home ownership rate, core measures of inequality, taxes, and welfare, as measured by compensating variation. The measures of inequality that we report are the Gini coefficients for both net wealth and consumption, as well as the inter-quartile range and the ratio of wealth between the third and first quartile. We report taxes both as a share of income, and its value as a discounted stream of payments, holding housing choices constant.

We measure welfare using (the negative) of compensating variation. For an individual in the counterfactual with state a, j, h, we find the value function in the baseline specification with the same productivity and housing state. This provides us with a level of assets a', where the difference with respect to a gives us our compensating variation. If a value function in the baseline specification cannot be found, the value function is interpolated such that a level of assets a' can be found. We then compare this to expected lifetime income inclusive of current assets.

	Tax 1%	Tax 2%	Tax 3%
Home ownership rate	0.6075	0.5554	0.4985
Gini (Net W)	0.5361	0.5484	0.5652
Gini (Consumption)	0.2977	0.2969	0.2970
Inter-quartile Range	5.6076	5.5004	5.5004
Q3/Q1 Ratio	9.6450	12.3063	17.9595
Tax (Share of Income)	0.0236	0.0391	0.0501
Tax (Discounted Stream)	1.9576	3.2450	4.1614
Compensating Variation	0.0000	2.9354	5.6815
Compensating Variation / Expected Future Income	0.0000	0.0722	0.1407

Table 2.5 – Effects of higher Property Taxes

We first notice that an increase in property taxes is associated with a lower home ownership rate in equilibrium. In our model, there are three forces at play. Higher property taxes increase the cost of owner-occupied housing, which will tend to depress the home ownership rate. On the other hand, higher property taxes also leads to higher rents in equilibrium due to the user-cost formula we are applying which makes renting less attractive. A last effect is that higher property taxes make housing consumption more expensive compared to consumption of the non-durable good. The net effect is for property taxes to reduce the home ownership rate. The intuition can be understood from Table 2.6 which shows the fraction of households choosing each housing size at different levels of the property value tax. We find that individuals substitute towards smaller houses and renting as taxes increase.

Parameter	$t = 0.01$	$t = 0.02$	$t = 0.03$
h = 0.0	0.393	0.442	0.501
h = 2.2	0.275	0.380	0.401
h = 3.3	0.298	0.155	0.068
h = 4.0	0.035	0.023	0.029

Table 2.6 – Housing Distribution

As the smallest unit is only available for rent, this will tend to depress the home ownership rate.

From Figure 2.7 below we notice that it is especially households in the middle of the wealth distribution that switch to renting rather than owning in response to higher property taxes.

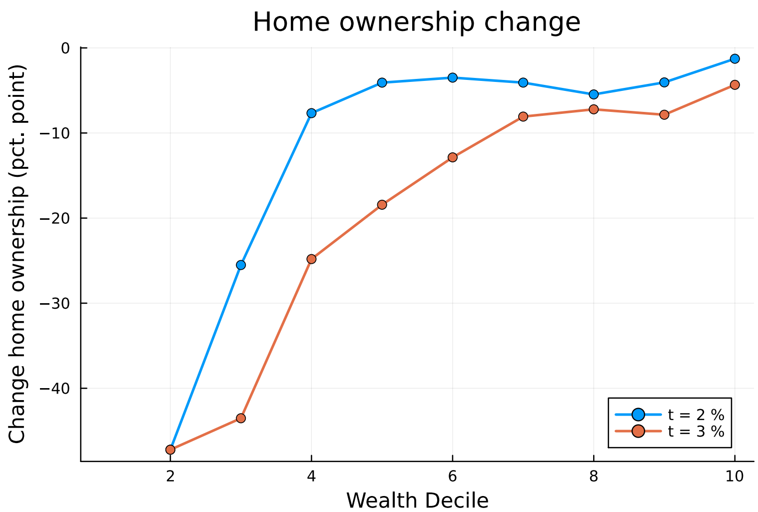

Figure 2.7 – Home ownership change

Interestingly, we find that the effect of property taxes on inequality is ambiguous. It depends on what variable to consider. We find that higher property taxes lead to an increase in wealth inequality as measures by the Gini coefficient. On the other hand, we find that the increase in property taxes *decrease* consumption inequality, though not by much, when taxes increase from 1 % to either 2 % or 3 %. To better understand why we get these two opposite effect, we show the change in non-durable consumption and net worth over the wealth distribution in Figure 2.8 below.

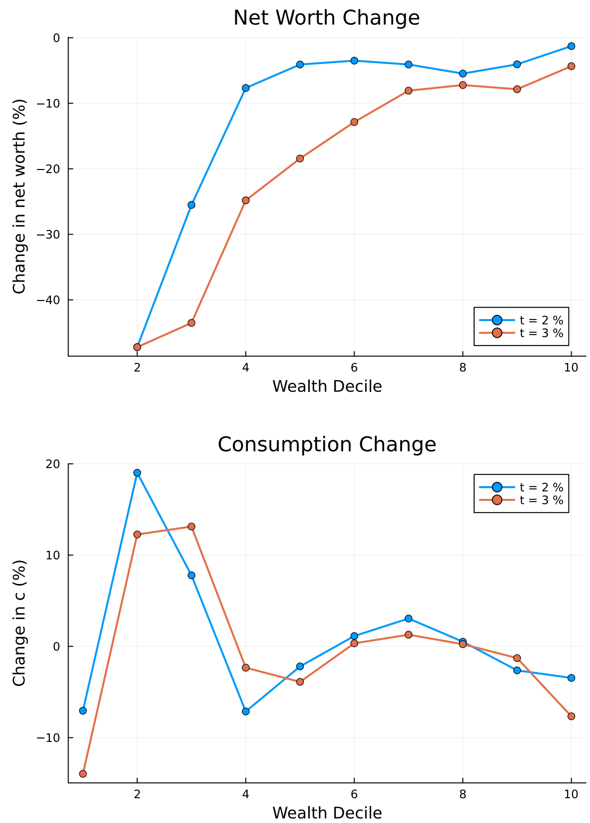

Figure 2.8 – Effects of Higher Property taxes

We see that net worth drops most for lower wealth deciles, while consumption actually increases for households in decile two and three. This is due to the fact that as housing has now become more expensive, the incentive to save for a down payment is reduced as fewer households will later choose to become home owners. Instead of entering the market for owner-occupied housing, poor households choose to increase their consumption level. From the figure below, we see that it is especially households in the middle class that choose to stay renters compared to the baseline scenario. But as households in the lower wealth decile are likely to be income poor households, they are on an upward income trajectory. As they realize that as they climb the wealth distribution, they are less likely to increase their housing consumption by becoming home owners, they choose to decrease their savings.

Following the increase in property taxes, we see that the average tax paid by the household increases, however, the change is smaller than compensating variation, indicating that individuals prefer to be in a world where taxes are lower, which may reflect the fact that individuals prefer choices to be undistorted.

In Figure 2.9, we take a closer look at the distributional consequences of property taxes. In particular, we report the average welfare effect (negative of compensating variation) as a share of future expected income, for each wealth decile in figure 2.9. We find that welfare decreases are larger for wealthier individuals as a share of future expected income, which could reflect three features of our model. First, wealthier individuals are more likely to be homeowners and are thus more likely to be impacted by an increase in property taxes. Second, wealthier individuals have a lower marginal utility of wealth, so they require a greater sum to compensate them for welfare losses. Third, wealthier individuals in our model may also be more likely to be of higher income, but higher income individuals may have shorter expected lifespans in a stochastic aging model.

