

Price discrimination, switching costs and welfare: Evidence from the Dutch mortgage market*

Jurre H. Thiel[†]

Job Market Paper

[Click here for the latest version.](#)

January 20, 2019

Abstract

In many markets with switching costs firms charge a lower price to new customers than to existing customers, a practice called history-based price discrimination. By exploiting a ban on history-based price discrimination in the Dutch mortgage market, this paper estimates the effects of history-based price discrimination on consumer surplus, profits and welfare. These effects are theoretically ambiguous because history-based price discrimination can make markets more or less competitive. I estimate a structural model, of which the supply side consists of a dynamic game. To deal with the curse of dimensionality, I employ techniques from machine learning to reduce the dimension of this game's state space. I implement these techniques in a new estimation method for dynamic games, which I justify with a new solution concept: Sparse Markov Perfect Equilibrium (SMPE). In an SMPE, firms optimally pay attention to a subset of state variables instead of the full state. Therefore, the state space is considerably smaller than under the standard assumption of Markov Perfect Equilibrium. I show that the Lasso identifies which variables firms pay attention to. For an average mortgage, banning history-based price discrimination increases welfare by €125 per year and consumer surplus by €415 per year, while bank profits drop by €290 per year.

*I thank Jaap Abbring, Victor Aguirregabiria, Eric Bartelsman, Elisabeth Honka, Alessandro Iaria, András Kiss, Alexander MacKay, Will Wang and Matthijs Wildenbeest and particularly Mauro Mastrogiacono and José-Luis Moraga-Gonzalez for their many comments and suggestions. This paper has benefitted from presentations at the University of Toronto, VU Amsterdam, EARIE 2017 (Maastricht), the Tinbergen Workshop on Empirical IO (Amsterdam), the 2018 RES Symposium of Young Researchers (Brighton), the University of Tilburg, the 2018 North American Summer Meeting of the Econometric Society (Davis), the 2018 Barcelona GSE Summer Forum Workshop on Consumer Search and Switching Costs, the TI Jamboree (Amsterdam) and EARIE 2018 (Athens). Views expressed are those of the author and do not necessarily reflect official positions of De Nederlandsche Bank. In this paper use is made of data of the DNB Household Survey.

[†]Vrije Universiteit Amsterdam and Tinbergen Institute. E-mail: j.h.thiel@vu.nl.

1 Introduction

In many markets with switching costs firms charge a lower price to new customers than to existing customers. This phenomenon, called history-based or behavior-based price discrimination, has been documented for credit markets (Ioannidou and Ongena 2010; Barone, Felici, and Pagnini 2011), cellular contracts (Alé 2013) and newspapers (Asplund, Eriksson, and Strand 2008). Whether history-based price discrimination increases or decreases welfare is theoretically ambiguous. Compared to uniform pricing, firms have an incentive to charge relatively high prices to renewing customers, because they face switching costs. On the other hand, they want to charge relatively low prices to new customers to entice them to switch. Since either effect can dominate, the effect of history-based price discrimination on welfare is an empirical question. Despite the fact that history-based price discrimination is common, this question has so far remained unanswered.

This paper exploits a natural experiment in the Dutch mortgage market to estimate the effects of history-based price discrimination on consumer surplus, bank profits and welfare. To do so, I develop a structural model of demand and supply of this market. The supply side of the model is a dynamic game and suffers from the curse of dimensionality. To estimate the model, I therefore introduce a new method to estimate dynamic games with large state spaces. The method is based around a new solution concept for dynamic games, Sparse Markov Perfect Equilibrium, in which firms display partial attention to the state. I show that estimation techniques from the machine learning literature can identify the variables firms pay attention to.

In the Dutch mortgage market, most households fix their interest rate for a certain period, most commonly ten years. When this period ends, they can either renew their mortgage at their current bank or switch to a different one. However, switching is costly: for the average household, switching costs are around €3500 in addition to the opportunity cost of time.¹ Therefore, banks have an incentive to charge high interest rates to existing customers, since they are “locked in” because of switching costs. This is called the *rent extraction effect*. On the other hand, poaching customers from other banks becomes more attractive compared to uniform pricing because history-based price discrimination allows banks to charge them lower interest rates without cannibalizing the profits on their captive customers. This downward pressure on interest rates is called the *competition effect*. Depending on whether the rent-extraction or the competition effect dominates, the average interest rate can be higher or lower under history-based price discrimination than under uniform pricing. Moreover, because the price difference between a consumer’s current firm and its competitors is larger than under uniform pricing, history-based price discrimination encourages switching. When switching is costly, this is socially wasteful. Finally, cross-segment inefficiencies may occur when a consumer purchases from an inefficient bank because it obtains a relatively low interest rate there (Stole 2007). For these reasons, the question whether history-based price discrimination increases

1. Typical costs include advice, taxation, notary and insurance fees. A calculation by the country’s largest mortgage broker suggests total costs of about €6000 for an average household. Since these costs are tax deductible and the average marginal income tax rate is 42%, the average household will incur monetary costs of about €3500. (<https://www.hypotheke.nl/jouw-woonsituatie/hypotheek-oversluiten/>, Accessed March 8, 2017.)

or decreases consumer surplus, profits and total welfare can only be answered empirically.²

Consistent with this theory, prior to 2013, Dutch banks offered higher interest rates to renewing customers than to new customers. Regulators were concerned that such history-based price discrimination was harmful to consumers. As a result, regulations were introduced stipulating that from 2013 onwards banks have to offer the same interest rate to new and to renewing customers if those customers have a similar risk profile. I exploit this ban to study the effects of history-based price discrimination on consumer surplus, bank profits and total welfare. I do so using administrative data from the Dutch central bank, which include the universe of mortgages from institutions under its supervision.³ To rule out as much as possible that interest rate differences between customers are caused by differences in risk, I focus on mortgages that are insured by the Dutch government.

I find that before the 2013 ban there were indeed significant interest rate differences between renewing and new customers. On average, a renewing household paid an interest rate that was .32 percentage points higher than a new customer. Since the average mortgage in my sample is about €150,000 this means that renewing households paid around €278 per year more in interest (after tax deductions).

To assess the effects of the ban on history-based price discrimination, I estimate a structural model of demand and supply of the Dutch mortgage market. Because higher sales today imply more locked-in consumers in the future, banks play a dynamic game. On the demand side, I allow for rich interactions between household and product characteristics in consumers' utility specification, since the curvature of the demand function is a crucial determinant of the effect of third-degree price discrimination on welfare (Holmes 1989). This demand-side heterogeneity implies that the state space of the dynamic game that banks play is very large: the pay-off relevant state contains the full joint density of previous market shares and household demographics. Traditional methods for estimation of dynamic games cannot deal with games with such large state spaces.⁴ Two solutions to this problem have been used in the literature. One is to restrict the amount of heterogeneity in the model, for example by having only a small number of types of consumers in the market.⁵ This is not an attractive option when studying price discrimination, since this limits the shapes the demand function can take. A second solution is to make ad hoc assumptions to reduce the dimension of the state space.⁶ However, there are typically many different ad hoc assumptions one

2. The theoretical literature on history-based price discrimination in the presence of switching costs (Chen 1997; Taylor 2003; Gehrig, Shy, and Stenbacka 2012; Rhodes 2012) generally finds that, compared to uniform pricing, history-based price discrimination can both increase and decrease consumer welfare, and is divided on the predicted effect on firm profits and total welfare.

3. These institutions have a combined market share of 75% - 80% (Mastrogiacomo and Van der Molen 2015).

4. Most methods, most prominently Bajari, Benkard, and Levin (2007), contain a first stage in which firms' actions are regressed on state variables. If the number of state variables is larger than the number of observations, this is impossible. So-called nested fixed point methods, in which an equilibrium is calculated for every candidate parameter vector, take too much time when the state space is large. The nested fixed point method of Abbring et al. (2017) is fast, but only applicable to models of firm entry and exit.

5. One recent example of this approach is Cosguner, Chan, and Seetharaman (2018).

6. For example, many papers group states together. A prominent example is Collard-Wexler (2013), who groups plants of different sizes together and ignores markets with too many firms to reduce the state space. Another common strategy is

could make. Moreover, I show that, at least for my application, some on the face reasonable ad hoc assumptions lead to very strange results.

I instead solve the curse of dimensionality by introducing a new method for estimating dynamic games. The method combines economic theory with techniques from machine learning. On the theory side, I dispense with the assumption that firms play a Markov Perfect Equilibrium (MPE). Instead, I introduce a new solution concept that is more amenable to estimation: Sparse Markov Perfect Equilibrium (SMPE). An SMPE is an MPE in which firms perform sparse maximization. Under sparse maximization, which was introduced by Gabaix (2014) and which I extend to dynamic games, firms optimally condition their policy functions only on a subset of the payoff-relevant state variables. This can be motivated by relaxing the usual implicit assumption that firms' information on the state is free. The relevant state in an SMPE is smaller than in an MPE so that an SMPE can be estimated in situations in which estimating an MPE is impossible. Which state variables firms pay attention to can be estimated using variable selection techniques from the machine learning literature.

Econometrically, I use a two-stage method as in Bajari, Benkard, and Levin (2007), where I use the Lasso in the first stage to estimate the policy functions. The Lasso selects the state variables firms pay attention to. For a given estimated model, SMPE imposes testable restrictions on the data. These restrictions can be used to test whether the econometric selection of state variables has a sensible economic interpretation. Thus, it is the interplay between theory and econometrics which allows me to deal with the curse of dimensionality.

I find that the ban on history-based price discrimination leads to economically significant increases in welfare, of about €125 per mortgage or €11.5 million in total per year. The ban causes a drop in interest rates, so that they are closer to marginal costs. Moreover, the ban leads to less socially wasteful switching. Finally, there is a reallocation effect: the ban causes the market share of low-cost banks to increase. The welfare increase accrues more than completely to consumers: consumer surplus increases with about €415 per year for an average mortgage. On the flip side, the average profit per mortgage decreases by €290.

This paper makes several contributions. First, I am, to the best of my knowledge, the first to empirically study whether history-based price discrimination increases or decreases welfare. The only other empirical study on history-based price discrimination in markets with switching costs is Cosguner, Chan, and Seetharaman (2017), who only consider the effect on profits. In particular, I find that cost asymmetries between firms can be of first-order importance for the welfare effects of history-based price discrimination. Although the importance of asymmetries has sometimes been recognized in the literature on price discrimination generally (e.g. Stole (2007)), the theoretical literature on history-based price discrimination has focused only on symmetric cost functions. Therefore, I show that this assumption is not innocuous and that cost asymmetries can be very important. This mirrors the recent finding of Asker, Collard-Wexler, and De Loecker (2017), who

to assume that firms' behavior only depends on some average or total value of the state instead of the full distribution. For example, Kalouptside (2014), assumes that ship manufacturers' value functions depend only on the total backlog in the market and Barwick and Pathak (2015) assume that real estate agent's commissions depend only on average housing market conditions. Another strategy is to assume the policy functions take a particular form (Wollmann 2018).

find that cost asymmetries are crucial to understand the effects of market power.

The second contribution of this paper is that it provides a tractable empirical model of supply for markets with switching costs. Switching costs, or demand inertia more broadly, have been documented in many markets.⁷ These include markets that are highly policy-relevant, such as health care (Nosal 2012; Handel 2013). Demand inertia also plays a key role in many macroeconomic models with so-called “customer markets” (Phelps and Winter 1970; Bils 1989), and was recently found to play an important role in inflation dynamics (Gilchrist et al. 2017). To assess potential policy interventions in markets with switching costs, it is important to understand how firms might respond to such measures. Some previous work, which I discuss in more detail below, has developed empirical models of firm behavior in markets with switching costs. However, these models use certain assumptions that limit their applicability. The method I introduce in this paper, SMPE, is generally applicable and can (contrary to previous models) accommodate rich heterogeneity, non-anonymous strategies and a large number of firms.

The final contribution of this paper is to provide a new method to estimate dynamic games with large state spaces. Here, I contribute to the literature on the estimation of dynamic games (Aguirregabiria and Mira 2007; Bajari, Benkard, and Levin 2007; Pakes, Ostrovsky, and Berry 2007; Pesendorfer and Schmidt-Dengler 2008). My method is most closely related to Bajari, Benkard, and Levin (2007). Like them, I follow a two-step method, where in the first step firms’ policy functions are regressed on state variables and in a second step the deep parameters of the model (typically marginal costs) are recovered. My main innovation is to use the Lasso in the first step. For this, I also give a micro-foundation in the form of SMPE.

The remainder of this paper is organized as follows. Section 2 discusses the related literature in more detail. Section 3 gives some background on the Dutch mortgage market and its ban on history-based price discrimination. Section 4 describes my data set and gives reduced-form evidence on the existence of history-based price discrimination in the Dutch mortgage market. Section 5 describes the structural model of demand and supply that I estimate. It also introduces the concept of Sparse Markov Perfect Equilibrium. Section 6 discusses identification and estimation of my model. Section 7 gives the estimation results, as well as an assessment of the welfare effects of the ban on history-based price discrimination. Section 8 discusses the robustness of my model, while Section 9 concludes.

2 Related literature

In addition to the literature summarized in the introduction, this paper is broadly related to two strands of literature: one on history-based price discrimination and empirical models of state-dependent demand and the other on simplifying solution concepts for dynamic games.

While there is a relatively large literature documenting the existence of history-based discrimination (Asplund, Eriksson, and Strand 2008; Ioannidou and Ongena 2010; Barone, Felici, and Pagnini

7. For an overview of papers estimating switching costs, see the literature review below.

2011; Alé 2013), there is not much work estimating its effects. The aforementioned Cosguner, Chan, and Seetharaman (2017) is an exception. They use counterfactual simulations to study whether history-based price discrimination would be profitable in the cola industry. They however do not have the exogenous variation in pricing (with and without history-based price discrimination) that I have. In addition, I also look at consumer surplus and total welfare and not just at profits.

My paper is related to various other papers that estimate dynamic models of firm pricing in the presence of state-dependent demand. Fleitas (2017) develops a dynamic model of insurer pricing in Medicare Part D. His model restricts strategies to be anonymous and symmetric and does not allow for rich consumer heterogeneity in the demand model; my model has neither restrictions. The anonymity assumption is particularly restrictive when studying oligopolies: in such markets it is difficult to imagine that the identity of the firm that consumers purchase from is irrelevant. Cosguner, Chan, and Seetharaman (2018) develop an empirical model of supply in the presence of switching costs of the cola market. Their estimator, based on forward iteration of the value function, is feasible because their market features only two firms. As the number of states for which forward iteration of the value functions must be performed increases exponentially in the number of firms, their estimator would quickly become infeasible for markets with more firms. My approach can (and, in this paper, does) handle a larger amount of firms. MacKay and Remer (2018) develop a model of supply with demand inertia in the context of gasoline markets. They require observations on firms' marginal costs to estimate their model, my approach does not. Rickert (2016) studies the German diaper market: his model has a finite horizon, which is not an attractive assumption in many markets. None of these papers feature markets with history-based price discrimination.

More broadly, my paper contributes to a growing literature documenting and estimating switching costs (Viard 2007; Dubé, Hitsch, and Rossi 2009; Cullen and Shcherbakov 2010; Miller and Yeo 2012; Nosal 2012; Handel 2013; Honka 2014; Cullen, Schutz, and Shcherbakov 2015; Ho 2015; Shcherbakov 2016; Raval and Rosenbaum, Forthcoming; Weiergräber 2017). This literature estimates the effects of (different levels of) switching costs on market outcomes. I take the level of switching costs as given and study a potential consequence of switching costs, namely history-based price discrimination.

Secondly, my paper contributes to a literature on simplifying solution concepts for dynamic games. In recent years, various solution concepts have been introduced, such as oblivious equilibrium (Weintraub, Benkard, and Van Roy 2008), experience based equilibrium (Fershtman and Pakes 2012) and moment-based Markov equilibrium (Ifrach and Weintraub 2017). These approaches make *computation* of equilibria of dynamic games easier. The notion of SMPE that I introduce makes *estimation* easier. SMPE is closest to the moment-based Markov equilibrium of Ifrach and Weintraub (2017). Ifrach and Weintraub (2017) make the assumption that firms pay attention to the full state of dominant firms and some moments of the state of fringe firms. In an SMPE, no such assumptions are necessary. Instead, firms choose which variables they pay attention to based on a cost-benefit analysis. This means that SMPE can also be used when the dimension of a common state or of the states of dominant firms is large.