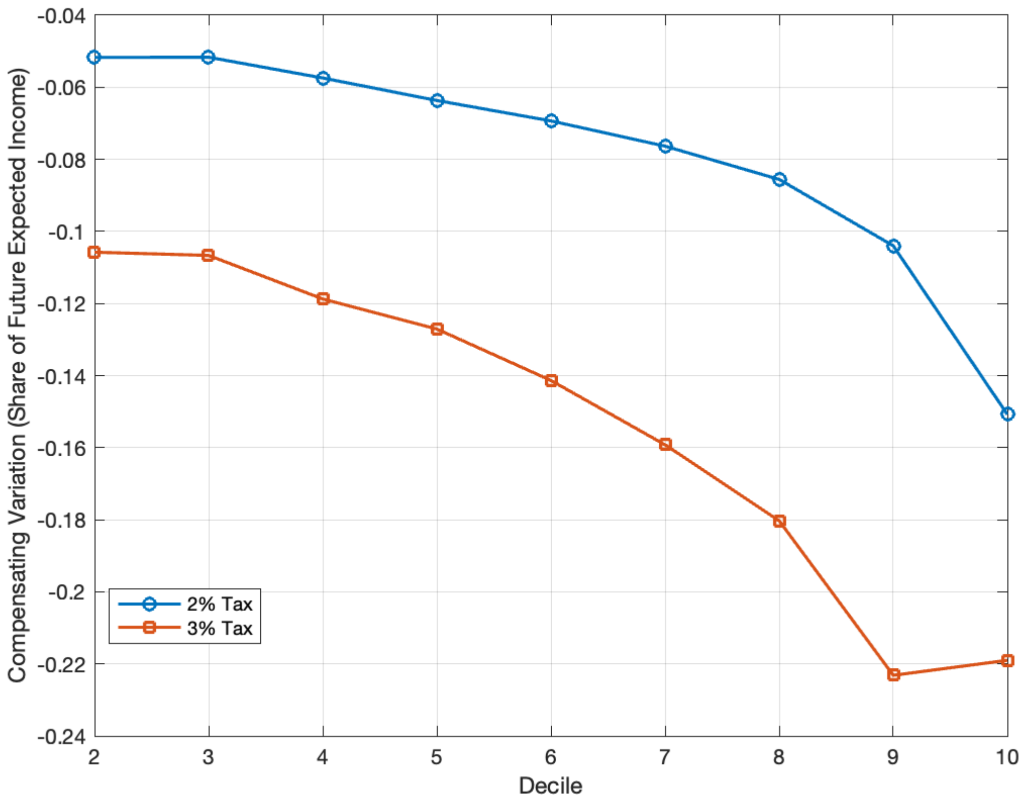

Figure 2.9 – Welfare Effects

Summarizing our results, we find that higher property taxes reduce the home ownership rate and increases wealth inequality, despite the fact that consumption inequality actually decreases, though it is rather constant across the different tax policies that we consider.

We now go on to test the sensitivity to our results to two key assumptions: non-homothetic preferences and fully competitive rental markets implying a user cost formula for the rental price.

2.6.2 Homothetic Preferences

We consider a model with homothetic preferences, which in our model, would be one where ϵ_h, ϵ_c are equal. To make our comparison as close as possible to our baseline specification, we re-calibrate $\epsilon_c, \epsilon_h,$ and α such that the home ownership rate with taxes at 1% is the same as in our model with non-homothetic preferences, under the restriction that $\epsilon_h = \epsilon_c$. Using simulated method of moments targetting only the aggregate home-ownership rate, we find that $\epsilon_h, \epsilon_c = 0.7285$, and $\alpha = 0.7695$ allows us to match the home ownership rate in the baseline model exactly. We show the results in Table 2.7 below.

	Tax 1%	Tax 2%	Tax 3%
Home ownership rate	0.6070	0.5213	0.4002
Gini (Net W)	0.5398	0.5353	0.5317
Gini (Consumption)	0.3131	0.3224	0.3303
Inter-quartile Range	7.1220	7.1220	7.4463
Q3/Q1 Ratio	15.6397	22.9595	23.9595
Tax (Share of Income)	0.0193	0.0320	0.0361
Tax (Discounted Stream)	1.6014	2.6603	2.9954
Compensating Variation	0.0000	2.3082	4.6258
Compensating Variation / Expected Future Income	0.0000	0.0562	0.1127

Table 2.7 – Results when preferences are homothetic

Compared to our results in table 2.5, we find that the home ownership rate has a greater decline as taxes increase, which we expect, as in a model with homothetic preferences, lower wealth and income individuals substitute more strongly to renting as taxes increase compared to a model with non-homothetic preferences where housing is a necessity good.

We find that the compensating variation is lower with homothetic preferences. Under non-homothetic preferences, for the same tax increase, the home ownership rate falls less and individuals have stronger preferences for larger homes, implying that the property tax payments made by households are higher.

2.6.3 Full tax incidence on landlords

One important assumption of our paper so far has been that the price of housing is fixed while the rent level is changing as we change the level of property taxes. To test the sensitivity of our results to this assumption, we consider a version where rents are fixed to the level they have when property taxes are 1 percent. We therefore suppose that all the incidence of higher property taxes fall on home owners and landlords. We report the results in Table 2.8.

	Tax 1%	Tax 2%	Tax 3%
Home ownership rate	0.6075	0.5106	0.4209
Gini (Net W)	0.5361	0.5686	0.5869
Gini (Consumption)	0.2977	0.2807	0.2684
Inter-quartile Range	5.6076	5.5004	5.1761
Q3/Q1 Ratio	9.6450	17.9595	16.9595
Tax (Share of Income)	0.0236	0.0364	0.0435
Tax (Discounted Stream)	1.9576	3.0242	3.6113
Compensating Variation	0.0000	2.8036	5.4535
Compensating Variation / Expected Future Income	0.0000	0.0692	0.1358

Table 2.8 – Results with full tax incidence on landlords

Compared to the model where rents fully to property taxes, we find that higher property taxes decreases the home ownership rate much more significantly. As explained before this is due to the fact that higher property taxes now only make owning less favorable while they do not change the utility of renting. We see that the decrease in consumption inequality is much stronger which is due to the fact that renters, who tend to be poor agents, are not directly affected by the policy. However, we find that compensating variation is slightly lower than with full tax pass through, which reflects the fact that individuals need to be compensated less if the rental option is unchanged. It is therefore of first order important to assess the degree to which the tax incidence falls on renters or landlords as this drastically impacts the equilibrium impact of property taxes.

2.7 Conclusion

This paper investigates the distributional effects of property value taxation through the lens of a heterogeneous-agent model with non-homothetic preferences over housing and non-durable consumption. Our results suggest that property taxes, though often considered progressive and efficient, may exacerbate wealth inequality. Middle-wealth households, who concentrate a large share of their wealth in housing, bear a disproportionate burden of the tax. Moreover, increased property taxation leads to higher rents, further disadvantaging renters and lower-wealth households.

References

Aiyagari, S. (1994). « Uninsured Idiosyncratic Risk and Aggregate Saving ». In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.

Atkinson, A.B. and J.E. Stiglitz (1976). « The design of tax structure: Direct versus indirect taxation ». In: *Journal of Public Economics* 6.1, pp. 55–75.

Bacher, Annika (2023). « Housing and Savings Behavior Across Family Types ». In: *Job Market Paper*.

Bontemps, Christian, Frédéric Cherbonnier, and Thierry Magnac (2023). « Reducing transaction taxes on housing in highly regulated economies ». Manuscript. TSE Working Paper.

Chamley, Christophe (1986). « Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives ». In: *Econometrica* 54.3, pp. 607–22.

Christiansen, Vidar (1984). « Which commodity taxes should supplement the income tax? » In: *Journal of Public Economics* 24.2, pp. 195–220.

Cioffi, Riccardo A. (2021). « Positive Long-Run Capital Taxation: Chamley-Judd Revisited ». In: *Job Market Paper*.

Dettling, Lisa J. et al. (2015). *Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances*. FEDS Working Paper 2015-086. Washington, D.C.: Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh (2017). « The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium ». In: *Journal of Political Economy* 125.1, pp. 140–223.

Fischer, Marcel and Natalia Khorunzhina (2019). « HOUSING DECISION WITH DIVORCE RISK ». In: *International Economic Review* 60.3, pp. 1263–1290.

Fisher, Jonas D. M. and Martin Gervais (2011). « WHY HAS HOME OWNERSHIP FALLEN AMONG THE YOUNG?* ». In: *International Economic Review* 52.3, pp. 883–912.

Flavin, Marjorie and Takashi Yamashita (2002). « Owner-Occupied Housing and the Composition of the Household Portfolio ». In: *American Economic Review* 92.1, pp. 345–362.

Halket, Jonathan, Lars Nesheim, and Florian Oswald (2020). « The Housing Stock, Housing Prices, And User Costs: The Roles Of Location, Structure, And Unobserved Quality ». In: *International Economic Review* 61.4, pp. 1777–1814.

Halket, Jonathan and Santhanagopalan Vasudev (2014). « Saving Up or Settling Down: Home Ownership over the Life Cycle ». In: *Review of Economic Dynamics* 17.2, pp. 345–366.

Heathcote, Jonathan (2005). « Fiscal Policy with Heterogeneous Agents and Incomplete Markets ». In: *The Review of Economic Studies* 72.1, pp. 161–188.

Iacoviello, Matteo and Marina Pavan (2013). « Housing and debt over the life cycle and over the business cycle ». In: *Journal of Monetary Economics* 60.2, pp. 221–238. ISSN: 0304-3932.

Iskhakov, Fedor et al. (2017). « The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks ». In: *Quantitative Economics* 8.2, pp. 317–365.

Judd, Kenneth L. (1985). « Redistributive taxation in a simple perfect foresight model ». In: *Journal of Public Economics* 28.1, pp. 59–83. ISSN: 0047-2727.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante (2020). « The Housing Boom and Bust: Model Meets Evidence ». In: *Journal of Political Economy* 128.9, pp. 3285–3345.

Norregaard, John (2013). *Taxing Immovable Property: Revenue Potential and Implementation Challenges*. IMF Working Paper 13/129. Prepared by the Fiscal Affairs Department; authorized by Victoria Perry and Michael Keen. Washington, DC: International Monetary Fund.

Piazzesi, M. and Martin Schneider (2016). « Housing and Macroeconomics ». In: vol. 2. Elsevier. Chap. Chapter 19, pp. 1547–1640.

Sommer, Kamila and Paul Sullivan (2018). « Implications of US Tax Policy for House Prices, Rents, and Homeownership ». In: *American Economic Review* 108.2, pp. 241–74.

Sommer, Kamila, Paul Sullivan, and Randal Verbrugge (2013). « The equilibrium effect of fundamentals on house prices and rents ». In: *Journal of Monetary Economics* 60.7, pp. 854–870.

Straub, Ludwig and Iván Werning (2020). « Positive Long-Run Capital Taxation: Chamley-Judd Revisited ». In: *American Economic Review* 110.1, pp. 86–119.

Appendices

2.A Robustness of the results to the chosen sample

We report the results for the sample used in this paper as well as for the full SCF sample.

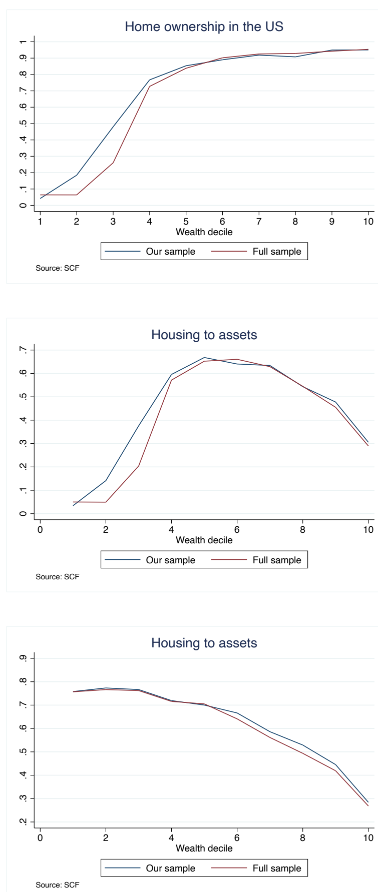

Figure 2.10 – Results with our sample vs full sample

In general, we find that the results are quite similar across the two samples.

2.B Data definitions

In the figure below, we report the flowchart of net worth provided by the Survey of Consumer Finances.

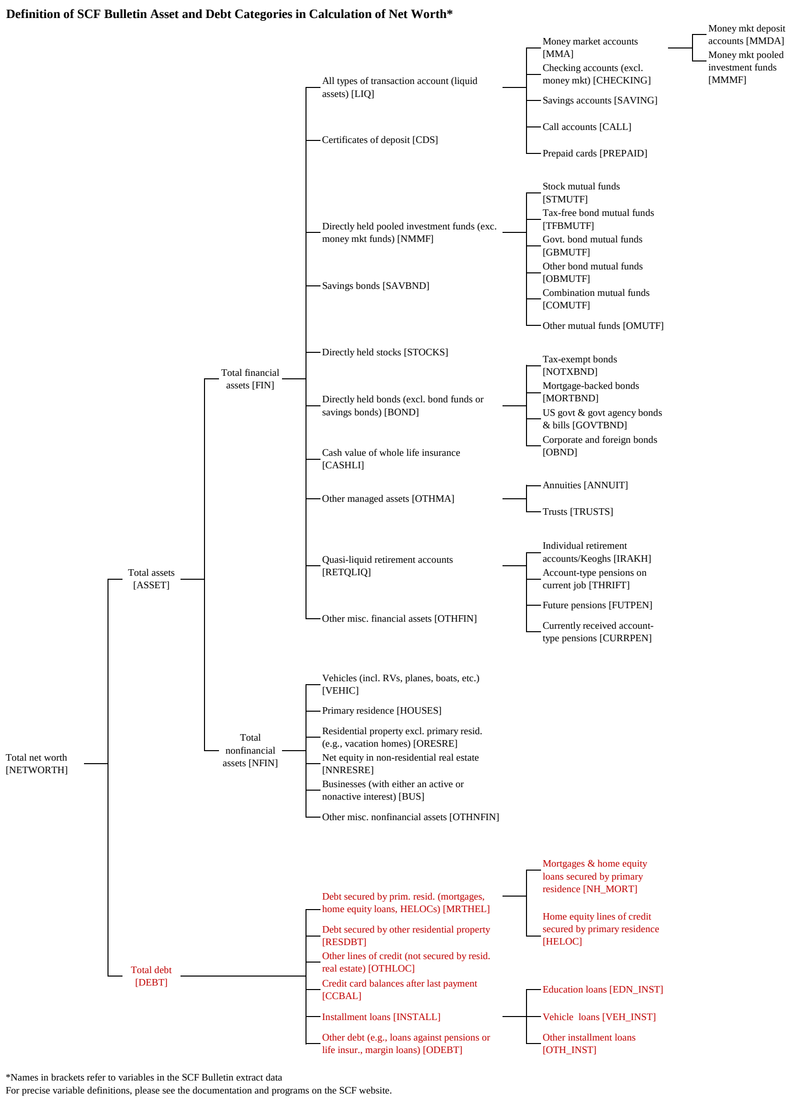

2.C Regression tables

In this section, we present the full results of the regressions in the main text

Table 2.9 – Regression of home ownership on wealth

	homeowner
wealth decile=2	0.150***
	(0.016)
wealth decile=3	0.447***
	(0.020)
wealth decile=4	0.722***
	(0.016)
wealth decile=5	0.808***
	(0.015)
wealth decile=6	0.838***
	(0.015)
wealth decile=7	0.858***
	(0.014)
wealth decile=8	0.847***
	(0.016)
wealth decile=9	0.887***
	(0.016)
wealth decile=10	0.885***
	(0.017)
10 quantiles of income	-0.004*
	(0.002)
age	0.010***
	(0.003)
age_sq	-0.000*
	(0.000)
married	0.067***
	(0.009)
kids	0.026***
	(0.003)
educ	-0.010***
	(0.002)
race	-0.020***
	(0.004)
Constant	-0.193**
	(0.071)
Sample size	14,784
R^2	0.51

Table 2.10 – Regression of portfolio share on wealth

	hta_nwdecile_ho
10 quantiles of networth	-0.047***
	(0.000)
10 quantiles of income	-0.004***
	(0.000)
age	0.003***
	(0.001)
age_sq	-0.000***
	(0.000)
married	0.009***
	(0.002)
kids	-0.003***
	(0.001)
educ	-0.000
	(0.000)
race	0.000
	(0.001)
Constant	0.849***
	(0.015)
Sample size	10,127
R^2	0.87

Chapter 3

The Child Penalty and Statistical Discrimination Against Women

Emil Mortensen [1]

Abstract

This paper examines how statistical discrimination arises when firms cannot observe whether workers will assume future caregiving responsibilities. I develop a theoretical model in which firms, acting under asymmetric information, discriminate against women in hiring for high-investment, inflexible roles, anticipating a higher probability of labor force withdrawal. The resulting equilibrium is inefficient relative to a first-best world with observable caregiving status: caregiving women impose a negative externality on non-caregiving women, who are unjustly excluded from high-productivity positions. Supporting the model's key assumptions, I provide empirical evidence from the German Socio-Economic Panel (SOEP) showing that the birth of the first child leads to sharp declines in full-time employment for women, but not for men, and increases the probability of job transitions for women, with no effect observed for men.

Keywords: Asymmetric information, gender inequality, discrimination
JEL classification: D80, J70

1. Toulouse School of Economics. Email: emitm@outlook.dk. This chapter is a revised version of my memoire that was written under the supervision of Christian Hellwig. I have received useful suggestions and comments from Fabrice Collard, Eugenia Gonzalez-Aguado, Christian Hellwig, Javier Gonzalez-Morin, En Qi Teo, and Nicolas Werquin.

3.1 Introduction

A large body of empirical research has documented the persistence of gender wage gaps across labor markets, and much of this gap can be attributed to the child penalty (see, for instance, Kleven, Landais, and Søgaard (2019)). In Germany, I find that the birth of the first child reduces a woman's gross income by 40 percent three years after the birth. This paper investigates, theoretically, whether gender inequality may emerge even earlier. Specifically, I ask the following question: does the anticipation of future caregiving responsibilities after the birth of a child contribute to gender inequality in labor market outcomes prior to parenthood? Understanding whether expectations about future child-related labor market decisions influence pre-parenthood labor market outcomes is crucial for identifying the full extent of gender disparities and for designing effective policy interventions.