3 The Dutch Mortgage Market

Because the Dutch mortgage is in some aspects quite different from mortgage markets in other countries, I begin with a brief overview of its most important features. In the Netherlands, mortgages are primarily sold by banks. The market is reasonably concentrated, with an HHI of 2100. Three banks (ABN Amro, ING and Rabobank) dominate the market, with a competitive fringe consisting of smaller banks and pension funds. Approximately 55% of households own their house.⁸

In the Dutch mortgage market, many different types of mortgages are sold. Two categories can be distinguished. The first consists of non-amortizing mortgages—mortgages where the principal is paid in a lump sum at the end date. Such mortgages include bullet, savings, life and investment mortgages; the latter three are sold together with a financial product, the returns of which are used to pay off the principal at the mortgage's end date. The second type of mortgage are amortizing mortgages: for these regular payments towards the principal are made. Amortizing mortgages, which are more common in most other countries, exist in the form of annuity and linear mortgages.⁹

Mortgages in the Netherlands are, contrary to what is typical in the United States, with recourse so that consumers are personally liable for any outstanding mortgage debt in case of default. However, mortgages smaller than the average national house price are typically insured by the government through the so-called national mortgage guarantee (Nationale Hypotheek Garantie, or *NHG*, in Dutch). The NHG pays off the remaining balance if a household defaults on its mortgage because of divorce, disability or unemployment.¹⁰ Enrollment in the NHG costs 1% of the loan sum. Banks view mortgages with NHG as low risk and offer significant interest rate discounts if consumers choose to enroll.¹¹

Households tend to fix their interest rate for relatively long periods, most commonly for ten years (Table 2). However, this fixed interest rate period is shorter than the typical duration of a mortgage, which is thirty years. When the fixed interest rate period ends but the principal is not yet due, a household's current bank offers it to renew its mortgage. At this point, however, it is also possible to switch to a different bank.¹² However, switching is costly. Switching costs typically include the costs of a notary, the cost of appraisal, and, if the mortgage qualifies for NHG insurance, insurance fees. The country's largest mortgage broker estimates that these costs are around €3500 for an average household, accounting for the fact that these costs are tax-deductible.¹³

History-based price discrimination, of which I give evidence below, existed in two forms. Some banks would explicitly have a different interest rate for renewing customers. Other banks, however, applied price discrimination in more implicit ways, through discounts that in practice were only

8. Statistics Netherlands. <http://statline.cbs.nl/StatWeb/publication/?VW=T&DM=SLNL&PA=71446ned>. Accessed January 31, 2017.

9. An annuity mortgage features constant mortgage payments. A linear mortgage, on the other hand, features constant amortization, so that mortgage payments are decreasing over time.

10. NHG. <https://www.nhg.nl/Consument/Wat-is-NHG>. Accessed January 31, 2017.

11. A typical discount is between .4 and .7 percentage points (Fransman 2017).

12. Switching to a different bank is also possible when the fixed interest rate period is still ongoing. However, a household incurs severe penalties when it ends the fixed interest rate period prematurely.

13. See footnote 1 for the source and further explanation.

available to new customers.¹⁴ One common example of the latter is a discount for first-time buyers: this discount typically was no longer available upon renewal.

3.1 Ban on history-based price discrimination

The origin of the Dutch ban on history-based price discrimination lies in the observation that after a large initial drop during the 2008 financial crisis, mortgage interest rates in the Netherlands quickly increased again from 2009 onwards, while they stayed low in other European countries (Dijkstra, Randag, and Schinkel 2014b). Various possible reasons have been given for this, from collusion to the large reliance of Dutch banks on external funding, the cost of which increased sharply after the financial crisis.¹⁵ After an investigation, the Dutch competition authorities saw no reason to intervene, but they did provide some recommendations to make the Dutch mortgage market more competitive (Nederlandse Mededingingsautoriteit 2011). One recommendation, which the government followed, was to ban history-based price discrimination, as the competition authorities believed the differences in interest rates between prolonging and first-time customers to be anti-competitive. As a result, the ban on history-based price discrimination came into effect on January 1, 2013.

The regulation states that institutions are legally obliged to offer the same interest rate to households at the end of their fixed interest rate period as to households who are first-time customers at that institution, if those two households have a similar risk profile. As the responsible regulator AFM later clarified, this does not mean that banks are barred from all types of price discrimination (Autoriteit Financiële Markten 2015). For example, they are allowed to (and in practice do) offer a discount if a household also has a deposit account at the same institution. However, such discounts have to be equally available to existing and first-time clients.

4 Data

This section describes my data. I begin by introducing my data set. Then I explain how I construct my sample and give some descriptive statistics. Finally, I show reduced-form evidence of history-based price discrimination in the Dutch mortgage market.

4.1 Data sources

My main data source is the Loan-Level Data (LLD) from the Dutch Central Bank (DNB). The LLD are a yearly panel containing micro-level data covering almost the complete Dutch mortgage market. Starting in 2013, participating banks hand in a yearly report of all their outstanding mortgages. The LLD contain detailed information on the loans, as well as some information on the

14. Dutch Competition Authority (in Dutch), pp. 29-30. "Sectorstudie Hypotheekmarkt." www.acm.nl/sites/default/files/old_publication/bijlagen/7091_Sectorstudie%20Hypotheekmarkt.pdf. Accessed September 19, 2018.

15. See Dijkstra, Randag, and Schinkel (2014a) and the other articles in the same issue of the *Journal of Competition Law and Economics* for an overview of the arguments.

underlying property and the household that purchased it. Not all banks are required to submit information to the LLD: small banks and some foreign banks do not have to report information. The LLD cover between 75% and 80% of the full market and aggregate statistics match those from other data sources (Mastrogiacomo and Van der Molen 2015). The banks in my sample are ABN Amro, Florius, ING, Obvion and Rabobank, as well as some fringe players which I aggregate into a single bank. The only bank with a significant market share that is missing is SNS Reaal, which does not provide information on whether or when a mortgage is renewed.¹⁶ To preserve banks' anonymity—a pre-condition for accessing the LLD—I will anonymize bank names for the remainder of this paper.¹⁷

This study is based on the LLD from 2013, 2014 and 2015, so that I have three years of observations. However, since I observe the *stock* of all mortgages in 2013 and most mortgages have a fixed interest rate period of at least five years, I can see virtually all purchases of the directly preceding years.

In addition to the LLD, I use the DNB Household Survey (DHS). The DHS is a yearly survey of a random sample of approximately 2000 Dutch households on their finances. I use the DHS because the LLD are a choice-based sample: it does not provide any information on households that purchase no mortgage. To say something about the extensive margin of the market, I add the demographic information on those households from the DHS. Since the LLD are a choice-based sample and the DHS is a random sample, I adjust for different sampling probabilities in the demand estimation.

4.2 Inferring switching

Because of the panel dimension of my data, I can observe a household's previous mortgage when it purchases a new mortgage, allowing me to observe whether it renewed its current mortgage or switched to a different bank. However, the panel dimension of my data is limited by the fact that banks use different schemes to encode household id's. This means that when a household switches to another bank, I cannot observe to which bank exactly. To address this issue, I probabilistically match switching households based on the birth year of the primary borrower, the type of mortgage (e.g. bullet or annuity), the outstanding balance of the mortgage at the moment of switching and the maturity year of the loan. Further details can be found in Appendix B.1.

Since the first wave of the LLD is 2013, I do not have a "previous" observation for households that make a purchasing decision before 2013. I *can* observe whether or not the mortgage is renewed or not as there is a field that indicates this. Therefore, before 2013, I know the previous bank of renewing households—as it equals their current bank—but not of switching households or households that purchase their first mortgage. I deal with this by estimating the model on the part of the data for which I can identify switching behavior well, i.e. for the post-2013 data, as I explain

16. SNS Reaal has a market share of approximately 8%. This is significantly smaller than the market shares of ABN Amro, ING and Rabobank, who all have a market share of more than 20%.

17. For this reason, I am also unable to provide market shares of the anonymized banks, because the market shares allow identification of the banks based on publicly available data.

Table 1: Comparison of mortgages with and without government insurance (NHG).

	NHG	No NHG
Age household head (years)	38.0	49.0
Household income (€)	44730	70537
Initial loan (€)	141668	113502
Property valuation (€)	188254	330344
Observations	436608	229621

Note: The table compares average household and loan characteristics for loans with and without NHG insurance, for mortgages with a fixed interest rate period starting in the period 2010-2015.

in more detail below.

4.3 Sample selection

My empirical strategy depends on comparing the interest rate that new and renewing households pay for the same mortgage. One reason other than purchasing history for such interest rate differences may be differences in risk. To rule out as much as possible that interest rate differences are caused by differences in risk, I focus on mortgages with NHG insurance. As explained in Section 3, such mortgages are insured by the government so that from the perspective of banks they are more or less risk free. However, mortgages are only eligible for NHG insurance when the purchase amount is at most €245,000.¹⁸ As a result, households that have a mortgage with NHG differ systematically from those that do not. Table 1 shows that on average, households with NHG are 11.5 years younger, their household income is €15,000 lower and they own properties valued €150,000 less than households without NHG. However, households with NHG mortgages have larger loans because households who do not enroll into the NHG scheme tend to make larger down payments. My results should thus be understood to be applicable to this segment of the overall mortgage market.

I further restrict my sample by only considering mortgages with a fixed interest rate period of around five, ten, fifteen or twenty years. If a mortgage has a fixed interest rate period within six months of any of these periods, I round off the duration to that length. For example, I round a fixed interest rate period of 57 months to six years. The reason for not considering other fixed interest rate periods is that they all have very small market shares.

In my data there remain interest rate differences between households purchasing the same type of loan in the same month even after controlling for history-based price discrimination. These differences are caused by the fact that I observe the date on which the loan deed was signed, which is different from the date the interest rate was offered: e.g. one household who signed their mortgage in March might have received their interest rate offer in January and another their (different) offer in February. A second reason for these differences is the existence of other types of price discrimination which I do not observe. Most banks for example offer a small discount

18. In 2015. The maximum purchase amount was somewhat higher in previous years, up to €276,190 in 2012.

Table 2: Distribution of fixed interest period durations.

	2010	2011	2012	2013	2014	2015
5 years	11.29	22.16	21.82	27.15	21.07	8.46
10 years	81.12	69.14	66.96	67.71	68.73	65.87
15 years	1.96	3.12	3.60	2.60	3.50	7.18
20 years	5.63	5.58	7.63	2.54	6.70	18.49
Observations	73923	69363	64721	65843	78240	84809

Note: The table shows the proportion (%) of fixed interest rate durations by start year of the fixed interest rate period. Fixed interest rate durations within six months of one the stated durations are counted towards the market share of that duration. Mortgages with other fixed interest rate durations are discarded.

on the interest rate if a household opens a deposit account at the same bank. Since I do not have information on the source of this residual price dispersion, I take as a mortgage’s interest rate its modal interest rate in a given month, and assume that every household paid this interest rate. Some mortgages have zero sales in certain months. For those loans I impute the interest rate based on a simple linear regression.¹⁹

4.4 Descriptive statistics

Table 2 shows the market shares of the fixed interest rate durations of the loans in my sample. A fixed interest rate period of ten years is by far the most common, with a market share between 67% and 82%. A five year period is the second most common, with relatively small market shares for fifteen and twenty years. In 2015 the market share of twenty year fixed interest rate periods increases dramatically, probably because consumers wanted to “lock in” the historically low interest rates in 2015.

Table 3 describes the main characteristics of loans in my sample. The average interest rate decreases from 2010 (4.66%) until 2015 (2.77%). The average fixed interest rate period is approximately ten years, comparable to the mode. The proportion of loans that is renewed rather than a new purchase increases dramatically after the ban on history-based price discrimination in 2013, from around 10% before to between 23% and 40% after. This means that households switch less often to a different bank after their fixed interest rate period ends. The average switching rate in my sample is 6%.

Table 4 shows that before 2013, non-amortizing loans such as bullets and savings mortgages were by far the most popular. This is because, before 2013, the interest rate on mortgages was fully deductible, so that non-amortizing mortgages were very attractive. From 2013 onwards, new non-amortizing mortgages no longer qualify for interest rate deductibility. Therefore, the market share of non-amortizing loans becomes much smaller. Since this change in taxes happens at the same time as the ban on history-based price discrimination, it is a possible confounder. In the

19. I run separate regressions for the interest rates for new and renewig customers. I include bank, loan type, fixed interest rate period and month fixed effects. The regressions have an R^2 of .827 (new customer interest rate) and .886 (renewing customer interest rate).

Table 3: Loan characteristics.

	2010	2011	2012	2013	2014	2015
Interest rate	4.66 (0.45)	4.66 (0.49)	4.49 (0.51)	4.11 (0.55)	3.47 (0.52)	2.77 (0.44)
Current balance (€)	114616 (51666)	116587 (52965)	114927 (50929)	114228 (52504)	116261 (54354)	118380 (51275)
Fixed interest rate period (months)	121 (36)	115 (41)	118 (45)	108 (36)	117 (43)	141 (53)
Renewed mortgage	0.099 (0.299)	0.132 (0.339)	0.106 (0.308)	0.233 (0.423)	0.350 (0.477)	0.398 (0.489)
Observations	73923	69363	64721	65843	78240	84809

Note: The table shows average characteristics of loans with NHG insurance by start year of the fixed interest rate period. Standard deviations are in parentheses.

Table 4: Distribution of payment methods.

	2010	2011	2012	2013	2014	2015
Annuity	2.58	3.29	5.96	33.40	48.41	43.81
Linear	0.37	0.45	0.68	2.34	4.25	4.47
Bullet	20.85	18.32	18.20	16.12	16.58	22.18
Savings mortgage	66.80	68.00	68.48	34.30	13.77	14.12
Life mortgage	7.41	8.15	5.32	11.70	14.36	12.39
Investment mortgage	1.99	1.78	1.36	2.14	2.64	3.03
Observations	73923	69363	64721	65843	78240	84809

Note: The table shows the distribution of payment types of loans with NHG insurance by start year of the fixed interest rate period.

demand model, the tax change amounts to a change in households' choice set for which I can control.²⁰ It also also possible that the tax change causes a change in banks' pricing. In Section 8.1 I show evidence that this is not the case.

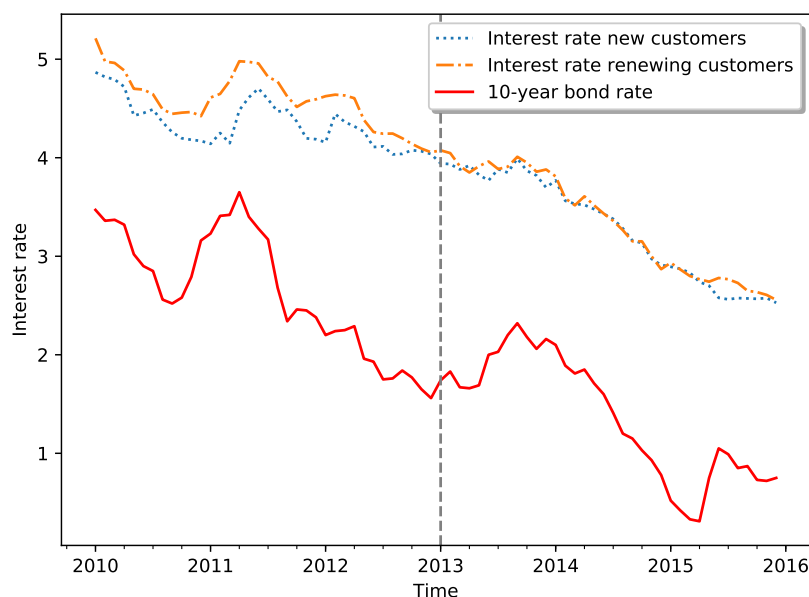
4.5 Evidence of history-based price discrimination

In this section, I present reduced-form evidence on the extent of interest rate differences between new and existing customers and the effect of the ban on history-based price discrimination. Figure 1 plots the average interest rates paid by renewing and new customers. Before the ban, renewing customers clearly pay higher interest rates. These differences largely disappear in 2013, after the ban.

To get at the magnitude of price discrimination, I regress the interest rate consumers pay on month by loan fixed effects and whether the loan is renewed or not. Thus, I compare the interest

²⁰ The reason the tax change can be interpreted as a change in choice sets is because the tax change made non-amortizing mortgages so unattractive that in practice banks do not even offer them any more to first-time buyers.

Figure 1: Interest rates for new customers and renewers.



Note: The figure shows the average interest rates for households with a NHG mortgage. Both interest rates are weighed by the total market share, i.e. the sum of the market shares across new and renewing customers. The interest rate on 10-year Dutch government bonds is included for comparison.

paid by renewing and new customers for the same loan in the same month.²¹ Table 5 shows that, before the 2013 ban, renewing consumers paid statistically higher interest rates than new consumers. This difference is also economically significant. Since the average mortgage's starting balance is around €150,000, the differences in interest rates imply that renewing consumers paid between €228 (in 2012) and €348 (in 2011) after tax in yearly interest payments more for the same type of loan as new consumers. These estimates conform to previous estimates made by the AFM, the regulatory agency charged with upholding the ban, in the run-up to the 2013 ban on history-based price discrimination.²²

After the 2013 ban on history-based price discrimination, the interest rate difference between renewing and new consumers drops but remains statistically significant (Table 5). In economic terms, a renewing customer with an average mortgage of €150,000 pays between €70 (in 2014) and €167 (in 2015) more for the same type of loan than a new customer. Thus, the ban on history-based price discrimination significantly reduced the interest rate differences between renewing and new households, but did not eliminate them completely. The large increase in interest rates differences in 2015 is entirely due to one of the larger banks having differences between renewing and new customers comparable to pre-ban levels.²³ This finding is consistent with statements made by the AFM that there are some loopholes banks use to partially get around the ban.²⁴ Since, in particular

21. In Figure 1, I do not control for loan fixed effects. Therefore, differences in Figure 1 might be caused by renewing and new customers purchasing different mortgages.

22. See for example the newspaper article "Straf voor verlengen hypotheek moet van tafel", *Het Financieele Dagblad*, 25 september 2010. This article reports pre-ban interest rate differences between 0.2 and 0.4 percentage points.

23. That is, removing this bank from the sample gives differences comparable to 2013 and 2014.

24. For example, many banks offer a discount if the mortgage contract is certified by a notary within a certain time frame.

Table 5: Average interest rate difference between renewing and new customers.

	2010	2011	2012	2013	2014	2015
Renewed mortgage	0.328*** (0.0167)	0.401*** (0.0190)	0.263*** (0.0140)	0.0925*** (0.00886)	0.0806*** (0.00829)	0.193*** (0.0159)
Constant	4.631*** (0.00166)	4.604*** (0.00251)	4.461*** (0.00149)	4.089*** (0.00207)	3.438*** (0.00290)	2.698*** (0.00635)
Loan \times month fixed effects	✓	✓	✓	✓	✓	✓
Observations	73923	69363	64721	65843	78240	84809

Note: The table shows the average difference in interest rates between renewing and new customers for the same loan. A loan is defined by the originating bank, the length of the fixed interest rate duration and the payment method. Standard errors are in parentheses. *** $p < 0.001$.

for the years 2013 and 2014, the interest rate differences between renewing and new households are much smaller after than before the ban, I will assume that banks charge a single interest rate after the ban for simplicity.

5 A model of mortgage demand and supply

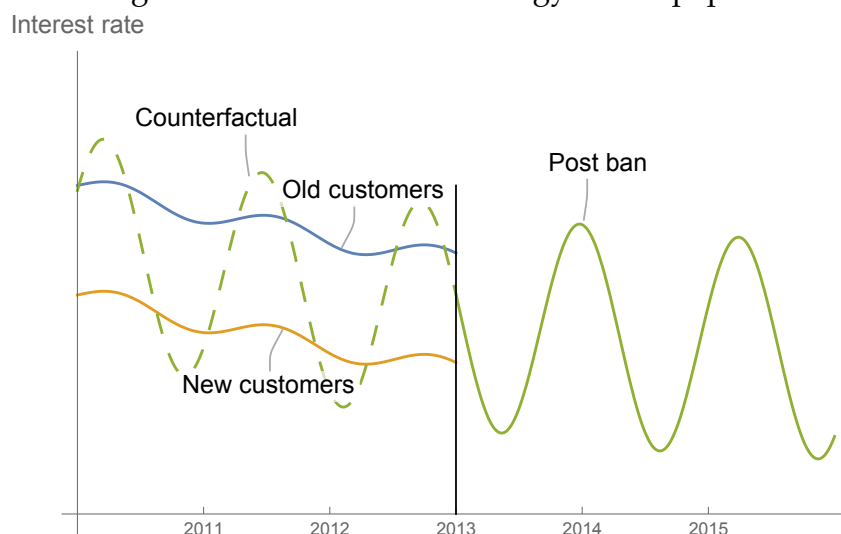
In this section, I develop a structural model of the Dutch mortgage market. I begin by describing the empirical strategy I use to identify the effects of history-based price discrimination. Then I describe the demand and supply side of the model. I conclude this section by introducing the new solution concept I use for the supply side, Sparse Markov Perfect Equilibrium.

5.1 Empirical strategy

To estimate the effect of the ban on history-based price discrimination, I use the empirical strategy depicted in Figure 2. I estimate a structural model of demand and supply. Since, as I explain in Section 4.2, I am better able to follow switching households over time after 2012 and switching behavior is a crucial element of my model, I estimate demand and supply on the data covering the *post-ban* period (2013-2015). Given estimates of banks' policy functions after the ban, I can predict what interest rates they would have counterfactually set *pre-ban* had history-based price discrimination already been banned then. Comparing these counterfactual interest rates with observed interest rates gives an estimate of the effect of the ban on interest rates. I then use the estimated demand and supply model to calculate the effect of the ban on consumer surplus and bank profits, respectively. This strategy estimates a causal effect to the extent that I am able to control for all relevant changes in the mortgage market around the time of the ban. Given that the market underwent many changes around the ban, this is a strong assumption. I describe the most

Since prolonging a mortgage does not require certification by a notary, such discounts are automatically not available to prolonging households (Autoriteit Financiële Markten 2015). Such discounts are now banned, but were still sometimes offered during my sample.

Figure 2: The estimation strategy of this paper.



Note: The model is estimated for the post-ban period. The reduced-form policy function estimates are then extrapolated to the pre-ban period. Comparing with observed pre-ban interest rates then gives the effect of the ban.

important of these changes in Section 8.1. Here I also argue that these changes cumulatively did not have a large impact on mortgage pricing.

This approach differs from the more typical approach in empirical industrial organization where the estimated model is used to compute a counterfactual equilibrium. Because I observe both equilibria of interest (with and with history-based price discrimination), I can instead extrapolate banks' post-ban policy functions to the pre-ban period and compare with observed interest rates. This method has three advantages over computing a counterfactual. The first is that it is much simpler. Since the supply side of my model is a dynamic game, computing a counterfactual equilibrium is computationally challenging. Moreover, dynamic games typically have many equilibria. Thus, computing a counterfactual equilibrium raises the issue of equilibrium selection. Here, I can simply assume that banks keep playing the observed equilibrium. Finally, this method is arguably more credible since it does not require me to make assumptions on bank conduct when history-based price discrimination is legal.

5.2 Demand

I estimate a standard discrete-choice model of demand.²⁵ Households choose which mortgage to purchase, I take the desired loan amount as given.²⁶ The relevant market in month t consists of

25. In practice, many households combine different loans, for example a savings mortgage with an annuity, into a single mortgage. I ignore this because the second mortgage almost always either has the same repayment method (but has for example a different maturity) or is an annuity. Modeling the choice of the second (and potentially third) loan would impose significant difficulties, since it would require a model for the decision how much of the total loan sum to allocate to what type of loan in addition to a model of loan choice. However, there is little variation in the type of second loan so that the gains do not outweigh the benefits.

26. For renewing households, the desired loan amount *is* in fact given since they need to refinance their outstanding debt. Loan-to-value ratios for NHG loans are on average around 95%, which indicates that most first-time borrowers

households whose mortgages' fixed interest rate periods end, as well as households who move in that period and may purchase a mortgage if they buy rather than rent.²⁷ Denote all households that make a purchasing decision at time t by \mathcal{H}_t . A household chooses from up to J possible mortgages in the choice set \mathcal{C}_i . The choice set differs between households because, from 2013 onwards, a household can only purchase a non-amortizing mortgage if it previously had a non-amortizing mortgage (as I explain in Section 4.4).

Banks sell multiple mortgages. Denote by \mathcal{J}_b the set of mortgages sold by bank b . I denote by $j = 0$ the outside option, which is sold by "bank" $b = 0$. The outside option includes not having a mortgage, as well as having a mortgage by an institution which is not in my sample.²⁸

The utility household i derives from mortgage j depends on the mortgage's characteristics X_{jt} , its interest rate r_{jtn} , household income D_{it} and potentially switching costs s_{it} . I subscript the interest rate with $n \in \{0, 1\}$ to indicate the interest rates banks charge pre-ban to old and new customers, respectively. New customers include switching households and households that purchase their first mortgage. A mortgage is defined by the bank that sells it, the fixed interest rate duration (as in Table 2) and the payment method (as in Table 4). In other words, X_{it} contains bank, fixed interest rate duration, and payment method dummies.

Consumers face switching costs. A household incurs switching costs if it purchases a mortgage from a different bank—switching to a different type of mortgage by the same bank is free. This is a simplification since switching to a different type of mortgage will typically involve some costs. However, the main costs of switching to a different a bank—taxation, a notary and NHG insurance fees—do not need to be paid when staying at the same bank. Switching costs measure the monetary, hassle and time costs of switching. In addition, the switching cost parameter will pick up other frictions in the market as well, such as search costs and inattention. Disentangling between these sources of consumer inertia requires either variation in the magnitude of the different frictions or direct data on consumer search, both of which are unavailable in my setting.²⁹

The utility household i derives from purchasing mortgage j is

$$u_{ijt} = X_{jt}\Pi_{Dit} - (\alpha_{ijt}r_{jt1} + s_{it})\Delta_{ijt} - \alpha_{ijt}r_{jt0}(1 - \Delta_{ijt}) + \zeta_{jt} + \varepsilon_{ijt}. \quad (1)$$

where Δ_{ijt} is a dummy denoting whether household i needs to pay switching costs to purchase mortgage j . Π_{Dit} and α_{ijt} are coefficients and ζ_{jt} and ε_{ijt} are error terms. The first term gives the

borrow close to as much as they can. (The maximum loan-to-value ratio is 102%). Therefore, the effect of the ban on loan amounts is not particularly interesting.

27. It is possible that a household ends its fixed interest rate period prematurely. This is costly, but it can sometimes be worthwhile if interest rates have dropped sufficiently since the start of the fixed interest rate period. However, ending a fixed interest rate period prematurely does not imply the need to switch: most households end up renewing the loan at their current bank. Therefore, I take the choice to end a fixed interest rate period as exogenous and model the mortgage choice of households that do so the same as of households whose fixed interest rate period ends contractually.

28. In the remainder of this paper, I will often refer to households who have previously purchased a mortgage from some bank b . Such statements should always be understood to be valid for $b = 0$ as well, i.e. such statements also apply to households who previously purchased the outside option.

29. See Kiss (2017) for more on disentangling attention and switching costs and Honka (2014) on disentangling search and switching costs.

utility the household derives from the mortgage characteristics X_{jt} . The second term contains the interest rate paid by new customers and switching costs s_{it} . Pre-ban, renewing customers pay a different interest and they never incur switching costs, which is reflected by the third term. The fourth term, ζ_{jt} , denotes the unobserved quality of mortgage j . As is usual in the empirical industrial organization literature, ζ_{jt} is allowed to be correlated with the interest rate. Since ζ_{jt} is unobserved, interest rates are endogenous, for which I will correct in the estimation of the demand model. The final term, ε_{ijt} , is an idiosyncratic error term which is assumed to be conditionally independent from all other variables in the utility specification.

The outside option gives utility

$$u_{i0t} = -s_{it}\Delta_{i0t} + \varepsilon_{i0t}.$$

In this specification, switching to and from the outside option is costly. The assumption that switching *from* the outside option is costly reflects that the same fees need to be paid when purchasing a new mortgage as when switching. Switching *to* the outside option is also costly as paying off a mortgage before its end date typically triggers a fine.

To make the demand model as flexible as possible, the model includes household-specific coefficients. The coefficients on mortgage characteristics, the coefficient on the interest rate and switching costs depend on household income D_{it} as follows:

$$\begin{aligned}\Pi_{Dit} &= \Pi_D D_{it}, \\ \alpha_{ijt} &= \exp\{\Pi_\alpha D_{it}\}, \\ s_{it} &= \exp\{\Pi_s D_{it}\}.\end{aligned}$$

The preceding formulation assumes that demand is static, despite the fact that households face an inherently dynamic problem: when fixing their interest rate, they have to form expectations about the interest rates they will face when they have to renew their interest rate. One reason I estimate a static demand model is precisely because I believe it is too complicated for most households to think ahead ten years and realize they might have to pay more then, let alone form correct expectations about the strategic behavior of banks in the future. In fact, survey evidence indicates that 60% of households do not even *consider* switching at the moment of renewal.³⁰ Therefore, it seems unreasonable to assume a significant fraction pays attention to dynamic considerations. A second reason for estimating a static demand function is that virtually all households in my sample make only a single decision during my sample period, so that any dynamic model of demand would be identified only by functional form.

30. The survey was conducted by the Dutch homeowners' association. Vereniging Eigen Huis. "Huisseigenaren laten massaal geld liggen bij nieuwe rente hypotheek", January 3 2018. <https://www.eigenhuis.nl/actueel/nieuws/2018/01/03/09/00/huisseigenaren-laten-massaal-geld-liggen-bij-nieuwe-rente-hypotheek>

5.3 Supply

Because the interest rate a bank sets today affects the number of captive consumers it has tomorrow, banks face a dynamic problem when setting their interest rates. I use a non-standard equilibrium concept, Sparse Markov Perfect Equilibrium, to solve the dynamic game between the banks. Before I introduce it below, I describe the constituent parts of the model: the state, how it evolves and banks' flow profit functions.

The state consists of two parts: cost shocks and previous sales. There is no private information: the complete state is known to all banks. At the beginning of every month, there is a shock to the common cost of funding i_t . In the Netherlands, mortgage funding comes from many different sources, including short-term and long-term deposits, money market funds and securitization. I use the interest rate on 10-year Dutch government bonds as a proxy for these costs, as this interest rate displays significant co-movement with observed mortgage interest rates (as can be seen in Figure 1).³¹

The second state variable consists of previous market shares, which banks need to take into account because households face switching costs. However, because of the heterogeneity in the demand function, banks not only need to take into account *the amount* of mortgages they sold in the past, but also *the type of consumers* they sold them to. In terms of the model, this is the case because the coefficients in households' utility, Π_{Dit} and α_{it} , as well as switching costs s_{it} depend on household income D_{it} . Thus, there is in effect a different demand function for every type of household. Because of switching costs, the demand of a particular type of household depends on past market shares for this type. Therefore, total demand depends on past market shares of every type of household. In other words, the pay-off relevant state variable is the joint density of previous sales and household income that are in the market at time t . Denote this density by $f_t(b, D)$, where b is a random variable denoting which bank a household in \mathcal{H}_t purchased from previously and D is the distribution of household income. Denote the set of all state variables at time t by $\sigma_t = \{f_t(\cdot), i_t\}$.

Let r be the vector of all interest rates. The demand for mortgage j from consumers who previously bought from bank b

$$d_j(b, r) = M_t \int p_j(b, r, D) f(b, D) dD,$$

where p_j is the probability implied by the demand model that a household of type D whose current bank is b purchases mortgage j when interest rates equal r and M_t is the total demand for loans (in euros) at time t .