I first use data from the German Socio-Economic Panel (SOEP) to study the effect that the birth of a child has on women's labor market outcomes. I apply an event-study design around the birth of the first child and document a sharp decline in earnings, full-time employment, and job stability for women, but not for men. I take this as suggestive evidence that women are more likely than men to face caregiving responsibilities causing them to either work part-time or look for a job that offers more flexibility.

Motivated by the fact that women are more likely to change jobs or to start working part-time after the birth of a child, I then develop a theoretical model where women face a higher probability of assuming caregiving responsibilities. I assume that workers with caregiving responsibilities value not only the wage they receive but also the amount of flexibility in the job. I suppose that workers have private information about future caregiving roles, i.e., firms cannot distinguish workers with future caregiving responsibilities from workers without them. I take the view that there are no productivity differences between women and men, so all of the results are driven by the assumption that women face a higher probability of assuming caregiving responsibilities. On the firm side, I consider an environment where there are two types of firms. Workers can be employed either by a low-productivity firm that is able to offer flexibility or by a high-productivity firm that cannot offer this, and also requires an upfront investment in the worker to provide the worker with the necessary skills to carry out the job. The fact that the high-productivity firm needs to pay an upfront investment implies that it is reluctant to employ workers that it believes will leave the firm upon the birth of a child. Combined with private information, this is the central mechanism causing gender inequality in the model.

I then show that, under profit-maximizing behavior and wage parity constraints, firms respond by hiring fewer women into high-investment roles under asymmetric information compared to the first-best. This form of statistical discrimination arises not from bias, but from rational responses to uncertainty. Intuitively, caregiving women impose a negative externality on non-caregiving women due to asymmetric information that causes firms to be reluctant to hire women for high-productivity jobs that require an initial upfront investment by the firm. This leads to inefficiencies, and reduced labor market earnings for the women. On the other hand, men with future caregiving responsibilities benefit as they are able to be hired by the high-productivity firms which is inefficient as those roles should ideally be given to women without future caregiving responsibilities. In that sense, the results shown in this model are fully consistent with an inefficient *glass ceiling* that prevents women from reaching top positions they would be better qualified for than the men who end up getting them.

The paper is organized as follows. Section 3.2 relates this paper to the existing literature. Section 3.3 introduces the German Socio-Economic Panel. Section 3.4 provides an event-study around the birth of the first child that is used for motivating the model. Section 3.5 introduces the model. Section 3.6 solves the equilibrium of the model. Section 3.7 features a discussion of the results. Section 3.8 concludes.

3.2 Related Literature

This paper contributes to the literature on *gender inequality*, the *child penalty*, and *asymmetric information*, by analyzing how imperfect private information about caregiving responsibilities leads to inefficient labor market outcomes and gender inequality.

This paper contributes to a vast and growing literature on gender inequality. Several studies emphasize that job flexibility is a key mechanism driving gender disparities. Goldin (2014) argues that labor markets disproportionately reward rigid, inflexible schedules that penalize women. These insights are consistent with the theory of *equalizing differences* formalized by Rosen (1986), where workers accept lower pay in exchange for non-wage benefits like flexibility. Andrés Erosa et al. (2022) show quantitatively how occupational structure and time constraints shape gendered labor outcomes. Doepke and Kindermann (2019) examine how intra-household bargaining over fertility and careers contributes to gender inequality. Albanesi and Olivetti (2009) further model how expectations of future home production contribute to wage gaps under asymmetric information.

Several macro-labor studies provide quantitative frameworks relevant to this model. Andres Erosa, Fuster, and Restuccia (2016) offer a life-cycle model linking labor supply choices to gender wage gaps. Sloane, Hurst, and Black (2021) show how college majors and occupational sorting contribute to early-career gender disparities, reinforcing labor market frictions. Similarly, Olivetti and Petrongolo (2017) and Blau and Kahn (2013) evaluate how family policy design and institutional structure shape female labor supply across countries. Azmat and Ferrer (2017) show that there is a wage gap among lawyers that can be accounted for by the fact that young male lawyers outperform their female counterparts, largely due to differences in aspirations and the presence of preschool-aged children.

The paper also contributes to the literature on the child penalty. A growing empirical literature has documented that the birth of a child results in substantial and persistent negative impacts on women's earnings and labor market attachment, with no comparable effect for men. Kleven, Landais, and Søgaard (2019) show that childbirth leads to sharp, long-lasting earnings declines for women, while men's careers continue largely unaffected using Danish register data. I replicate their study for Germany. The replication of Kleven, Landais, and Søgaard (2019) for Germany has also been done by Kleven, Landais, Posch, et al. (2019). My analysis extends theirs by additionally examining full-time employment and job transitions, and I further propose a theoretical model to think about the effects prior to childbirth. Nix and Andresen (2019) also find a substantial child penalty, and Adda, Dustmann, and Stevens (2017) use a dynamic life-cycle model to show how motherhood alters career paths and long-run earnings.

Finally, the paper is also related to the literature on *statistical discrimination* and *asymmetric information*. The foundational model by Akerlof (1970) describes how markets fail under quality uncertainty. Phelps (1972) and Arrow (1973) adapt these ideas to labor markets, showing how group averages are used when individual productivity is unobservable. In this tradition, the current paper models how gender operates as a noisy signal of future caregiving status, leading to rational but inefficient hiring discrimination. This contrasts with taste-based discrimination models as in Becker (1957), where disparities arise from employer preferences rather than information constraints. The modeling approach in this paper builds closely on Gayle and Golan (2012), who estimate a dynamic adverse-selection model in which statistical discrimination accounts for a large portion of the gender earnings gap. While their model centers on unobserved ability, the current paper shifts the focus to unobserved caregiving intent, offering a complementary channel. Finally, this paper is grounded in the informativeness principle of Holmstrom (1979), which shows that any observable variable correlated with unobserved productivity—such as gender in this case—will enter into optimal contracts. In the absence of observability, firms use gender as a proxy for caregiving risk, generating distortions that penalize even career-focused women. These externalities reduce both efficiency and equity in labor market allocations.

3.3 Data

I use data from the German Socio-Economic Panel (SOEP), a longitudinal survey that began in 1984 and interviews households annually. The SOEP provides rich information on household composition, occupational status, earnings, and other key socio-economic indicators.

A distinctive feature of the SOEP is its effort to follow all members of the originally sampled households over time. If household members move out—such as children forming their own households or households splitting due to divorce—the new households continue to be surveyed alongside the original ones. This design enables comprehensive tracking of individuals and household dynamics across time.

The SOEP comprises several subsamples. The initial sample, Sample A, includes individuals residing in West Germany in 1984, covering 4,528 households. In 1990, Sample C was introduced, adding 2,179 households from the former German Democratic Republic. In addition to these core samples, the SOEP has added various special samples focused on particular subpopulations (e.g., immigrants, refugees) that are underrepresented in general population samples. To maintain the representativeness of the survey over time and to compensate for attrition, several refreshment samples have also been added.

In this paper, I use data from the two original core samples (A and C) as well as all subsequent refreshment samples. Table 3.1 provides an overview of the samples used and the number of initially surveyed households. One household may include multiple individual-level observations.

Sample Letter	Sample Name	Start Year	Households
A	West German Sample	1984	4,528
C	East German Sample	1990	2,179
E	Refresher	1998	1,060
F	Refresher	2000	6,043
H	Refresher	2006	1,506
J	Refresher	2011	3,136
K	Refresher	2012	1,526
N	Refresher Sample (PIAAC-L)	2017	2,314

Table 3.1 – SOEP samples used in this paper

Sample restrictions: I restrict the analysis to individuals aged 17 to 50 and residing in private households. I further restrict attention to individuals that I observe two years prior to the birth of their first child, and then for the following five years, implying that I work with a balanced panel. In Table 3.2 below, I provide summary statistics for the sample used.

	Mean	SD	Min	Max	N
Woman	0.56	0.50	0.00	1.00	5,916
Age	29.32	7.22	17.00	50.00	5,916
Number of Children	1.01	0.63	0.00	3.00	5,916
Income	1,388.26	1,566.48	0.00	16,600.00	5,911
Full time employment	0.48	0.50	0.00	1.00	5,916
New job	0.24	0.43	0.00	1.00	3,999

Table 3.2 – Summary Statistics

The average age of the sample, consisting of individuals observed from two years before to three years after the birth of their first child, is approximately 29 years. The sample includes more women than men. The average gross labor income is €1,388 and the average number of children is slightly above one, indicating that some individuals have a second child shortly after the first one.