Denote by c_j the marginal cost of supplying a loan of €1. I let

$$c_{jt} = \gamma_{j0} + \gamma_{j1} + \gamma_{j2} i_t.$$

31. I have also tried using the marginal rate on deposits as a proxy for the cost of funding. However, this is a worse proxy: when I include both the rate on Dutch government bonds and the marginal rate on deposits in the policy function regression in Section 6.1.2, the Lasso always selects only the interest rate of Dutch government bonds.

γ_{j0} measures any marginal cost of supplying loan j in addition to the cost of funding, for example the implied costs of pre-payment risk (which differ across loan types). Some loan types, such as savings or investment mortgages, are commonly sold together with other high-margin products such as life insurance. γ_{j0} also measures this implicit cross-subsidy. I restrict γ_{j0} to be the same for all loans with the same payment method sold by the same bank, e.g. all annuities sold by bank b have the same γ_{j0} . γ_{j1} and γ_{j2} measure the cost of funding. Similarly, I restrict γ_{j1} and γ_{j2} to be the same for all loans with the same fixed interest rate duration sold by the same bank, e.g. all 10-year loans sold by bank b have the same γ_{j1} and γ_{j2} . Thus γ_{j1} and γ_{j2} measure the differences in cost of funding of loans with different fixed interest periods as a function of the average cost of funding.

Denote by r_b and r_{-b} the interest rates set by bank b and its competitors, respectively. The flow profits of bank b in state σ are

$$\pi_b(\sigma, r) = \sum_{j \in \mathcal{J}_b} \left\{ (r_{j0} - c_j) d_j(b, r_b, r_{-b}) + \sum_{b' \neq b} (r_{j1} - c_j) d_j(b', r_b, r_{-b}) \right\}.$$

The first term contains the profits from customers who do not switch, the second term from customers who do.

Given the state and banks' interest rates, the future state can be calculated as follows. At time t , the proportion of households with mortgage j whose mortgage will expire in τ months is equal to the weighed sum of the proportion at time $t - 1$ of households with that mortgage expiring in $\tau + 1$ months and the market share at time $t - 1$ of mortgage j among mortgages with a duration of τ months. The weights are the number of mortgages expiring in $\tau + 1$ months and the number of expired mortgages at time $t - 1$, respectively. The evolution of the joint density of previous purchases and household characteristics, evaluated at a point (j, D) , can be written as

$$f_{t+\tau}(b, D) = \frac{|\mathcal{H}_{t+\tau}|}{|\mathcal{H}_{t+\tau}| + |\mathcal{H}_t|} f_{t+\tau}(b, D) + \frac{|\mathcal{H}_t|}{|\mathcal{H}_{t+\tau}| + |\mathcal{H}_t|} \sum_{j \in \mathcal{J}_b} \sum_{b'} \phi_j(\tau) p_j(b', r, D) f_t(b, D),$$

where $\phi_j(\tau)$ is an indicator function that equals 1 if and only if product j has a fixed interest duration of τ months. $|\mathcal{H}|$ denotes the number of elements in \mathcal{H} , i.e. $|\mathcal{H}_t|$ is the number of households that make a purchasing decision at time t . Combining this transition function for all points (j, D) and all $\tau = 1, \dots$, gives the full transition function for f :

$$f_{t+\tau} = \Gamma_t(s, r).$$

5.4 Sparse Markov Perfect Equilibrium

The typical solution concept for dynamic games is Markov Perfect Equilibrium (Maskin and Tirole 2001, MPE hereafter). An MPE is a sub-game perfect equilibrium in which agents' strategies are constrained to be functions of only the payoff-relevant state.³² That is, an MPE consists of policy

32. That is, agents' strategies cannot be functions of play in previous' periods, except insofar that play changes the current period's state.

functions $\rho_b(\sigma)$ such that

$$\rho_b(\sigma) = \arg \max_{\rho} \pi_b(\sigma, r, \rho_{-b}(\sigma)) + \beta \mathbb{E}[V_b(\Gamma(\sigma, r))]$$

for all possible states σ . Here,

$$V_b(\sigma) = \pi_b(\sigma, \rho_b(\sigma), \rho_{-b}(\sigma)) + \beta \mathbb{E}[V_b(\Gamma(\sigma, r))]$$

is bank b 's Bellman equation and $\Gamma(\sigma, r)$ the evolution law.

In the supply model derived above, the payoff-relevant state is infinite-dimensional: banks have to keep track of the joint density of previous market shares and household demographics. This creates a challenge when estimating the model as existing methods (e.g. Bajari, Benkard, and Levin (2007)) require a first stage in which policy functions are regressed on state variables. When the number of state variables is larger than the number of observations of the policy function, such a regression is impossible. Since I have an infinite number of state variables, I cannot proceed as is standard.

In general, the pay-off relevant state will be large when the model, or parts of it, allows for rich heterogeneity. Typically, this kind of rich heterogeneity is required to match the patterns found in the data. The most common example of this is in demand estimation, where it is well known that the simple logit model cannot capture typical demand elasticities.³³ Normally, the researcher then faces a trade-off: richer heterogeneity in one part of the model (for example in the demand function) increases the size of the state space in another (for example in the supply side). Here, I could reformulate the model in such a way that the state space is finite-dimensional, for example by having only two types of households, rich and poor, in the demand function. However, this would reduce the flexibility of the demand function a lot.³⁴ Because it is well-known that the shape of the demand function is a crucial determinant of the welfare effects of third-degree price discrimination (Holmes 1989), it is important that I allow for sufficient flexibility of the demand function. By introducing techniques to deal with large state spaces in the estimation, I hope to relieve this tension: it is possible to have rich heterogeneity in one part of the model without creating insurmountable challenges when estimating the part of the model with the dynamic game.

To solve the challenge of a large state space, I combine techniques from machine learning with micro-economic theory. The theoretical part consists of a new solution concept that allows easier estimation of games with large state spaces: Sparse Markov Perfect Equilibrium (SMPE). As I explain in the section on estimation below, SMPE is particularly amenable to estimation using machine learning techniques. The main difference with a standard Markov Perfect Equilibrium is that I relax the assumption that information on the state is free. In an SMPE, agents then optimally pay attention to a subset of the state. As a result, the domains of their policy functions have a

33. See, for example, Akerberg et al. (2007) for a discussion of this issue.

34. Moreover, to obtain flexible estimates of policy functions, it is typically required to include functions of state variables, such as higher-order polynomials or interactins. Given the number of observations I have, this would already be impossible when the heterogeneity is restricted to two groups.

smaller dimension, making estimation and calculation much simpler. This concept is not only *computationally* attractive, but also *behaviorally*. Banks do not just know the state they are in, they need to perform some kind of market research. Such research is costly. Therefore, they will only try to figure out those state variables of which knowledge has a sufficiently large impact on their profits. For example, the MPE of the supply model implies that banks' policies are a function not just of past market shares, but also of market shares across households with an income of €30,000, of €31,000, and so on, as the full density of household demographics and past market shares is payoff-relevant. Most likely, it does not pay off for banks to invest in such detailed knowledge.³⁵ SMPE formalizes this intuition.

5.4.1 Sparse maximization

Before I introduce SMPE, I give a brief introduction to sparse maximization. Sparse maximization, introduced for a single decision maker by Gabaix (2014, for static settings) and Gabaix (2017, for dynamic settings), is a simple method to model inattention. The basic idea is that a decision maker pays attention to a variable if the benefits of paying attention exceed the costs. The sparse maximization operator then defines the benefits of paying attention in such a way that the model becomes tractable.

Sparse maximization works as follows. For notational simplicity, consider a single bank.³⁶ Denote by σ the vector of pay-off relevant state variables. The bank wants to maximize $v(\sigma, r)$, where r is the vector of interest rates the bank sets. In a static context, $v = \pi(\sigma, r)$, in a dynamic context $v = \pi(\sigma, r) + \beta \mathbb{E}[V(\Gamma(\sigma, r))]$, where $V(\cdot)$ is the continuation value. Let $m_i \in \{0, 1\}$ indicate whether the bank pays attention to state variable σ_i .³⁷ There is a fixed cost of paying attention of $\kappa \geq 0$ per state variable.³⁸ In Gabaix (2014, 2017), this cost is interpreted as a psychological cost. In my setting, however, a more natural interpretation is the real cost of obtaining or processing information on a state variable, for example due to cost of market research or revenue management systems.

Every period, the bank forms a *sparse state*, $\hat{\sigma}$. It knows the true value of every state variable it pays attention to. For variables it does not pay attention to, it assumes they equal some default value σ^d . (This default is specified by the researcher. Below I specify how I define the default state

35. Moreover, the manager in charge of setting interest rates may be cognitively limited and unable to process the full information, instead choosing to focus on the most important state variables. As shown by Gabaix (2014), such bounded rationality also leads to sparse maximization.

36. I cover the extension to games below.

37. Gabaix (2014, 2017) also considers values of m_i between 0 and 1. In this case, the bank would display partial attention. I do not consider this case, since it is not required to make the state space simpler. Moreover, partial inattention cannot directly be identified from the data using the methods I introduce in Section 6.1. However, once banks are restricted to display either full or no inattention, identification becomes particularly simple. This assumption is equivalent to a functional form assumption on the cost of attention, see also Footnote 38.

38. Gabaix (2014, 2017) also considers other cost functions. Many of these cost functions lead to partial inattention. Since the focus of this paper is not on the cost function (for example, I do not estimate the cost of attention), I pick the simplest cost function that leads to sparsity.

in this application.) Thus, the bank's sparse state is

$$\hat{\sigma}_i = \begin{cases} \sigma_i & \text{if } m_i = 1, \\ \sigma_i^d & \text{if } m_i = 0 \end{cases}$$

for every state variable σ_i . The bank chooses its interest rates to maximize its profits in the sparse states:

$$\rho(\hat{\sigma}) = \arg \max_r v(\hat{\sigma}, r) = \arg \max_r v(m\sigma + (1 - m)\sigma^d, r).$$

How to choose the attention vector m ? For a given attention vector m , the bank expects to lose

$$\mathbb{E}[v(\sigma, \rho(\hat{\sigma})) - v(\sigma, \rho(\sigma))] \quad (2)$$

from not paying attention. Here, $\rho(\sigma) = \arg \max_r v(\sigma, r)$ is the vector of optimal interest rates in the true state σ . Therefore, the bank should choose its attention vector m to minimize

$$\mathbb{E}[v(\rho(\sigma, \hat{\sigma})) - v(\sigma, \rho(\sigma))] + \kappa \sum_i m_i.$$

However, this problem is more complicated than maximizing $v(\cdot)$ under full attention. Therefore, it does not offer a simplification. The idea of sparse maximization is to replace this (very intractable) problem by a (more tractable) second order Taylor expansion around the default state σ^d . In the default state, the bank's profit maximizing interest rates are

$$r^d = \rho(\sigma^d) = \arg \max_r v(\sigma^d, r).$$

Gabaix (2014, 2017) calls r^d the "default action". The second order Taylor expansion of (2) with respect to m_i is³⁹

$$\Lambda_i \equiv -\frac{1}{2} \mathbb{E}[(\sigma_i - \sigma_i^d)^2] r_{\sigma_i}^T v_{rr} r_{\sigma_i}, \quad (3)$$

with

$$r_{\sigma_i} \equiv \left. \frac{\partial \rho(\hat{\sigma})}{\partial \hat{\sigma}_i} \right|_{\hat{\sigma}=\sigma^d},$$

$$v_{rr} \equiv \left. \frac{\partial^2 v(\sigma^d, r)}{\partial r^2} \right|_{r=r^d}.$$

The bank pays attention to σ_i if the cost of inattention Λ_i is greater than the cost of attention κ . The expression for Λ_i shows that the cost of inattention is larger when

1. the state variable σ_i shows more variation, or,
2. knowing the true value of σ_i has a greater impact on the optimal interest rates set by the bank.

39. The first term of the expansion equals zero because it is equivalent to the first order condition of the bank. The expectation follows from rewriting $\hat{\sigma}_i = m_i \sigma_i + (1 - m_i) \sigma_i^d$.

The sparse max operator is then defined as follows.⁴⁰

Definition 1 (Sparse max operator (Gabaix 2017)) *The sparse max, $\text{smax}_r v(r, \sigma)$, is defined by the following procedure.*

1. Choose the attention vector m^* :

$$m^* = \arg \min_m \sum_i (1 - m_i)^2 \Lambda_i + \kappa \sum_i m_i,$$

Given m^* , a bank forms a sparse state $\hat{\sigma} = m^* \sigma + (1 - m^*) \sigma^d$.

2. Choose the interest rates

$$r = \arg \max_r v(\hat{\sigma}, r).$$

5.4.2 SMPE: Definition

I now extend sparse maximization to dynamic games. An SMPE is an MPE, except using sparse maximization. In every state, every bank maximizes its profits as if it were in the corresponding sparse state, given its beliefs on the interest rates its competitors set.

I define the default state σ^d as follows. I assume that banks substitute the long-run average for state variables they do not pay attention to.⁴¹ Therefore, bank b 's sparse state can be defined as follows.

Definition 2 (Sparse state) *Bank b 's sparse state σ_b is*

$$\hat{\sigma}_{bi}(\sigma_i) = \begin{cases} \sigma_i & \text{if } m_{b,i} = 1, \\ \mathbb{E}[\sigma_i] & \text{if } m_{b,i} = 0, \end{cases}$$

for all state variables i , where $\mathbb{E}[\sigma_i]$ is the long-run average of state variable i and m_b is b 's attention vector.

Since different banks can pay attention to different variables and form different sparse states, the question arises what banks' beliefs should be on their competitors' actions. A natural assumption is that when a bank behaves as if the true state equals its sparse state, it assumes that its competitors do so as well. Moreover, beliefs have to be consistent in equilibrium.

Assumption 1 (Consistency of beliefs) *A bank's beliefs on its competitors actions are consistent with equilibrium and with the structure of its own sparse state. Bank b believes its competitors behave as if the true state is equal to b 's sparse state:*

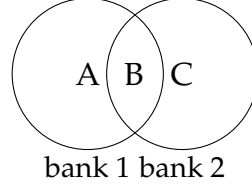
$$\hat{\rho}_{-b}(\sigma) = \rho_{-b}(\hat{\sigma}_b(\sigma)),$$

where $\hat{\rho}_{-b}(\cdot)$ is bank b 's belief on its competitors actions and $\rho_{-b}(\cdot)$ denotes b 's competitors' equilibrium interest rates.

40. To be precise, this is what (Gabaix 2017) calls the "sparse max operator without budget constraint". I do not require the version with budget constraint.

41. I rescale state variables that are past market shares so that the sum of market shares equals one.

Figure 3: Illustration of belief formation in SMPE.



Note: In this example, there are three variables in the true state—A, B and C. Bank 1 pays attention to A and B, bank 2 to B and C. Therefore, changes in C do not change bank 1’s belief on bank 2’s interest rates. Similarly, changes in A do not change bank 2’s belief on bank 1’s interest rates.

As an example of belief formation, consider the situation in Figure 3. Bank 1 pays attention to state variables A and B (and not C). Therefore it assumes that the state is equal to $(A, B, \mathbb{E}[C])$, i.e. it substitutes the long-run average of C for its true value. By Assumption 1, bank 1 believes that bank 2 behaves as if the state is $(A, B, \mathbb{E}[C])$ as well. To be consistent with equilibrium, Assumption 1 requires that bank 1 believes bank 2’s action in (A, B, C) is equal to its equilibrium action when the *true* state is $(A, B, \mathbb{E}[C])$.⁴²

To define an equilibrium, I require that for every possible state σ , banks’ policy functions are sparse maximizers of their discounted expected profits. Thus, in equilibrium bank b ’s Bellman equation is

$$V_b(\hat{\sigma}_b) = \operatorname{smax}_r \pi_b(r, \rho_{-b}(\hat{\sigma}_b), \hat{\sigma}_b) + \beta \mathbb{E} [V_b(\Gamma(\hat{\sigma}_b, r, \rho_{-b}(\hat{\sigma}_b)))], \quad (4)$$

where ρ_{-b} are the policy functions of b ’s competitors and β is the discount rate. Let $\hat{\Sigma}_b$ be the set of bank b ’s possible sparse states. An SPME is then defined as follows.

Definition 3 (Sparse Markov Perfect Equilibrium) *A Sparse Markov Perfect Equilibrium consists of policy functions $\rho_b : \hat{\Sigma}_b \rightarrow \mathbb{R}^{|\mathcal{J}_b|}$, such that*

$$\rho_b(\hat{\sigma}_b) = \arg \operatorname{smax}_r \pi_b(\hat{\sigma}_b, r, \rho_{-b}(\hat{\sigma}_b)) + \beta \mathbb{E} [V_b(\Gamma(\hat{\sigma}_b', r, \rho_{-b}(\hat{\sigma}_b')))],$$

for all banks $b \in \mathcal{B}$ and all states $\hat{\sigma}_b \in \hat{\Sigma}_b$, where $V_b(\cdot)$ is given by (4).

To formally investigate SMPE’s theoretical properties is outside the scope of this paper, where the focus is on the empirical application.⁴³ However, it seems likely that the usual multiplicity problem of MPE’s is even worse for SMPE’s: not only can there be many equilibria for given attention vectors, but it is also possible that there are equilibria which differ with respect to the state variables banks pay attention to. Since I assume that banks keep playing the estimated equilibrium, this is not an issue in this application. A second question is whether SMPE can be thought of as

42. Note that bank 2’s sparse state does not include A and that therefore, the requirement that A is at the true value is superfluous: holding B and C constant, changes in A do not change bank 2’s action. Thus, a state variable changes the belief of one bank on another bank’s action only if that variable is in both banks’ sparse states. In this case, since B is the only state variable that is in both banks’ sparse states, only changes in B change banks’ beliefs. However, changes in A do change bank 1’s behavior (since A is in its sparse state) and changes in C change bank 2’s behavior.

43. One hurdle in deriving the theoretical properties is that banks do not have rational expectations. As stressed by Gabaix (2017), this creates methodological challenges because backwards induction cannot be used.

an approximation to MPE. There is some reason to believe this might be the case: Gabaix (2017) shows that for finite-horizon single agent problems, the policy and value functions under sparse maximization differ only in second terms from the policy and value functions under rationality. I have not investigated whether this result extends to infinite-horizon games, but this single player finite-horizon result at least gives some indication that SMPE might be an approximation to MPE.

6 Identification and Estimation

In this section, I discuss the identification and estimation of my structural model. I start with the main methodological contribution of this paper: the identification and estimation of inattention. Then, I discuss the estimation of the demand and supply side of my model.

6.1 Sparse states and policy functions

6.1.1 Identification of sparse states

The intuition behind the identification of sparse states is simple: a bank pays attention to a state variable if and only if its policy function is a function of that state variable.⁴⁴ This is true since the sparse max operator picks precisely those state variables to pay attention to that have the largest impact on a bank's optimal action. Conversely, if one observes that a bank's interest rates do not vary with a certain state variable, it must not have paid attention to that variable: there is no point in a bank paying attention to a state variable if it knowing the value does not have an impact on its interest rates. Indeed, the cost of inattention Λ_i in (3) is increasing in r_{σ_i} , the derivative of a bank's optimal interest rates with respect to the state, so that a bank pays attention to a state variable only when knowing its true value sufficiently changes its actions.

Another way to look at this is by observing that the sparse max operator in effect truncates the policy functions. Gabaix (2014) shows that if policy functions are linear when agents pay full attention, with the coefficient on σ_i equal to a_i , the coefficient on σ_i under sparse maximization is simply⁴⁵

$$a_i \mathcal{I} \left(a_i^2 \geq 2 \frac{\kappa^2}{\text{Var}(\sigma_i)} \right).$$

Here, $\mathcal{I}(\cdot)$ is the indicator function. Therefore, a bank pays attention to a state variable σ_i if the policy function when attention is free varies "enough" with σ_i .

Since policy functions are observed by the econometrician, all that is required to identify sparsity is to select those state variables which have the largest effect on a bank's policy function. This is precisely what variable selection methods from the machine learning literature do: estimating a banks' policy function with a suitable penalized regression method is sufficient to identify its sparse state.

44. For legibility, I write paying attention to a state variable. It should be understood that by this I also mean any function of state variables, such as particular moments of the previous market share distribution f .

45. The same holds for the first-order Taylor approximation.

6.1.2 Estimation of sparse states and policy functions

To estimate which variables are in a bank's sparse state, I employ the Lasso. The Lasso augments the standard ordinary least squares (OLS) objective function with a penalty term on the coefficients. The result of this penalty term is that many coefficients will be set at exactly zero. Thus, the Lasso gives a sparse regression, where only the most important variables will have non-zero coefficients associated with them.

Banks sell multiple types of mortgages, each having a different interest rate. An SMPE implies that the policy functions of all the mortgages of the same bank should depend on the same sparse state. This is because paying attention has a fixed cost per state variable per firm: thus, it never pays for a bank to ignore a state variable for some mortgages but not for others. In other words, it is required to select the same variables for all policy functions of the same bank. To accomplish this, I use what the machine-learning literature calls a multi-task Lasso (Zhang 2006; Liu, Ji, and Ye 2009; Obozinski, Taskar, and Jordan 2010). The multi-task Lasso simultaneously estimates the coefficients of multiple models, imposing the same sparsity structure on all of them.⁴⁶

The multi-task Lasso works as follows. Group bank b 's interest rates in a $T \times |\mathcal{J}_b|$ matrix R_b , where the (t, j) -element of R_b is r_{jt} , the interest rate charged for mortgage j .⁴⁷ Let X be a $T \times K$ matrix of possible regressors, where K can be larger than the number of observations T . The multi-task Lasso is defined as

$$\hat{B} = \arg \min_B \frac{1}{2T} \|R - XB\|_{\text{Fro}}^2 + \lambda \|B\|_{21},$$

where $\|A\|_{\text{Fro}} = \sqrt{\sum_{ij} a_{ij}^2}$ is the Frobenius norm and $\|A\|_{21} = \sum_i \sqrt{\sum_j a_{ij}^2}$ is the $\ell_1 \ell_2$ norm. $\lambda > 0$ is the regularization parameter that generates sparsity. The larger λ , the sparser the set of selected variables. When $\lambda = 0$, the multi-task Lasso reduces to a separate OLS regression for each policy function. The j 'th column of the estimated coefficients \hat{B} , call it \hat{B}_j , contains the estimated parameters for r_j . The multi-task Lasso constrains these estimates such that all \hat{B}_j 's of the same bank have the same sparsity structure. In other words, \hat{B} will contain a zero row for every covariate that is not included.⁴⁸

The most common method to select the penalty parameter λ is k -fold cross-validation.⁴⁹ This

46. An additional advantage of using the multi-task Lasso is that it increases the number of effective observations in the Lasso regression. This can matter since I only have 36 observations (three years of monthly data) per mortgage. However, every bank sells 24 different mortgages, so that when I use a multi-task approach I have $24 \times 36 = 864$ observations per bank to identify the sparsity of its policy functions.

47. I only estimate policy functions on the post-ban data, so there is no need to distinguish between the interest rates charged to new and renewing customers.

48. One downside of this approach is that the multi-task Lasso also chooses the same functional form for every policy function when higher-order polynomials and interactions of variables are included. I have experimented with running a separate Lasso for every policy function to estimate its functional form after using the multi-task Lasso to pick which variables are in its domain, but this did not lead to significantly different results.

49. Cross-validation is a data-driven approach to choosing λ . The data set is split into a training and test set. For different values of λ , the model is estimated on the training set. The estimated model that gives the best out-of-sample fit on the test set is chosen. In k -fold cross validation, this is done k times and the model with the best average out-of-sample prediction over the k folds is chosen.

procedure however assumes that the data are i.i.d. and banks' interest rates display serial correlation (see Figure 1). Therefore, I use time series cross-validation: I use the first t observations to estimate the model and the $t + 1$ 'th observation to calculate the out-sample-fit. I do this for every $t = 1, \dots, T - 1$ and take the value of λ that gives the best average out-of-sample fit over all values of t .

After using the Lasso to estimate the sparsity of banks' policy functions, I use OLS to estimate their coefficients as in Belloni and Chernozhukov (2013). Using post-Lasso OLS is important, since the Lasso alone shrinks coefficients towards zero too much, leading to implausible counterfactual interest rates.⁵⁰

In X I include (functions of) the payoff-relevant state variables i_t and $f_t(b, D)$. The Lasso only works well if all independent variables are on the same scale. Therefore, I standardize them. One problem is that $f_t(b, D)$ is a density and thus infinite-dimensional. Therefore, I discretize this state variable by including as variables the market shares of households within certain brackets as follows. First, I include the overall market shares, $\int f_t(b, D)dD$. Then, I split the sample into two groups (rich and poor), and add the conditional market shares $\int_{d>E[D]} f_t(b, D)dD$ and $\int_{d\leq E[D]} f_t(b, D)dD$. Then, I divide the sample in four groups, etcetera. These variables are however highly correlated. When the independent variables are highly correlated, the Lasso can become unstable and pick a somewhat arbitrary variable from the set of correlated variables. For example, it might pick the market share of bank A among the poorest quarter of the population but not the overall market share of bank A. This is largely inconsequential for the fit, but makes the interpretation of the estimates somewhat difficult: it does not make a lot of sense for a bank to pay attention to the market share amongst a subsection of the population but not the overall market share.

Therefore, I employ the following procedure. First, I include the marginal market shares of deciding households, $\int f_t(b, D)dD$, as variables. I run the multitask post-Lasso OLS as described above using a second-degree polynomial of these variables. Then, I split the sample into two groups (rich and poor), and add the conditional market shares $\int_{d>E[D]} f_t(b, D)dD$ and $\int_{d\leq E[D]} f_t(b, D)dD$. In the spirit of the Frisch-Waugh-Lovell theorem, I run a multitask Lasso regression of the residuals of the first regression on the residuals of a regression of the split market shares on the non-split market shares to see whether they offer additional explaining power. If they do, I split every group again—so you get the market shares for the four income quantiles—and use the same procedure until no variables are added.

6.2 Demand

I estimate the demand model using maximum likelihood. To control for the endogeneity of interest rates I use the control function approach (Petrin and Train 2010). I use this approach instead of the more typical approach of Berry, Levinsohn, and Pakes (1995, BLP hereafter) for the following reasons. First, I want to exploit that I have so-called "micro-level" data, instead of the "market-level"

50. I have also tried using the fitness-based threshold Lasso combined with post-Lasso OLS, which according to Belloni and Chernozhukov (2013) is superior to just post-Lasso OLS. However, the thresholding did not induce any additional sparsity in my case so that the fitness-based threshold Lasso is equivalent to ordinary Lasso.

data that BLP requires. BLP can be augmented to incorporate micro-level data as well (Berry, Levinsohn, and Pakes 2004), but the control function approach is more efficient since it is based on maximum likelihood rather than the generalized method of moments like BLP. Second, many mortgages are not sold every month. Thus, would I use BLP, I would face the “zero market share” problem.⁵¹ However, zero market shares are not problematic for the control function approach.

To apply the control function method, I replace the unobserved product quality ζ_{jt} in consumers’ utility specification (1) by so-called control functions, $\hat{\mu}_{jt}$:

$$\tilde{u}_{ijt} = X_{jt}\Pi_D D_{it} - (\alpha_{it}r_{jt1} + s_{it})\Delta_{ijt} - \alpha_{it}r_{jt0}(1 - \Delta_{ijt}) + \psi\hat{\mu}_{jt} + \varepsilon_{ijt}.$$

Here ψ is an additional parameter to be estimated. The control functions are residuals from a first stage regression. I use the policy function regressions from the previous section as the first stage. This means I use the state variables of my model, banks’ common cost of funding and previous market shares, as instruments. The cost of funding is a valid instrument since it shifts banks’ supply curves but does not directly affect demand. Previous market shares are a valid instrument if they are uncorrelated with current unobserved product qualities. Since market shares from previous periods are obviously correlated with unobserved product qualities from that same period, this requires that there is not too much autocorrelation in unobserved product qualities. Most renewers purchased their mortgage ten years prior. Therefore, there should be no autocorrelation at a lag of ten years.⁵²

These instruments identify the effect of interest rates on demand. The identification of the other parameters of the demand model is standard. Differences in market shares between mortgages identify the mean contribution to utility of every product characteristic. Differences in purchases between households with different incomes identify how the coefficients in the demand vary with household income.

To estimate the model, I make the usual assumption that the idiosyncratic utility shocks ε_{ijt} follow a Type I Extreme Value distribution, so that the probability that household i purchases mortgage j can be written as

$$p_{ijt} = \frac{\exp\{\tilde{u}_{ijt}\}}{\sum_{k \in \mathcal{C}_i} \exp\{\tilde{u}_{ikt}\}}.$$

As explained in Section 4, I have a choice-based sample: households that purchase the outside option have a smaller probability of being in my sample than households that purchase an inside good. Manski and Lerman (1977) show that reweighing the likelihood function gives consistent estimates in this case. Let q_i be the sampling probability of household i . The log-likelihood is then

$$\log \mathcal{L}(\Pi_D, \Pi_{ac}, \Sigma) = \sum_t \sum_{i \in \mathcal{H}_t} \frac{\log p_{ijt}}{q_i}.$$

When calculating standard errors, I account for first-stage estimation error using the methods of

51. The problem is that BLP is based on inverting a system of equations based on the logarithms of the market shares. Since the logarithm of zero is undefined, BLP does not work when some products have zero market shares.

52. There is no issue if there is autocorrelation between one year’s quality and the next, as there will undoubtedly be.

6.3 Supply: marginal cost estimation

I estimate the supply side to get banks' marginal costs. The estimation of the supply side game consists in two stages. In the first stage, I estimate policy functions and which variables banks pay attention to—this is explained in Section 6.1. I now show how to use the first stage policy function estimates to estimate banks' marginal costs.

To estimate marginal costs, I use banks' stochastic Euler equations. It turns out that banks' number of controls (interest rates) is larger than the dimension of their sparse states. As a result, I can generate "moments" using only data from the current period. To see why, consider an example of a single agent that has two controls r_1, r_2 and one (perceived) state variable σ . Its Bellman equation is

$$V(\sigma) = \max_{r_1, r_2} \pi(\sigma, r_1, r_2) + \beta \text{EV}(\Gamma(\sigma, r_1, r_2)).$$

As is typical for dynamic models, it is possible to reformulate the Bellman equation so that the bank directly chooses the future state σ' :⁵³

$$V(\sigma) = \max_{\sigma'} \tilde{\pi}(\sigma, \sigma') + \beta \text{EV}(\sigma'), \quad (5)$$

$$\tilde{\pi}(\sigma, \sigma') = \max_{r_1, r_2 \text{ s.t. } \Gamma(\sigma, r_1, r_2) = \sigma'} \pi(\sigma, r_1, r_2). \quad (6)$$

The first order conditions of the maximization problem (6) are

$$\frac{\partial \pi}{\partial r_i} - \mu^T \frac{\partial \Gamma}{\partial r_i} = 0 \quad (7)$$

for $i = 1, 2$, where μ is the Lagrange multiplier. Except for μ and any unknown parameters of the model, all quantities in this expression are observed or can be calculated. Since there are two first order conditions and one multiplier, one parameter of the model can be estimated without using the first order conditions implied by the Bellman equation (5).

The intuition for this result is as follows. Say that $\frac{\partial \Gamma}{\partial r_1} = \frac{\partial \Gamma}{\partial r_2}$. If $\frac{\partial \pi}{\partial r_1} > \frac{\partial \pi}{\partial r_2}$, the bank has the following profitable deviation. It can increase r_1 and simultaneously decrease r_2 by ε . Because $\frac{\partial \Gamma}{\partial r_1} = \frac{\partial \Gamma}{\partial r_2}$, this leaves next period's state unaltered. However, since $\frac{\partial \pi}{\partial r_1} > \frac{\partial \pi}{\partial r_2}$, this deviation strictly increases today's profits. Therefore, if $\frac{\partial \Gamma}{\partial r_1} = \frac{\partial \Gamma}{\partial r_2}$, it is required that $\frac{\partial \pi}{\partial r_1} = \frac{\partial \pi}{\partial r_2}$. Generalizing this logic, optimality requires that in every month and for every interest rate a bank sets, the derivative of its profit function with respect to that interest rate is proportional to the derivative of the evolution law.