3.4 Motivating Evidence

This section documents that women are more likely than men to reduce full-time employment, start working part-time or to switch job following the birth of their first child. In particular, the increased incidence of part-time work among women provides motivating evidence for the model assumption that primary caregivers value job flexibility in addition to wages. I further replicate the findings of Kleven, Landais, and Søgaard (2019) on the child penalty in Germany, showing that women experience a substantial decline in labor income around the time of their first childbirth.

I apply the methodology of Kleven, Landais, and Søgaard (2019) to investigate how the birth of the first child affects the probability of full-time employment, part-time employment, and the likelihood of changing jobs. I further estimate the child penalty separately for men and women, using the birth of the first child as the event of interest.

I estimate the following event-study regression separately by gender $g \in \{\text{men}, \text{women}\}$:

$$Y^g_{ist} = \sum_{j \neq -1} \alpha^g_j \mathbb{1}(t = j) + \sum_a \beta^g_a \mathbb{1}(\text{age}_{is} = a) + \sum_s \gamma^g_s \mathbb{1}(\text{year} = s) + \epsilon^g_{ist}$$

where Y^g_{ist} denotes the outcome for individual i of gender g observed in year s and event time t relative to the birth of the first child (with $t = 0$ at the birth). The outcomes Y are the probability of full-time employment, the probability of part-time employment, the probability of a job-to-job transition, and gross labor income. Age and calendar year fixed effects are included to control for life-cycle and macroeconomic trends. Coefficients are normalized to zero at event time $t = -1$, so the α^g_j coefficients measure changes relative to the year before childbirth.

Figure 3.1 shows a marked decrease in the probability of full-time employment for women following childbirth, with no effect for men. This is relevant for the model that I later propose as firms may be more reluctant to hire women in positions that require an upfront investment in the worker as they face a higher probability of her leaving than a male applicant for the same job.

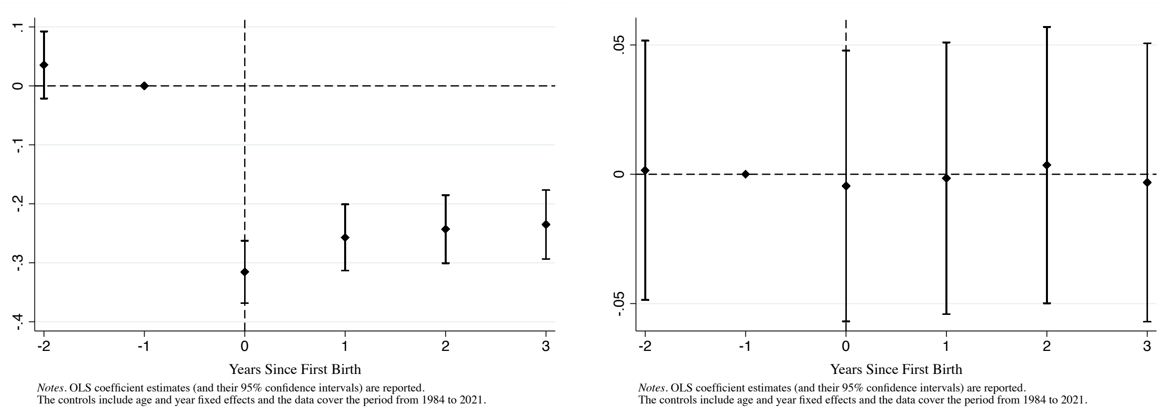

(a) Women

(b) Men

Figure 3.1 – Birth and the probability of full-time employment

In Figure 3.2, we see that some of decrease in full-time employment is driven by women switching to part-time positions. This may indicate that after the birth of the first child, workers value flexibility in the job in addition to the wage they get paid.

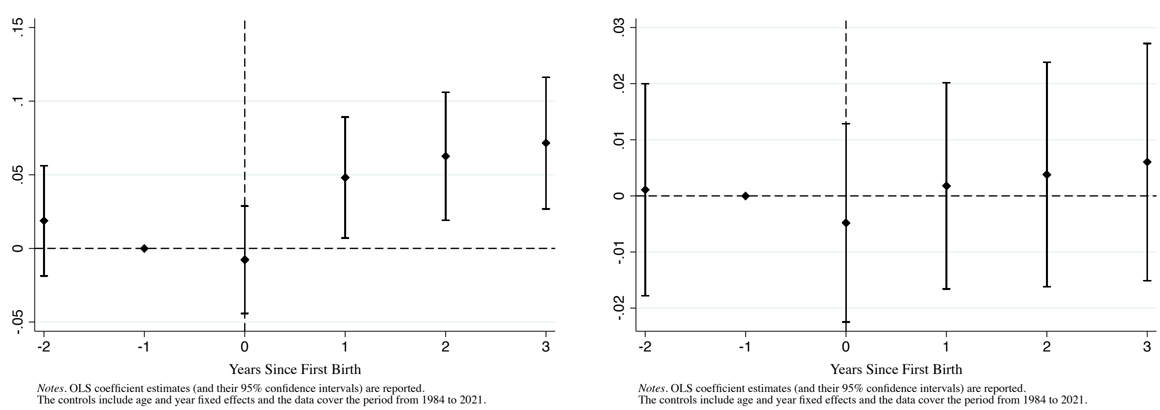

(a) Women

(b) Men

Figure 3.2 – Birth and the probability of part-time employment

Figure 3.3 reports estimates on job-to-job transitions. The birth of the first child significantly increases the likelihood that women change jobs, whereas no significant effect is found for men. Similar to the drop in full-time employment, it indicates that firms face a more elevated risk of a worker leaving post the birth of the first child if it is a female worker.

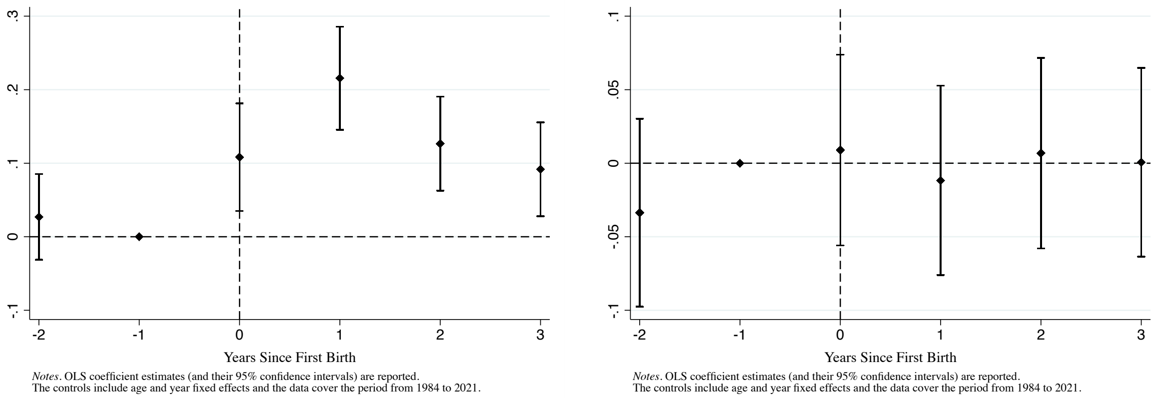

(a) Women

(b) Men

Figure 3.3 – Birth and the probability of job-to-job transitions

Finally, I replicate the results of Kleven, Landais, and Søgaard (2019), and I show the results in Figure 3.4. It displays the estimated effects of the birth of the first child on gross labor income. For these estimates, I normalize the coefficients by dividing by the average gross labor income at event time −1. This normalization facilitates interpretation since regressions are run in levels rather than logs, allowing inclusion of individuals with zero earnings (e.g., individuals out of the labor force).

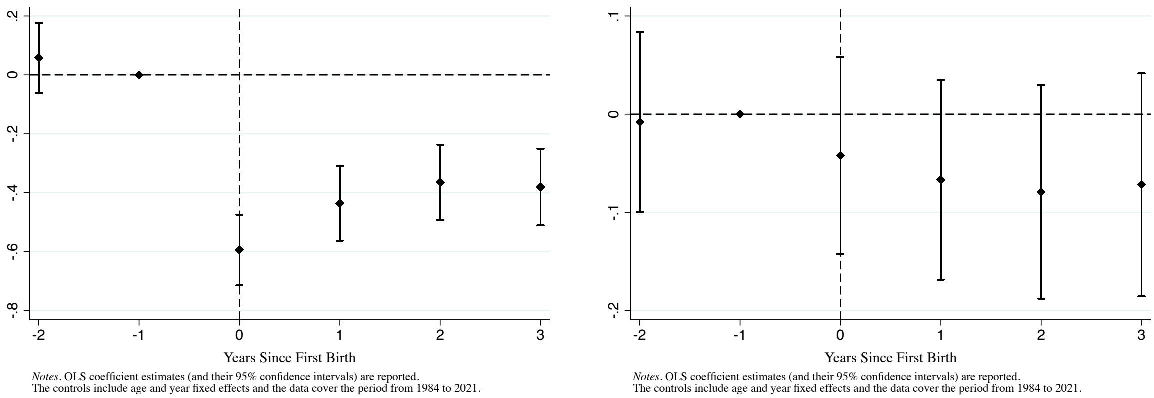

(a) Women (b) Men

Figure 3.4 – Birth and gross labor income (The Child Penalty)

A sharp decline in earnings occurs around childbirth for women, while no significant change is observed for men. This decline may be partially explained by the large drop in full-time employment documented earlier as well as the increase in part-time employment.