The same result can be derived for the full model. Let S_b be the number of variables in bank b 's

53. The equivalence becomes immediately obvious by taking the first order conditions of both formulations.

sparse state $\hat{\sigma}_b$. The equivalent of (7) is

$$g_{bt}(\gamma, \lambda) \equiv \frac{\partial \pi_b}{\partial r_b} - \frac{\partial \Gamma_b}{\partial r_b} \mu_{bt} = 0, \quad (8)$$

where μ_{bt} contains the S_b multipliers of bank b in month t . Post-ban, there are $|\mathcal{J}_b|$ first order conditions: one for every interest rate. If $S_b < |\mathcal{J}_b|$, as is the case, these first order conditions contain additional information that can be used to estimate banks' marginal costs.⁵⁴ I derive the functional forms of these first order conditions in Appendix B.2. It is important to note that, following the definition of an SMPE, a bank maximizes its profits conditional on the policy functions of its competitors in its own sparse state. Thus, when calculating the first order conditions of bank b , one must use the estimated policy functions of its competitors evaluated in b 's sparse state.

Using only the first order conditions implied by (6) and not those implied by (5) has various benefits. First, it does not require the direct calculation of expectations, nor their estimation by substitution of future observed values. This is especially important in my application as in an SMPE banks have non-rational expectations and using this method allows estimation without specifying how banks *do* form expectations.⁵⁵ Also, it is possible to estimate marginal costs without specifying banks' discount factor β , which is typically not identified without further exclusion restrictions. Of course, these benefits come at the cost of a loss of efficiency if the model is specified correctly and the discount factor is identified or known.

I estimate marginal costs as follows. For every candidate parameter vector γ_b and period t , I find the Lagrange multipliers $\mu_{bt}(\gamma_b)$ —subscripted to indicate their dependence on the trial parameters—that solve bank b 's first order conditions (8). Since this an overdetermined system of equations, it will not be possible to find μ_{bt} that solve (8) for all (or any) γ_b . Therefore, I find $\mu_{bt}(\gamma_b)$ by OLS. I then search over parameters γ_b to minimize the average (over months) first order conditions, i.e. I solve the minimum distance objective

$$\min_{\gamma_b} \left(\frac{1}{T} \sum_{t=1}^T g_{bt}(\gamma_b, \mu_{bt}(\gamma_b)) \right)^T \left(\frac{1}{T} \sum_{t=1}^T g_{bt}(\gamma_b, \mu_{bt}(\gamma_b)) \right).$$

The value of γ_b that minimizes this objective function gives estimates of bank b 's marginal cost parameters.

54. I stress that there is nothing in the definition of an SMPE that requires this to be the case; it is what I find empirically. For example, one could also estimate the SMPE of a model with one control, so that there are never any additional degrees of freedom whatever the sparsity structure of banks' policy functions. When this method is not applicable, value function iteration as in Bajari, Benkard, and Levin (2007) can be used instead.

55. In the definition of an SMPE, I have specified that banks form expectations according to the value function induced by (4). This means that banks are aware that they are sparse maximizers. However, as (Gabaix 2017) discusses, it is also possible that banks are "naive" and expect to obtain the rational value function (without sparse maximization). Estimating marginal costs in this way is thus robust to this assumption.

Table 6: Structure of banks' policy functions.

	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Cost of funding	✓	✓	✓	✓	✓	✓
Market share Bank A	✓	✓	✓	✓	✓	✓
- Market share Bank A among richest 50%						
- Market share Bank A among poorest 50%						
Market share Bank B	✓	✓		✓		✓
- Market share Bank B among richest 50%						
- Market share Bank B among poorest 50%						
Market share Bank C	✓	✓				✓
- Market share Bank C among richest 50%						
- Market share Bank C among poorest 50%						
Market share Bank D						
- Market share Bank D among richest 50%						
- Market share Bank D among poorest 50%						
Market share Bank E						
- Market share Bank E among richest 50%						
- Market share Bank E among poorest 50%						
Market share Bank F	✓	✓		✓		✓
- Market share Bank F among richest 50%						
- Market share Bank F among poorest 50%						
R^2	0.911	0.777	0.709	0.823	0.781	0.807

Note: A check mark indicates that the policy function of a bank depends (possibly non-linearly) on the variable. All other market shares conditional on being in a certain income bracket drop out.

7 Results

This section presents my results. I start with the estimates of the structural model: policy functions, demand functions and marginal costs. Then, I answer the main question of this paper: what is the effect of the ban on history-ban price discrimination on consumer surplus, bank profits and welfare?

7.1 Model primitives

7.1.1 Sparse states and policy functions

Table 6 contains the estimated structure of banks' policy functions, and therefore also the variables in their sparse states. The multitask Lasso leads to a significant reduction of the size of the state space: the actual state is infinite-dimensional, the sparse states contain at most five variables. Yet, the fit is excellent. The R^2 's imply that these few variables are able to explain most variance in observed interest rates. Remember that the residuals from these policy function estimates are used as control functions in the demand estimation. Therefore, the high R^2 's moreover imply that the model's state variables are highly relevant instruments.

Every bank's policy function depends on the cost of funding, as expected. In addition, their policy functions depend on the overall market shares of one to three other banks. No bank's policy function depends on market shares conditional on households falling into a certain income bracket. Although I cannot report the results on which this is based because of data confidentiality, banks are more likely to pay attention to the largest bank in the market or to price-fighting banks that tend to offer low interest rates.

7.1.2 Demand

Table 7 contains estimates of the demand model. The estimates imply that the demand for mortgages is somewhat elastic, with an average own-interest rate elasticity of -1.25 . The results confirm that demand side heterogeneity is important. All interaction effects with household income are statistically significant at conventional levels. Heterogeneity is particularly important for households' sensitivity to interest rates. The estimated parameter is three times as large for a household with an income one standard deviation above the average than for a household with an average income. As expected, households' sensitivity to interest rates decreases in their income.

In addition, the estimates imply that there are significant switching costs. For example, a household on average only switches to a mortgage giving it the same utility if its interest rate is at least 3.4 percentage points lower. This large estimate has two explanations. First, as I argued in Section 5, these switching costs measure many different frictions that I cannot differentiate in the data. For example, it is known that in this market there is significant inattention as well: survey evidence indicates that 60% of households that receive an offer to renew their mortgage do not consider switching.⁵⁶ This fact is also picked up by the high switching cost estimate. Note moreover that since both the coefficient on the interest rate and switching costs differ by household income, there is significant heterogeneity in the propensity to switch. For example, a household with an income just one standard deviation below the average switches to an otherwise equivalent mortgage when the interest rate is 1.4 (instead of 3.4) points lower.

7.1.3 Marginal cost estimates

Table 8 contains the average marginal costs of the different types of mortgages sold in the Dutch mortgage market. For comparison, I also estimate the supply model under the assumption that banks are not forward-looking and optimize their flow profits. On the whole, the patterns make sense. Longer fixed interest rate periods are associated with higher marginal costs. The marginal costs of bullet, savings and investment mortgages are lower than of other types of loans. This reflects that these type of mortgages cross-subsidize other (often mandatory) products sold by the same bank, such as investment products. Finally, the interest rate on demand deposits, an important source of funding for Dutch banks, was around .5 percentage points in this period. This is consistent with the estimated marginal costs for 5-year loans, which are only somewhat higher.

56. Vereniging Eigen Huis. "Huiseigenaren laten massaal geld liggen bij nieuwe rente hypotheek", January 3 2018. <https://www.eigenhuis.nl/actueel/nieuws/2018/01/03/09/00/>

Table 7: The estimated demand model.

	Constant term		Interaction with household income	
	Coefficient	Std. error	Coefficient	Std. error
Bank				
Bank A	-0.342	(0.008)	-0.888	(0.004)
Bank B	-0.220	(0.009)	-0.721	(0.004)
Bank C	-1.866	(0.012)	-0.610	(0.005)
Bank D	-0.841	(0.009)	-0.868	(0.004)
Bank E	-0.196	(0.010)	-0.869	(0.004)
Bank F	-0.435	(0.010)	-0.898	(0.004)
Payment type				
Annuity	-	-	-	-
Linear	-2.641	(0.002)	0.327	(0.001)
Bullet	0.009	(0.017)	0.325	(0.003)
Savings mortgage	0.311	(0.019)	0.331	(0.004)
Life mortgage	-0.195	(0.014)	-0.256	(0.004)
Investment mortgage	-1.569	(0.014)	-0.102	(0.003)
Duration fixed interest rate				
5 years	-	-	-	-
10 years	2.042	(0.006)	-0.338	(0.003)
15 years	-0.395	(0.012)	-0.440	(0.003)
20 years	0.600	(0.014)	-0.280	(0.003)
Interest rate				
Π_α	-0.423	(0.006)	-0.825	(0.002)
Switching costs				
Π_c	0.807	(0.003)	0.008	(0.001)
Control functions	✓			
Bank dummies	✓			
Log likelihood	-101387237			
Observations	283396			

Note: The table reports the coefficients of households' utility specification (1). The baseline product is a 5 year annuity. Standard errors, in parentheses, account for first stage estimation error (Karaca-Mandic and Train 2003). Household income is standardized.

Table 8: Average marginal costs by payment method and fixed interest rate duration.

	Dynamic				Static			
	5 yr.	10 yr.	15 yr.	20 yr.	5 yr.	10 yr.	15 yr.	20 yr.
Annuity	0.87	1.49	1.82	1.89	0.87	1.42	1.80	1.83
Linear	0.64	1.26	1.59	1.66	0.66	1.21	1.59	1.62
Bullet	0.49	1.11	1.44	1.51	0.48	1.03	1.41	1.44
Savings	0.52	1.14	1.47	1.54	0.51	1.07	1.44	1.47
Life	0.97	1.59	1.91	1.99	0.94	1.49	1.87	1.90
Investment	0.76	1.38	1.70	1.78	0.78	1.33	1.70	1.74

Note: The table shows average marginal costs per mortgage type. The average is unweighed over the banks in my sample. Marginal costs are calculated for an interest rate on Dutch government bonds of $i_t = 1.37$, the average in the post-ban period. The left-hand side of the table shows marginal cost estimates for the model described in the main text. The right-hand side of the table shows marginal cost estimates under the assumption that banks are not forward-looking.

The estimates of marginal costs are on the whole somewhat larger under the assumption that banks are forward-looking under static maximization. For the most common fixed interest rate duration, ten years, the difference is about 5%. This means that would I not account for banks' dynamic incentives, I would overestimate banks' profits.⁵⁷

Figure 4 shows the dispersion of marginal costs across banks. Because the dispersion is significant, there is potential for history-based price discrimination to introduce cross-segment inefficiencies—consumers purchasing from inefficient banks because they are locked-in there. An important mechanism to keep in mind when evaluating the ban will thus be the (re-)allocation of consumers across banks.

7.2 The effect of history-based price discrimination on consumer surplus, profits and welfare

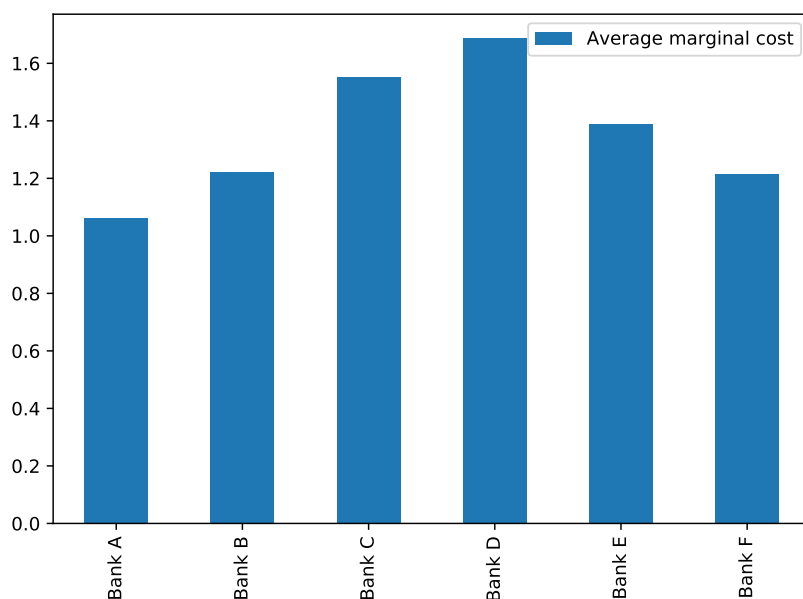
Figure 5 shows the predicted average interest rate implied by the estimated policy functions. The post-2013 prediction is in-sample and shows the fit of the model. The pre-2013 prediction is out-of-sample. Therefore, it shows the counterfactual interest rates implied by the model, had the ban on history-based price discrimination been instituted before 2013. The counterfactual uniform interest is typically below the interest rates for renewing and for new customers.⁵⁸ This means that the

huiseigenaren-laten-massaal-geld-liggen-bij-nieuwe-rente-hypotheek

57. This mirrors the recent finding of MacKay and Remer (2018) that it is important to account for forward-looking behavior of firms in markets with customer inertia: firms have a so-called "investment motive" in that they might want to charge a lower price than under static profit optimization to attract future locked-in customers. Here, in a static setting the interest rates banks set can only be explained by relatively low marginal costs. However, once I account for banks' forward-looking behavior estimated marginal costs are higher. However, the differences I find are not as large as in MacKay and Remer (2018).

58. One may expect that the uniform interest rate must lay between those two. However, Corts (1998) shows that in oligopolies third-degree price discrimination may lead to prices for *all* consumer segments to be below or above the uniform price. This can happen when firms have different "strong" markets, i.e. markets where the elasticity of demand they face is comparatively low. When there are switching costs, this the case: a bank's locked-in customers form its strong

Figure 4: Marginal costs across banks.



Note: The figure shows a histogram of average marginal costs per bank, where the average is taken over all mortgages a bank offers. The marginal costs are calculated for the interest rate on 10-year Dutch bonds equal to $i_t = 1.37$, the average in the post-ban period.

rent-extraction effect is stronger than the competition effect.

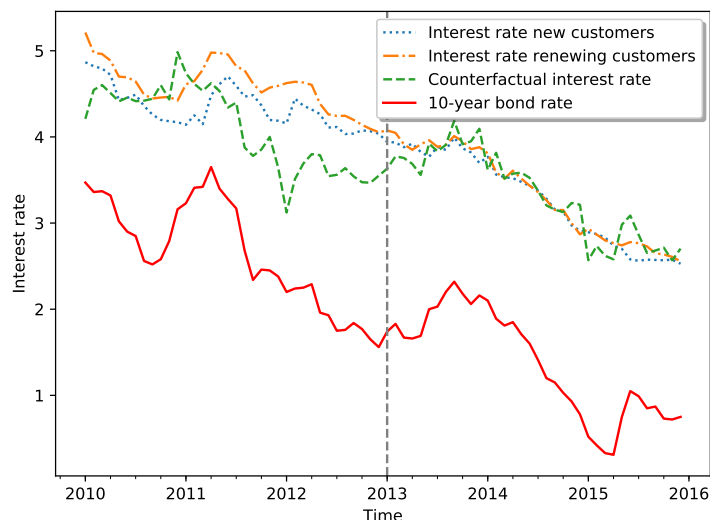
Figure 6 shows the counterfactual interest rates when I do not estimate an SMPE and do not use the Lasso, but instead make the ad hoc assumption that banks care about the market shares of all banks, but not how different types of households are distributed amongst them. This is an assumption that in principle seems very reasonable, but Figure 6 shows clearly that the resulting counterfactual interest rates are not: the counterfactual interest rates are erratic and often quite high or low. This shows the danger of ad-hoc assumptions and the importance of data-driven methods: even assumptions that a priori seem reasonable run the risk over over-fitting and can create unrealistic outcomes when extrapolated to out-of-sample states.

Given the counterfactual interest rates, I can calculate the effect of the ban on consumer surplus and bank profits. Table 9 shows the estimated effects on consumer surplus. For an average mortgage of €150,000, expected consumer surplus increases by €415 per year. Households whose mortgages are up for renewal gain the most, €588 per year for an average mortgage. This is particularly because the consumer surplus of renewing at their current bank is more attractive, since interest rates are lower under the ban: the expected consumer surplus of renewing increases by €848. Households are also about 3.3 percentage points less likely to switch. Since switching is costly, this means that the consumer surplus of switching also increases, by €434 per year.

While consumers gains, banks lose. Table 10 shows that, for an average mortgage of €150,000, the ban on history-based price discrimination causes a loss in profits of €290 per mortgage per year. This loss can be split up into two parts: the first, which I call the reallocation effect, calculates

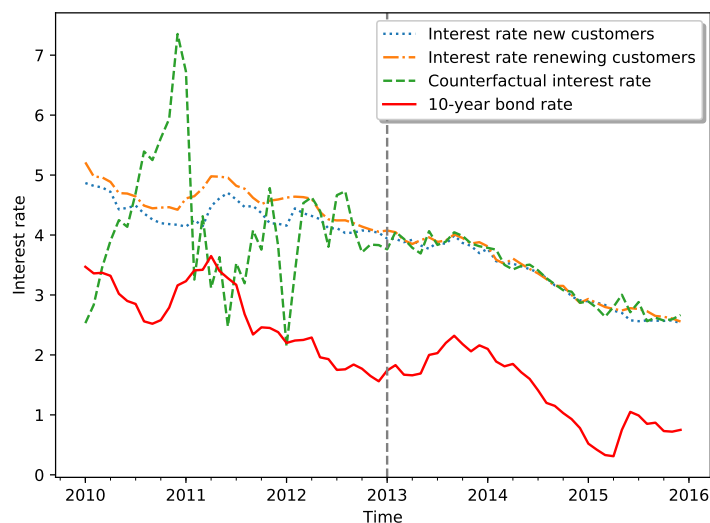
market.

Figure 5: Counterfactual and observed interest rates.



Note: The figure shows observed interest rates and counterfactual interest rates as predicted by the structural model. All are weighed by observed market shares. The interest rate on 10-year Dutch bonds is included for comparison.

Figure 6: Counterfactual and observed interest rates when the Lasso is not used.



Note: The figure shows observed interest rates and counterfactual interest rates as predicted by the structural model, when the Lasso is not used to estimate policy functions. Instead the ad-hoc assumption is made that policy functions depend on the market shares of all banks, but not on the distribution of household incomes per bank. All are weighed by observed market shares. The interest rate on 10-year Dutch bonds is included for comparison.

Table 9: The effect of the ban on history-based price discrimination on consumer surplus.

	Potential first-time buyers	Renewing households	Total
Probability of switching (%)		-3.259	
Consumer surplus of switching (€)		430	
Consumer surplus of renewing (€)		848	
Total consumer surplus (€)	185	588	415

Note: The table shows the yearly effect of the ban on history-based price discrimination on consumer surplus for a mortgage with a balance of €150,000. The effects are calculated by evaluating the estimated structural model at observed and counterfactual interest rates, then taking the difference. The effects are an average over the post-ban period 2010-2013, as well as an average over all households who were active in the market during that period, weighed by loan sum.

Table 10: The effect of the ban on history-based price discrimination on bank profits.

	Difference (€)
Profits	-290
– Reallocation effect	229
– Price effect	-519

Note: The table shows the yearly effect of the ban on history-based price discrimination on bank profits for a mortgage with a balance of €150,000. The effects are calculated by evaluating the estimated structural model at observed and counterfactual interest rates, then taking the difference. The effects are an average over the post-ban period 2010-2013, as well as an average over all households who were active in the market during that period, weighed by loan sum. The reallocation effect calculates the effect on profits of counterfactual market shares, holding interest rates constant at observed levels. The price effects measures the effect on profits of counterfactual interest rates, holding market shares constant at counterfactual levels.

profits under counterfactual market shares but observed interest rates. The reallocation effect thus measures if consumers purchase from more efficient banks because of the ban.⁵⁹ My results indicate that this is the case. For an average mortgage of €150,000, average bank profits increase by €229 per year because of consumer reallocation. Of course, this average belies some heterogeneity: more efficient banks gain and less efficient banks lose.

The second effect is the price effect. The price effect measures the change in bank profits going from observed to counterfactual interest rates. Since interest rates on the whole are lower under the ban, the price effect is negative: holding the market shares fixed at the counterfactual market shares, lower interest rates cause banks' profits to decrease with €519 per year for a mortgage of €150,000. As it turns out, the price effect dominates the reallocation effect and the ban on history-based price discrimination causes average bank profits to decrease.

Adding up, the ban on history-based price discrimination causes an increase of total welfare of €125 per year for a mortgage of €150,000. For the whole NHG segment of the market, this implies a welfare increase of about €11.5 million per year.

To summarize, why does the ban on history-based price discrimination increase welfare? As it turns out, it reduces the following three inefficiencies. First, as mentioned above, the average interest rate is lower without than with history-based price discrimination. Because interest rates

59. Theory suggests that price discrimination might cause cross-segment inefficiencies (Stole 2007), that is, might cause consumers to purchase from inefficient firms.

are closer to marginal costs, efficiency increases. Second, there is less switching without history-based price discrimination. Since switching costs are a deadweight loss, this also increases welfare. Finally, I find that reallocation towards more efficient banks is crucial to understanding the welfare increase: without the reallocation effect, the effect of the ban on welfare would be $\text{€}125 - \text{€}229 = -\text{€}104$. Therefore, the focus of the theoretical literature on history-based price discrimination on symmetric firms is misguided: cost asymmetries and reallocation can be of first-order importance to understanding the effects of history-based price discrimination.

8 Discussion

In this section I discuss the robustness of my results. First, I discuss other changes in the Dutch mortgage market around the time of the ban on history-based price discrimination and how those changes might bias my results. Then, I try quantify this potential bias by computing pre-ban interest rates from the post-ban model. Finally, I discuss the empirical content of SMPE and how this can be used to test whether estimated attention vectors are consistent with sparse maximization.

8.1 Other changes in the Dutch mortgage market

The estimates of the effects of the ban on history-based price discrimination are causal to the extent that there are no exogenous changes to the Dutch mortgage market that cause interest rates to differ before and after the ban. My estimates control for changes in the cost of funding and previous market shares. However, various other changes happened in the Dutch mortgage market around the ban on history-based price discrimination. I now discuss some of these changes and how they might impact my results.

The general consensus is that after 2013 the Dutch mortgage market became more competitive. The market share of smaller banks grew and margins decreased (Fransman 2017). My results are consistent with this change—I find that the ban on history-based price discrimination made the mortgage market more competitive. Indeed, this was the main aim of the Dutch competition authorities in proposing the ban. However, because I do not have a control market, I cannot definitively rule out that this increase in competitiveness was caused by other factors than the ban. This would imply that my estimates of the effects of the ban on consumer surplus are biased upwards and my estimate of effect of the ban on bank profits is biased downwards.

A second change is the ban on non-amortizing mortgages for new home purchases in 2013. This leads to a change in households' choice sets, which as I describe in Section 5, I control for. This change in choice sets could as a secondary effect also change mortgage pricing. If banks' pricing changes as a result of this ban, one would expect the difference in interest rates between amortizing and non-amortizing mortgages to change. I test this implication in Table 11, where I regress the difference in interest rates between annuity and bullet loans on loan fixed effects and a dummy indicating the post-ban period. I find no statistically or economically significant change in the difference between the interest rates of annuities and bullets after 2013. Therefore, I conclude that

Table 11: The interest rate difference between annuity and bullet mortgages, pre- and post-ban.

	(i)	(ii)
After 2013	0.034 (0.018)	-0.115 (0.063)
Bank \times fixed interest rate duration fixed effects	✓	✓
Month fixed effects		✓
Observations	1728	1728

Note: The table displays the change in the interest rate difference between annuities and bullets after the ban on history-based price discrimination. HAC-robust standard errors are in parentheses.

the ban on non-amortizing mortgages for new purchasers after 2013 has no effect on mortgage pricing.

Another potential confounder might be an exogenous change in the demand for mortgages. Given that the economy and housing market markedly improved after 2013, it is natural to assume that, if anything, the demand for mortgages increased during this period. Since this would imply that interest rates increase, this is difficult to reconcile with the general observation that the market became *more* competitive after 2013. However, to the extent that the demand for mortgages increased after 2013, ignoring this in my analysis would only offset the fact the market became more competitive.

A final change to the Dutch mortgage market is that regulations for mortgage brokers became stricter over time. Perhaps the most important change is the 2013 ban on commission payments from banks to brokers. This could have an important effect on the mortgages brokers recommend, and since about half the mortgages are sold through brokers, on the overall market. As Thiel (2018) shows, banning commissions could lead to lower interest rates as commissions soften competition. Ignoring this could therefore lead to a further overestimation of the welfare improvement of the ban.

8.2 Predicting pre-ban interest rates

As a further check on the model, I compute pre-ban interest rates from the estimated post-ban model. The idea is that, if the other changes in the Dutch mortgage market around the time of the ban are inconsequential, I should be able to predict pre-ban interest rates from the post-ban model by allowing banks to engage in history-based price discrimination. I do this by following Hortacsu and Puller (2008): given estimated policy functions of its competitors, I calculate a bank's response and see how close it is to its actual interest rates.

To do so, I first re-estimate banks' policy functions based on the pre-ban data. The reason I do this is that when history-based price discrimination is possible banks may find it optimal to pay attention to different state variables than when it is not. Using time series cross-validation as in Section 6.1.2 gives estimated sparse states with significantly more variables than in the post-ban period (i.e. more than in Table 6). Because calculating a bank's best response falls prey to the curse

Table 12: Predictions of pre-ban interest rates from the post-ban model.

	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
% deviation interest rate renewing customers	4.985	28.8	-11.395	8.995	-2.305	-0.915
% deviation interest rate new customers	-29.69	3.345	-11.135	-5.275	-41.065	-26.435

Note: The table shows percentage deviations of computed best response interest rates from observed interest rates, averaged over all mortgages a bank sells.

of dimensionality when it pays attention to many variables, I choose a larger penalty parameter than indicated by time series cross-validation in the Lasso estimation to make calculation of best responses feasible.⁶⁰

Then, I solve bank b 's dynamic profit maximization problem. Bank b sets its interest rates in sparse state $\hat{\sigma}_b$, where the interest rates of bank b 's competitors are calculated from their estimated policy functions. That is, for every sparse state $\hat{\sigma}_b$, bank b maximizes

$$\pi_b(\hat{\sigma}_b, r, \hat{r}_{-b}(\hat{\sigma}_b)) + \beta \mathbb{E}[V(\Gamma(\hat{\sigma}_b, r, \hat{r}_{-b}(\hat{\sigma}_b)))]. \quad (9)$$

Here, \hat{r}_{-b} are the estimated policy functions of b 's competitors, evaluated in the sparse state $\hat{\sigma}_b$. I calculate a bank's best response using value function iteration. Full details of the algorithm can be found in Appendix B.3. Here, I mention only two important simplifying assumptions I make during this calculation.

First, I assume that banks only sell ten-year mortgages. Table 2 shows that almost 70% of mortgages are of this duration. This greatly simplifies the calculation of (9), since otherwise the future states at five, fifteen and twenty years also have to be taken into account. Given this assumption, I use a discount factor of 2% per year (around the average of the cost of funding in the post-ban period), so that $\beta = .98^{10}$.

A second simplifying assumption I make is that there is no uncertainty around the cost of funding. Because the estimated policy functions (Figure 5) vary almost one-to-one with the cost of funding, I assume cost of funding is constant. This reduces the computational time greatly since it removes the need to calculate expectations.

Table 12 displays the percentage deviation of the computed best response interest rates from the observed interest rates. The first pattern that emerges is that the model does an adequate job when predicting the interest rates charged to renewing customers: for three out of six banks, the model predicts interest rates within 5% of the observed interest rates, for five out six within 12%. Moreover, for half of the banks the model predicts higher interest rates than observed, for half lower. This suggests that the other changes in the Dutch mortgage market around the time of the ban on history-based price discrimination did not systematically increase or decrease interest rates.

The computed interest rates for new customers show larger deviations from observed interest rates than the computed interest rates for renewing customers. For three out six banks, the model does an adequate job when predicting observed interest rates. For the remaining three banks, the

60. To be precise, I choose the penalty parameter such that every bank pays attention to two or three state variables.

computed interest rates are significantly lower than observed. There are two possible interpretations of this fact. The first is that the divergence between computed and observed interest rates is caused by the other changes in the Dutch mortgage market around the time of the ban on history-based price discrimination. This is however difficult to reconcile with the size of the divergence and the fact that this divergence only occurs for the interest rates for new customers. If, for example, the very low interest rates computed for bank A's new competitors were due to the increased post-ban competitiveness of the market, it is difficult to explain why A's computed interest rates for renewing customers are higher than observed. A second interpretation is that economic models of history-based price discrimination do not capture the mechanisms of the Dutch mortgage market well. In this interpretation, the calculated interest rates overestimate the importance of the so-called investment motive. The investment motive is the incentive of banks to set a low interest to capture a large market share. My computations predict a much larger incentive motive than observed in the data. This shows the importance of having exogenous variation in pricing: a counterfactual simulation, such as in Cosguner, Chan, and Seetharaman (2018), would imply much lower interest rates for new consumers than in fact observed.

8.3 The empirical content of SMPE

In this section, I discuss the empirical content of SMPE. SMPE puts restrictions on the data. These restrictions can be used as a specification test, to choose the Lasso penalty parameter or to estimate the cost of attention.

SMPE has empirical content because under sparse maximization banks pay attention to a state variable if the cost of inattention towards that variable is sufficiently large. Because the cost of inattention can be calculated from an estimated model, this is a testable implication. To be precise, recall that a bank pays attention to state variable σ_i if and only if

$$\Lambda_i \geq \kappa,$$

i.e. when the cost of inattention is greater than the cost of attention. Since I take the default state equal to the long-run average, i.e. $\sigma_i^d = \mathbb{E}[\sigma_i]$, one of the following inequalities must hold for every σ_i :

$$\begin{aligned} \text{Var}[\sigma_i] r_{\sigma_i}^T v_{rr}^b r_{\sigma_i} &\geq \kappa \text{ if } m_i = 1, \\ \text{Var}[\sigma_i] r_{\sigma_i}^T v_{rr}^b r_{\sigma_i} &\leq \kappa \text{ if } m_i = 0, \end{aligned}$$

where

$$v^b(\hat{\sigma}, r) = \pi_b(\hat{\sigma}_b, r, \rho_{-b}(\hat{\sigma}_b)) + \beta \mathbb{E}[V_b(\Gamma(\hat{\sigma}_b, r, \rho_{-b}(\hat{\sigma}_b)))].$$

The testable implication of SMPE is the existence of a scalar κ that solves the system of inequalities for all state variables. Intuitively, an estimated model is only consistent with sparse maximization if there exists a cost of attention that generates the estimated attention vectors.

To implement the test, note that if a bank pays attention to σ_i but not σ_j , it is possible to add their respective inequalities to get

$$\Lambda_i - \Lambda_j \geq 0$$

or

$$\frac{\text{Var}[\sigma_i]}{\text{Var}[\sigma_j]} \geq \frac{v_{rr}^b r_{\sigma_j} r_{\sigma_j}}{v_{rr}^b r_{\sigma_i} r_{\sigma_i}}. \quad (10)$$

Because κ has dropped out, this inequality can be calculated from an estimated model. The intuition behind this inequality is that the more a state variable varies, the higher the cost of inattention becomes. Therefore, if a bank pays attention to σ_i but not to σ_j , σ_i must have a relatively large variance.