In summary, the birth of the first child has no statistically significant impact on men's labor market outcomes with respect to income, full-time employment, part-time employment, or job transitions. In contrast, for women, childbirth leads to a substantial reduction in earnings, a significant decline in the probability of full-time work, and an increased probability of changing jobs.

3.5 Model

This section presents a model to study statistical discrimination against women. The model is motivated by the observation that women are less likely to work full-time and more likely to change jobs or work part-time following the birth of their first child as documented in Section 3.4. I interpret this as evidence that women are more likely than men to assume the role of primary caregiver. The purpose of the model is to better understand how this caregiving role leads to discrimination against women not only *ex* *post*, i.e., after assuming caregiving responsibilities, but also *ex ante*, i.e., before the birth of the first child.

3.5.1 Economic Environment

I consider an economy that lasts for two periods. It consists of four different types of workers and two types of firms. A worker's type is defined by two characteristics: their sex (male or female) and whether they have caregiving responsibilities. I assume asymmetric information: the worker knows their own type, but the firm only observes the worker's sex and cannot observe caregiving responsibilities. No worker has caregiving responsibilities in the first period. However, women are more likely than men to have caregiving responsibilities in the second period. There are two types of firms: 1) an L-type firm, which performs a low-productivity task that can be done flexibly, and 2) an H-type firm, which performs a high-productivity task that cannot be carried out flexibly.

3.5.2 Workers

Workers without caregiving responsibilities derive utility solely from consumption, while workers with caregiving responsibilities derive utility from both consumption and job flexibility. Male and female workers share the same preferences conditional on caregiving status. Let γ_F be the fraction of female workers who have caregiving responsibilities in the second period, and γ_M be the corresponding fraction of male workers. No workers have caregiving responsibilities in the first period. There are no financial markets, so workers consume their entire earnings. The utility of the outside option to the worker of not working is normalized to 0, and the utility that workers derive from a contract that specifies an income (w) and flexibility (f) is summarized as follows:

$$u(w, f) = w + \beta_i f, \quad i \in \{C, NC\}$$

where C denotes caregiving responsibilities and NC denotes no caregiving responsibilities. I assume $\beta_{NC} = 0$, so workers without caregiving responsibilities do not value flexibility, whereas $\beta_C > 0$. The parameter β is private information that is known only to the worker. The probability of $\beta > 0$ is γ_F for a female worker and γ_M for a male worker. Motivated by the fact that women appear to be the ones to drastically reduce their labor market participation, I am going to assume that women are more likely than men to have caregiving responsibilities in the second period.

Assumption 3.5.1. $\gamma_F > \gamma_M$

This assumption is key for statistical discrimination to arise against women as firms are only able to observe sex, but not future caregiving status.

3.5.3 Firms

There are two types of firms in this economy, each producing a different task: 1) L-type firms, which perform routine tasks (task L) that allow job flexibility, and 2) H-type firms, which require physical presence and long hours, making flexibility impossible, producing task H. Below, I describe each task in detail.

Task L

An L-type firm produces an output θ_L when it opens a vacancy, employs a worker, and carries out task L. I assume no search frictions and free entry into vacancy openings. The firm's profit depends on the contract it offers, which specifies the wage w and the amount of flexibility f. The firm's profit function is:

$$\pi(w, f) = \theta_L - w - \frac{1}{2}af^2$$

where a is a cost parameter, and I assume that the cost of providing flexibility is convex. If the firm does not open a vacancy, profits are zero.

Task H

H-type firms cannot offer flexibility, but the task yields a higher surplus. The task requires an upfront investment κ in the worker to equip them with the necessary skills to perform the job. If the firm opens a vacancy, it offers a wage contract. The firm's profit in period 1 is:

$$\pi_1 = \theta_H - \kappa - w$$

where θ_H is the surplus generated, w is the wage paid to the worker, and κ is the investment cost.

In period 2, profits depend on whether the worker stays or leaves: If the worker stays, the firm produces output and pays a wage w, yielding profits:

$$\pi_2^S = \theta_H - w$$

since the firm has already equipped the worker with the necessary skills, period two profits are higher than period one profits if the worker stays. This gives the firm an incentive to hire workers that they believe will not switch jobs.

If the worker leaves, the profits are simply 0.

$$\pi_2^L = 0$$

The firm's problem is to design contracts that maximize total profits over both periods, subject to the IR constraints of the workers, i.e.

$$w \geq 0, \quad (IR_{NC})$$

for non-caregiving workers, and

$$w + \beta f \geq 0, \quad (IR_C)$$

for caregiving workers.

The IC constraint will be described more carefully later as it is relaxed compared to the standard case as firms are allowed to reject workers that are willing to work for the proposed contract.

I will make the following three assumptions.

Assumption 3.5.2. $\theta_H - \frac{1}{2}\kappa > \theta_L$ *and* $\theta_H - \kappa < \theta_L$

and

Assumption 3.5.3. $\theta_H - \frac{1}{2}\kappa < \theta_L + \frac{1}{2}\frac{\beta^2}{a}$ *and* $\theta_H - \frac{\kappa}{2-\gamma_F} > \theta_L$

Assumption 3.5.2 ensures that κ is such that it is efficient for workers without caregiving responsibilities to be employed by the high-productivity firm, while it is inefficient for workers with future caregiving responsibilities to be employed by the high-productivity firm.

Assumption 3.5.3 ensures that workers with caregiving responsibilities prefer to work for the low-productivity firm that can offer flexibility. If this assumption were to be violated, the results of the model would be trivial as all types of workers would prefer to work for the high-productivity firm. Assumption 3.5.3 further ensures that women without caregiving responsibilities would prefer to work for the high-productivity firm when there is asymmetric information.

3.6 Equilibrium

Each worker chooses their preferred job contract from the available options. A job contract specifies a wage and the amount of flexibility offered. The task L firm and the task H firm offer the contracts that maximize their profits.

I will first solve for the first-best equilibrium to have a benchmark. I will then move on to solving for the equilibrium with asymmetric information, where firms are not able to distinguish workers with caregiving responsibilities from workers who do not have caregiving responsibilities.

3.6.1 First-best Equilibrium

In a first-best equilibrium where firms can observe the type of a worker, the proposed wage will be the one that maximizes total surplus. Due to free entry the firms make zero profit, so all of the surplus from the match goes to the workers.

Type H firm. The type H firm will propose a wage to the different types that is equal to the total surplus of the match. A worker without caregiving responsibilities will generate a surplus of θ_H in each period, while the firm will have to pay an upfront investment cost of κ in period one. Conversely, a worker with caregiving responsibilities in period 2 will only generate a surplus of θ_H in the first period, while they will start working for the low-productivity firm in period two as that firm can offer flexibility, which is more attractive for caregiving workers according to assumption 3.5.3. The proposed wage by the high-productivity firm will therefore be:

$$w_{NC} = \theta_H - \frac{1}{2}\kappa$$
$$w_C = \theta_H - \kappa$$

Hence, workers with no caregiving responsibilities are proposed a higher wage as they will stay with the firm for two periods (and thereby have a lower per-period cost of the upfront investment), whereas workers with caregiving responsibilities will only stay with the firm for one period and therefore generate a lower surplus implying that they get a lower wage offer.

Type L firm. The type L firm would instead offer a contract with a wage of θ_L to individuals without caregiving responsibilities. This is due to the fact that these workers do not value flexibility so they have no incentive to offer flexibility to this type of worker as it is costly.

For workers that value flexibility, they will offer the contract that maximizes total surplus. Total surplus is equal to the surplus generated by the match and the utility that the worker gets from flexibility subtracted by the cost to the firm of providing flexibility. Hence, the amount of flexibility is the solution to the following maximization problem.

$$f^* = \max_f \{\theta_L - \tfrac{1}{2}af^2 + \beta f\}$$

Maximizing this with respect to flexibility yields $f^* = \frac{\beta}{a}$, and using the zero-profit condition the wage becomes $\theta_L - \frac{1}{2}\frac{\beta^2}{a}$. Hence, the contract proposed to workers with caregiving responsibilities is

$$(w, f) = \left(\theta_L - \tfrac{1}{2}\tfrac{\beta^2}{a}, \tfrac{\beta}{a}\right)$$

Due to assumption 3.5.2, workers without caregiving responsibilities will prefer the contract offered by the high-productivity firm, while workers with caregiving responsibilities will prefer the contract by the low-productivity firm. Therefore, we have a first-best equilibrium that is characterized by the following (w, f) contracts shown in table 3.3 below

Type	Task	Period 1 Contract	Period 2 Contract
Non-caregiving women	H	$(\theta_H - \frac{1}{2}\kappa, 0)$	$(\theta_H - \frac{1}{2}\kappa, 0)$
Caregiving women	L	$(\theta_L, 0)$	$\left(\theta_L - \frac{1}{2}\frac{\beta^2}{a}, \frac{\beta}{a}\right)$
Non-caregiving men	H	$(\theta_H - \frac{1}{2}\kappa, 0)$	$(\theta_H - \frac{1}{2}\kappa, 0)$
Caregiving men	L	$(\theta_L, 0)$	$\left(\theta_L - \frac{1}{2}\frac{\beta^2}{a}, \frac{\beta}{a}\right)$

Table 3.3 – First-Best Contracts

We see that in the first-best equilibrium, there is no gender inequality conditional on type. However, men will, on average, earn a higher income due to assumption 3.5.1 that states that women are more likely to be the primary caregiver.

3.6.2 Second Best Equilibrium

The model's equilibrium comprises the optimal contracts proposed by both firm types, given the workers' optimal responses. The problem of the worker is to choose the contract that maximizes utility. Male and female workers are identical in their preferences, so I will just solve the decision problem for the caregiving and non-caregiving workers.

Non-caregiving workers. The problem for the non-caregiving worker is trivial as they only value consumption. In each period, they will therefore simply choose the contract that offers the highest wage. Let $\mathcal{C} = \{c_1, c_2, \ldots, c_n\}$ be the set of available contracts in a given period, and let each contract $c_i \in \mathcal{C}$ offer a wage $w_i = w(c_i)$ and a level of flexibility $f_i = f(c_i)$. The worker's problem is therefore:

$$\max_{c_i \in \mathcal{C}} \; u(w_i)$$

Since u is strictly increasing, the optimal contract c^* is the one that offers the highest wage:

$$c^* = \arg\max_{c_i \in \mathcal{C}} \; w_i$$

Caregiving workers. The problem of the caregiving worker is to choose the contract that maximizes utility.

$$\max_{c_i \in \mathcal{C}} \; u(w_i, f_i) = \max_{c_i \in \mathcal{C}} \{w_i + \beta f_i\}$$

I can now solve for the optimal contracts proposed by type L and type H firms given the behavior of workers.

Optimal contract for type L firm

Due to the free-entry condition, we immediately deduce that profits must be 0 in equilibrium. Hence, we have

$$\theta_L = w + \tfrac{1}{2}af^2$$

i.e., the revenue generated by the task (θ_L) must be equal to the cost of the contract which consists of wage payments and the cost of offering flexibility.

Firms face the problem of how to optimally design the contract that they offer to workers. To solve this problem, the firm realizes that it faces different types of workers. It faces workers that value flexibility and workers that don't value it. The workers that don't value flexibility have the following IR constraint

$$w \geq 0, \quad (IR_{NC})$$

whereas the workers that value flexibility have the following one:

$$w + \beta f \geq 0, \quad (IR_C)$$

The first result that follows from the model is the following.

Proposition 3.6.1. *There cannot exist a pooling equilibrium where firms carrying out task L only offer one contract.*

The proof is in appendix 3.A. The intuition of this result is that if there were to be a pooling equilibrium where all workers were offered a positive level of flexibility, it would be possible to offer an alternative contract that would propose zero flexibility and increase the wage by a small ϵ. That contract would be preferred by workers without caregiving responsibilities, while the firm would make a positive profit. Similarly, if there were a pooling equilibrium where no contract offered a positive level of flexibility, a firm could attract workers with caregiving responsibilities by offering a small amount of flexibility while reducing the wage.

Separating equilibrium. As there cannot be a pooling equilibrium, I will solve for a separating equilibrium where firms offer two different contracts and people self-select into their preferred one. The contract for workers that do not have caregiving responsibilities is trivial. Given the zero-profit condition which we will write as the condition that the contract must lie in the set defined below

$$A = \{(w, f) : w + \tfrac{1}{2}af^2 = \theta_L\}$$

and the IC constraint for the workers in period 1

$$(w, f) = \arg\max_{\tilde{w},\tilde{f}\in A} u(w, f)$$

We immediately obtain

$$w = \theta_L$$

due to the fact that $u(w, f) = w$.

For workers that value flexibility, the firm would like to minimize the cost yet satisfy the IC constraint. From the worker problem, we know that the IC constraint is

$$(w, f) \in \arg\max_{\tilde{w},\tilde{f}\in A} \{\tilde{w} + \beta\tilde{f}\}$$

Using this to replace for the wage in the profit function and letting $\bar{u}$ denote the utility to the worker, we get

$$\pi^E = \theta_L - \bar{u} + \beta f - \tfrac{1}{2}af^2$$

The first-order condition yields

$$\beta = af$$

Intuitively, β captures the gain to the firm by marginally increasing the amount of flexibility, permitting them to lower wages by β. In equilibrium, this must be equal to the marginal cost of providing more flexibility.

Solving for f, we obtain

$$f = \frac{\beta}{a}$$

We see that the amount of flexibility provided by firms is an increasing function of β. The intuition is that a higher value of β implies that the firm can reduce more the wage it has to offer when increasing the level of flexibility. We also find that the amount of flexibility offered is a decreasing function of a. The higher the cost of providing flexibility, the less flexibility the firm will offer.

Using the zero-profit condition, we can determine the wage which will be

$$w = \theta_L - \frac{1}{2}\frac{\beta^2}{a}$$

Hence, the firms producing task L will offer the following two contracts

$$(w_1, f_1) = (\theta_L, 0) \text{ and } (w_2, f_2) = \left(\theta_L - \frac{1}{2}\frac{\beta^2}{a}, \frac{\beta}{a}\right)$$

That is, there will be a separating equilibrium where the firms will offer two different contracts. They will offer a contract that provides 0 flexibility which will be preferred by people just wanting to consume, and there will be another contract that provides a positive level of flexibility which will be preferred by caregiving workers with children.

It remains to be shown that these two contracts satisfy the IR constraints. For people not valuing flexibility, the contract yields a payoff of $\theta_L > 0$ so the IR constraint is indeed satisfied. The workers who value flexibility will obtain a utility of

$$U_C = \theta_L - \frac{1}{2}\frac{\beta^2}{a} + \frac{\beta^2}{a} = \theta_L + \frac{1}{2}\frac{\beta^2}{a} > 0$$

Hence, the IR constraints are indeed satisfied.

3.6.3 Optimal contracts for type H firm

The firm knows that there are four types of workers, yet the firm is only able to distinguish a man from a woman. It cannot distinguish a worker with a preference for flexibility from a worker who does not have such a preference.

Due to assumption 3.5.3 a worker with caregiving responsibilities will prefer the contract offered by the type L firm to the most favorable contract that can be offered by the type H firm.

To solve for the optimal contract, we will thus focus on the case where the firm only wants to attract workers with no caregiving responsibilities. The relevant IR and IC constraints are therefore

$$w > 0, \qquad \text{(IR)}$$
$$w > \theta_L, \qquad \text{(IC)}$$

We will first guess that IC is satisfied, and then we will verify ex post. IC trivially implies IR. To solve for the contract, we first use that the expected profits must be zero due to free entry. Hence, we have

$$\theta_H - w - \kappa + (1 - \gamma_i)(\theta_H - w) = 0, \quad i = \{m, f\}$$

where we have used the assumption that workers with caregiving responsibilities will leave the firm in period 2. We can then solve for the wage that the firm must offer to men and women

$$w_M = \theta_H - \frac{\kappa}{2 - \gamma_M}$$
$$w_F = \theta_H - \frac{\kappa}{2 - \gamma_F}$$

and we notice that the IC constraint is satisfied due to assumption 3.5.3.

We see that the wage that men and women will receive will depend on the likelihood that they will have caregiving responsibilities in period 2. Under assumption 3.5.1 we have that women are more likely than men to have caregiving responsibilities. Hence, if firms are forced to offer the same wage to men and women for performing the same task, we have the following result

Proposition 3.6.2. *If firms are not able to condition wages on gender, no women will be employed in task H if* $\gamma_F > \gamma_M$.

The proof is in appendix 3.A.

Women will, of course, prefer the high-paying contract, but firms will have to reject them as employing them is associated with an expected loss. We will make the assumption that it is feasible for firms to undertake this kind of discrimination as it will be very difficult to detect if someone was not chosen for a job due to the gender of the person. Women will after this rejection apply to the contract proposed by the firms performing task L that pays less. All women will be employed to perform task L, which will be inefficient as the women with non-caregiving responsibilities should be employed in task H, while the men with caregiving responsibilities should be employed in task L.

3.7 Discussion

Based on the previous section, we can state the equilibrium contracts in period 1 in the first-best (FB) and second (SB) equilibrium in table 3.4. Hence, these are the contracts that the workers will have prior to the birth of a child.

Type	FB contract	SB contract
Non-caregiving women	$\left(\theta_H - \frac{1}{2}\kappa, 0)\right)$	$(\theta_L, 0)$
Caregiving women	$(\theta_L, 0)$	$(\theta_L, 0)$
Non-caregiving men	$(\theta_H - \frac{1}{2}\kappa, 0)$	$(\theta_H - \frac{\kappa}{2-\gamma_M}, 0)$
Caregiving men	$(\theta_L, 0)$	$(\theta_H - \frac{\kappa}{2-\gamma_M}, 0)$

Table 3.4 – Equilibrium contracts period 1

We see that there is no difference in the income received by men and women conditional on caregiving responsibilities in the first-best equilibrium in period 1. This is intuitive as men and women are equally productive, and equally likely to stay in the job conditional on caregiving responsibilities.

However, this does not hold true in the case of asymmetric information. This follows from the informativeness principle of Holmstrom (1979): gender is a signal about caregiving in the future. Thus, the optimal contract would incorporate all the informative signals about payoff relevant variables. Gender is irrelevant if caregiving responsibility is observed, that is why it is not included in the first-best.

Caregiving women are indifferent between the first-best equilibrium under full information, and the second best equilibrium where firms are not able to distinguish workers with caregiving responsibilities from workers without. Men with caregiving responsibilities on the other hand are better off. In the first-best equilibrium, firms are able to see that they will eventually have caregiving responsibilities, and they will therefore end up in the low-productivity firm that does not require an initial investment in the worker. In the second best, however, firms are not able to tell men with future caregiving responsibilities from men without. Hence, men with future caregiving responsibilities benefit from this as the high-productivity firm is willing to hire them with the hope that they are a worker without future caregiving responsibilities. On the other hand, men without caregiving responsibilities are worse off.

We see that non-caregiving women are worse off compared to the first-best equilibrium. The reason for this is that caregiving women exert a negative externality on them. As firms cannot tell that a woman does not have future caregiving probabilities with certainty, they will have to make a judgment based on the probability only conditional on gender. As women are more likely than men to face caregiving responsibilities, the women without will therefore appear, under asymmetric information, to be a worse investment for the initial training cost than a man with caregiving responsibilities. This inefficiency trivially does not arise in a first-best. Women without future caregiving responsibilities would therefore benefit if they could more easily signal this or if firms could better screen.

3.8 Conclusion

This paper demonstrates that hiring discrimination against women can arise in environments characterized by imperfect information, wherein firms are unable to distinguish between workers with caregiving responsibilities and those without. Discrimination emerges *ex ante*, i.e. prior to the birth of a child, if firms hold the belief that women are more likely than men to assume primary caregiving roles. Consequently, even women without caregiving responsibilities may be excluded from positions involving substantial upfront training investments, despite being optimal candidates in a first-best equilibrium with full information. The problem stems from the fact that prospective caregivers initially mimic the behavior of non-caregivers, thereby rendering it impossible for firms to differentiate between the two types.

References

Adda, Jérôme, Christian Dustmann, and Katrien Stevens (2017). « The Career Costs of Children ». In: *Journal of Political Economy* 125.2, pp. 293–337.

Akerlof, George A. (1970). « The Market for 'Lemons': Quality Uncertainty and the Market Mechanism ». In: *The Quarterly Journal of Economics* 84.3, pp. 488–500.

Albanesi, Stefania and Claudia Olivetti (2009). « Home production, market production and the gender wage gap: Incentives and expectations ». In: *Review of Economic Dynamics* 12.1, pp. 80–107.

Arrow, K. J. (1973). « The Theory of Discrimination ». In: Princeton University Press.

Azmat, Ghazala and Rosa Ferrer (2017). « Gender Gaps in Performance: Evidence from Young Lawyers ». In: *Journal of Political Economy* 125.5, pp. 1306–1355.

Becker, G. S. (1957). « The economics of discrimination ». In: Chicago: University of Chicago Press.

Blau, Francine D. and Lawrence M. Kahn (2013). « Female Labor Supply: Why Is the United States Falling Behind? » In: *American Economic Review* 103.3, pp. 251–56.

Doepke, Matthias and Fabian Kindermann (2019). « Bargaining over Babies: Theory, Evidence, and Policy Implications ». In: *American Economic Review* 109.9, pp. 3264–3306.

Erosa, Andres, Luisa Fuster, and Diego Restuccia (2016). « A quantitative theory of the gender gap in wages ». In: *European Economic Review* 85, pp. 165–187.

Erosa, Andrés et al. (2022). « Hours, Occupations, and Gender Differences in Labor Market Outcomes ». In: *American Economic Journal: Macroeconomics* 14.3, pp. 543–90.

Gayle, George-Levi and Limor Golan (2012). « Estimating a Dynamic Adverse-Selection Model: Labour-Force Experience and the Changing Gender Earnings Gap 1968-1997 ». In: *The Review of Economic Studies* 79.1, pp. 227–267.

Goldin, Claudia (2014). « A Grand Gender Convergence: Its Last Chapter ». In: *American Economic Review* 104.4, pp. 1091–1119.

Holmstrom, Bengt (1979). « Moral Hazard and Observability ». In: *Bell Journal of Economics* 10.1, pp. 74–91.

Kleven, Henrik, Camille Landais, Johanna Posch, et al. (2019). « Child Penalties across Countries: Evidence and Explanations ». In: *AEA Papers and Proceedings* 109, pp. 122–26.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard (2019). « Children and Gender Inequality: Evidence from Denmark ». In: *American Economic Journal: Applied Economics* 11.4, pp. 181–209.

Nix, Emily and Martin Eckhoff Andresen (2019). *What Causes the Child Penalty? Evidence from Same Sex Couples and Policy Reforms*. Discussion Papers 902. Statistics Norway, Research Department.

Olivetti, Claudia and Barbara Petrongolo (2017). « The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries ». In: *Journal of Economic Perspectives* 31.1, pp. 205–30.

Phelps, Edmund S. (1972). « The Statistical Theory of Racism and Sexism ». In: *The American Economic Review* 62.4, pp. 659–661.

Rosen, Sherwin (1986). « Chapter 12 The theory of equalizing differences ». In: vol. 1. Handbook of Labor Economics. Elsevier, pp. 641–692.

Sloane, Carolyn M., Erik G. Hurst, and Dan A. Black (2021). « College Majors, Occupations, and the Gender Wage Gap ». In: *Journal of Economic Perspectives* 35.4, pp. 223–48.

Appendices

3.A Proofs

Proof of Proposition 3.6.1. By contradiction. Assume that there exists a pooling equilibrium where the firm offers only one contract to the workers. Consider first the case where this contract offers some positive level of flexibility $f > 0$.

Necessarily, this contract must offer a wage that is less than θ_L to yield non-negative profits. This cannot be an equilibrium as it would be possible to offer another contract with $f = 0$ and $w' = w + \epsilon$. For ϵ chosen sufficiently small this contract will provide the firms with positive profits. The workers that do not value flexibility will trivially prefer this zero-flexibility contract to the other one. Hence, it is not feasible to have a contract in a pooling equilibrium that delivers $f > 0$.

Assume for now that there exists a pooling equilibrium where the firm offers the following contract $(w, f) = (\theta_L, 0)$. It would be possible to offer a contract that offers slightly more flexibility. The benefit would be that the firm could offer the wage

$$w' = w - \beta \Delta f$$

and the marginal cost for doing this would be

$$-\frac{\partial \pi}{\partial f}(\Delta f) = af(\Delta f)$$

where for f small enough this is less than $\beta(\Delta f)$.

Hence, there cannot exist a pooling equilibrium with $f = 0$ as it would be possible for the firm to offer an alternative contract that would yield positive profits. □

Proof of Proposition 3.6.2. Given that firms have to pay the same wage to men and women, they will need to base the contract on the average probability that a worker leaves the firm. But then firms would have an incentive to discriminate in hiring rates against women. A firm employing a man would then make positive profits while a firm employing a woman would make a negative profit. The unique equilibrium is that the firm proposes the wage based on the probability that a man leaves

$$w = \theta_H - \frac{\kappa}{2 - \gamma_M}$$

and it employs all the men that apply to the job. It will decide to refuse all the women that apply as hiring them would cause negative profits. □