Equations (10) give a system of moment inequalities for every bank. These can be tested using recent results from the econometrics literature, for example the two-step bootstrap of Romano, Shaikh, and Wolf (2014). One application is as a specification test of the model: are the attention vectors the Lasso chooses actually consistent with sparse maximization? This test can also be used to choose the Lasso penalty parameter λ , instead of cross-validation. One would then estimate the model for various values of λ and choose the value that gives a model that is most consistent with sparse maximization, for example the one with the highest p -value.

I perform these two exercises in Appendix A. First, I find that my model is consistent with sparse maximization for three of out six banks. While I am able to reject sparse maximization for the other three banks, the extent of the violation is not large. For those banks, no more than half of the sample moment inequalities (10) are violated. Second, I use this test to find the Lasso penalty parameter λ . This results in attention vectors that are consistent with sparse maximization for every bank. However, choosing λ this way does not lead to a meaningfully different estimates of the effects of the ban on history-based price discrimination.

The empirical content of SMPE has further applications. For example, it can also be used as a check when computing a counterfactual equilibrium using the estimated attention vectors. Calculating a counterfactual SMPE under the assumption that attention vectors do not change is convenient: one benefits from the reduction of the state space, but does not required to re-calculate attention vectors. However, it is possible that in a counterfactual situation the estimated attention vectors are no longer optimal. Whether this is the case can be assessed using the test derived above. Simulating from a computed equilibrium gives an empirical distribution of $\text{Var}[\sigma_i]$. This empirical distribution can be used to implement the two-step bootstrap described above.

Second, a similar test statistic can be inverted to estimate the cost of attention κ . An application of this is the calculation of counterfactual attention vectors. The inequalities

$$\Lambda_i \geq \kappa$$

for the state variables a bank pays attention to (and the reverse inequalities for those it doesn't pay attention to), can likewise be interpreted as moment inequalities. Testing these inequalities for different values of κ gives a confidence interval for the cost of attention. Note however, that under

standard asymptotics, in particular holding the number of state variables fixed, κ will only be set identified.

9 Conclusion

This paper studied the effects of a ban on history-based price discrimination in the Dutch mortgage market. In this market, households that renew their mortgage paid on average between €228 and €348 more in interest than new customers for the same loan. Such interest rates differences are possible because of high switching costs.

I estimate the effect of the ban by developing a structural model of demand and supply of the Dutch mortgage market. The supply side is a dynamic game with an infinite-dimensional state space. To deal with this, I introduce a new solution concept, Sparse Markov Perfect Equilibrium (SMPE). In an SMPE, banks only pay attention to the most important state variables. I show how these state variables can be identified using the (multitask) Lasso. My results imply that banks pay attention to between two and five variables, these being a common interest rate and the market of share of some, but not all, other banks. Thus, SMPE reduces the dimension of the problem from infinity to a maximum of five.

Estimates of the demand function and marginal costs then allow me to calculate the effect of history-based price discrimination on consumer surplus, bank profits and welfare. Consumer surplus increases with €415 per year for an average mortgage. This is mainly because locked-in renewing households no longer pay higher interest rates than new customers. Because interest rates are lower, bank profits decrease by €290 per year for an average mortgage, despite the fact that the market share of more efficient banks increases. Adding up, this means that welfare increases with €125 per mortgage per year, or about €11.5 million per year for the whole market.

The analysis has shown the value of SMPE in empirical work. For future research, it would be interesting to investigate its theoretical properties further. In particular, it would be useful to devise an algorithm to calculate an SMPE for a given model.

References

- Abbring, Jaap H., Jeffrey R. Campbell, Jan Tilly, and Nan Yang. 2017. "Very Simple Markov-Perfect Industry Dynamics: Empirics." *Mimeo*.
- Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes. 2007. "Econometric tools for analyzing market outcomes." *Handbook of Econometrics* 6:4171–4276.
- Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential estimation of dynamic discrete games." *Econometrica* 75 (1): 1–53.
- Alé, Jorge. 2013. "Switching Costs and Introductory Pricing in the Wireless Service Industry." *Mimeo*.

- Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2017. "Market Power, Production (Mis)Allocation and OPEC." *Mimeo*.
- Asplund, Marcus, Rickard Eriksson, and Niklas Strand. 2008. "Price discrimination in oligopoly: evidence from regional newspapers." *Journal of Industrial Economics* 56 (2): 333–346.
- Autoriteit Financiële Markten. 2015. "Rentekorting bij hypothekeken." Accessed December 8, 2016. <https://www.afm.nl/nl-nl/professionals/nieuws/2015/mei/rentekorting-hypothekeken>.
- Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating dynamic models of imperfect competition." *Econometrica* 75 (5): 1331–1370.
- Barone, Guglielmo, Roberto Felici, and Marcello Pagnini. 2011. "Switching costs in local credit markets." *International Journal of Industrial Organization* 29 (6): 694–704.
- Barwick, Panle Jia, and Parag A. Pathak. 2015. "The costs of free entry: an empirical study of real estate agents in Greater Boston." *RAND Journal of Economics* 46 (1): 103–145.
- Belloni, Alexandre, and Victor Chernozhukov. 2013. "Least squares after model selection in high-dimensional sparse models." *Bernoulli* 19 (2): 521–547.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile prices in market equilibrium." *Econometrica* 63 (4): 841–890.
- . 2004. "Differentiated products demand systems from a combination of micro and macro data: The new car market." *Journal of Political Economy* 112 (1): 68–105.
- Bils, Mark. 1989. "Pricing in a customer market." *Quarterly Journal of Economics* 104 (4): 699–718.
- Chen, Yongmin. 1997. "Paying customers to switch." *Journal of Economics & Management Strategy* 6 (4): 877–897.
- Collard-Wexler, Allan. 2013. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica* 81 (3): 1003–1037.
- Corts, Kenneth S. 1998. "Third-degree price discrimination in oligopoly: all-out competition and strategic commitment." *RAND Journal of Economics* 29 (2): 306–323.
- Cosguner, Koray, Tat Y Chan, and PB Seetharaman. 2017. "Behavioral Price Discrimination in the Presence of Switching Costs." *Marketing Science* 36 (3): 426–435.

- Cosguner, Koray, Tat Y Chan, and PB Seetharaman. 2018. "Dynamic pricing in a distribution channel in the presence of switching costs." *Management Science* 64 (3): 1212–1229.
- Cullen, Joseph, Nicolas Schutz, and Oleksandr Shcherbakov. 2015. "Welfare analysis of equilibria with and without early termination fees in the US wireless industry." *Mimeo*.
- Cullen, Joseph, and Oleksandr Shcherbakov. 2010. *Measuring consumer switching costs in the wireless industry*. Technical report.
- Dijkstra, Mark A, Fleur Randag, and Maarten Pieter Schinkel. 2014a. "High mortgage rates in the low countries: an introduction." *Journal of Competition Law and Economics* 10 (4): 819–841.
- . 2014b. "High mortgage rates in the Low Countries: what happened in the Spring of 2009?" *Journal of Competition Law and Economics* 10 (4): 843–859.
- Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. 2009. "Do switching costs make markets less competitive?" *Journal of Marketing research* 46 (4): 435–445.
- Fershtman, Chaim, and Ariel Pakes. 2012. "Dynamic games with asymmetric information: A framework for empirical work." *The Quarterly Journal of Economics* 127 (4): 1611–1661.
- Fleitas, Sebastian. 2017. "Dynamic Competition and Price Regulation When Consumers Have Inertia: Evidence from Medicare Part D." *Mimeo*.
- Fransman, Robin. 2017. "Nationale Hypotheek Garantie, een terugblik." *Economisch Statistische Berichten* 102 (4749): 38–42.
- Gabaix, Xavier. 2014. "A sparsity-based model of bounded rationality." *The Quarterly Journal of Economics* 129 (4): 1661–1710.
- . 2017. "Behavioral Macroeconomics Via Sparse Dynamic Programming." *Mimeo*.
- Gehrig, Thomas, Oz Shy, and Rune Stenbacka. 2012. "A welfare evaluation of history-based price discrimination." *Journal of Industry, Competition and Trade* 12 (4): 373–393.
- Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek. 2017. "Inflation dynamics during the financial crisis." *American Economic Review* 107 (3): 785–823.
- Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review* 103 (7): 2643–82.
- Ho, Chun-Yu. 2015. "Switching cost and deposit demand in china." *International Economic Review* 56 (3): 723–749.

- Holmes, Thomas J. 1989. "The effects of third-degree price discrimination in oligopoly." *The American Economic Review* 79 (1): 244–250.
- Honka, Elisabeth. 2014. "Quantifying search and switching costs in the US auto insurance industry." *The RAND Journal of Economics* 45 (4): 847–884.
- Hortacsu, Ali, and Steven L. Puller. 2008. "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market." *The RAND Journal of Economics* 39 (1): 86–114.
- Ifrach, Bar, and Gabriel Y. Weintraub. 2017. "A Framework for Dynamic Oligopoly in Concentrated Industries." *Review of Economic Studies* 84 (3): 1106–1150.
- Ioannidou, Vasso, and Steven Ongena. 2010. "'Time for a change': loan conditions and bank behavior when firms switch banks." *The Journal of Finance* 65 (5): 1847–1877.
- Kalouptsi, Myrto. 2014. "Time to build and fluctuations in bulk shipping." *American Economic Review* 104 (2): 564–608.
- Karaca-Mandic, Pinar, and Kenneth Train. 2003. "Standard error correction in two-stage estimation with nested samples." *Econometrics Journal* 6 (2): 401–407.
- Kiss, András. 2017. "Salience and Switching." *Mimeo*.
- Liu, Jun, Shuiwang Ji, and Jieping Ye. 2009. "Multi-task feature learning via efficient l_2, l_1 -norm minimization." In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, 339–348. AUAI Press.
- MacKay, Alexander, and Marc Remer. 2018. "An Empirical Model of Consumer Affiliation and Dynamic Price Competition." *Mimeo*.
- Manski, Charles F., and Steven R. Lerman. 1977. "The estimation of choice probabilities from choice based samples." *Econometrica* 45 (8): 1977–1988.
- Maskin, Eric, and Jean Tirole. 2001. "Markov perfect equilibrium: I. Observable actions." *Journal of Economic Theory* 100 (2): 191–219.
- Mastrogiacomo, Mauro, and Remco Van der Molen. 2015. "Dutch mortgages in the DNB loan level data." *DNB Occasional Studies* 13 (4).
- Miller, Daniel P., and Jungwon Yeo. 2012. "Estimating Dynamic Discrete Choice Models of Product Differentiation: An Application to Medicare Part D with Switching Costs." *Mimeo*.

- Nederlandse Mededingingsautoriteit. 2011. *Sectorstudie Hypotheekmarkt: Een onderzoek naar de concurrentieomstandigheden op de Nederlandse hypotheekmarkt*. Accessed December 8, 2016. <https://www.acm.nl/nl/publicaties/publicatie/7091/Sectorstudie-Hypotheekmarkt/>.
- Nosal, Kathleen. 2012. "Estimating Switching Costs for Medicare Advantage Plans." *Mimeo*.
- Obozinski, Guillaume, Ben Taskar, and Michael I. Jordan. 2010. "Joint covariate selection and joint subspace selection for multiple classification problems." *Statistics and Computing* 20 (2): 231–252.
- Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)." *RAND Journal of Economics* 38 (2): 373–399.
- Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. "Asymptotic least squares estimators for dynamic games." *Review of Economic Studies* 75 (3): 901–928.
- Petrin, Amil, and Kenneth Train. 2010. "A control function approach to endogeneity in consumer choice models." *Journal of Marketing Research* 47 (1): 3–13.
- Phelps, Edmund S., and Sydney G. Winter Jr. 1970. "Optimal Price Policy under Atomistic Price Competition." In *Microeconomic Foundations of Employment and Inflation Theory*. New York: WW Norton Co, edited by Edmund S. Phelps, 309–337. New York: W. W. Norton & Co.
- Raval, Devesh, and Ted Rosenbaum. Forthcoming. "Why Do Previous Choices Matter for Hospital Demand? Decomposing Switching Costs from Unobserved Preferences." *Review of Economics and Statistics*.
- Rhodes, Andrew. 2012. "History-based price discrimination in markets with switching costs: who gains and who loses?" *Mimeo*.
- Rickert, Dennis. 2016. "Consumer Switching Costs, Dynamic Pricing, and Antitrust Analysis. Evidence from the Market for Diapers."
- Romano, Joseph P., Azeem M. Shaikh, and Michael Wolf. 2014. "A practical two-step method for testing moment inequalities." *Econometrica* 82 (5): 1979–2002.
- Shcherbakov, Oleksandr. 2016. "Measuring consumer switching costs in the television industry." *The RAND Journal of Economics* 47 (2): 366–393. doi:10.1111/1756-2171.12131.
- Stole, Lars A. 2007. "Price Discrimination and Competition." In *Handbook of Industrial Organization*, edited by Mark Armstrong and Robert K. Porter, 2221–2299. Elsevier.

- Taylor, Curtis R. 2003. "Supplier surfing: Competition and consumer behavior in subscription markets." *RAND Journal of Economics*: 223–246.
- Thiel, Jurre H. 2018. "Fee-based versus commission-based advice." *Mimeo*.
- Viard, V. Brian. 2007. "Do switching costs make markets more or less competitive? The case of 800-number portability." *RAND Journal of Economics* 38 (1): 146–163. doi:10.1111/j.1756-2171.2007.tb00049.x.
- Weiergräber, Stefan. 2017. "Network effects and switching costs in the US wireless industry: Disentangling sources of consumer inertia." *Mimeo*.
- Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2008. "Markov perfect industry dynamics with many firms." *Econometrica* 76 (6): 1375–1411.
- Wollmann, Thomas G. 2018. "Trucks without bailouts: Equilibrium product characteristics for commercial vehicles." *American Economic Review* 108 (6): 1364–1406.
- Zhang, Jian. 2006. "A Probabilistic Framework for Multi-Task Learning." PhD diss., Carnegie Mellon University.

A Testing the implications of SMPE

In this appendix, I further describe how the implications of SMPE can be used to test the model. If the estimated model is correct, equation (10) holds for every combination σ_i and σ_j such that bank b pays attention to state variable i but not j . Moreover, such a set of inequalities exists for every bank. Since equation (10) only contains quantities that can be estimated or that can be computed based on the estimated demand and supply model, it is easy to test whether it holds in the data. To that end, I note that (10) can be interpreted as a moment inequality. To test the moment inequalities I use the two-step bootstrap of Romano, Shaikh, and Wolf (2014). Full details, including how to calculate the right-hand side of (10), can be found in Appendix B.4.

In Table 13, I test whether the attention vectors identified by the multi-task Lasso are consistent with sparse maximization. Although, as I argued in Section 6.1, the Lasso identifies banks' attention vectors, this does not imply that the Lasso selects banks' true attention vectors with probability one. There are two reasons for this. The first is that the Lasso only selects the true non-zero coefficients as the sample size goes to infinity (Belloni and Chernozhukov 2013). In a finite sample, the Lasso can select the "wrong" state variables. The second reason is that identification of the attention vectors is dependent on using the correct functional form (or a superset thereof) in the Lasso regression. There are usually multiple, isomorphic, ways to define the pay-off relevant state and the true policy functions may depend on the state in a highly non-linear way. Therefore, it is easy to misspecify the functional form and it is important to check whether the selected model is consistent with sparse maximization. This also shows the practical benefit of using SMPE: in addition to

Table 13: Consistency of estimated policy functions with sparse maximization.

	p -value	Fraction of violated moment inequalities
Bank A	.270	12.5%
Bank B	.000	25%
Bank C	1.000	0%
Bank D	.000	33%
Bank E	1.000	0%
Bank F	.000	50%

Note: The first column contains p -values of a test that the estimate model is consistent with sparse maximization. The reported p -value is the largest significance level of the two-step bootstrap test of Romano, Shaikh, and Wolf (2014) for which the null cannot be rejected, where the size of the of the first stage bootstrap is set equal to 10% of the overall size of the test. The second column displays the fraction of sample moments for which (10) doesn't hold.

providing a micro-foundation for using the Lasso during the policy function estimation, it provides a method to check whether the statistical selection of attention vectors has a well-defined economic interpretation.

Table 13 shows mixed evidence that the estimated model is consistent with sparse maximization. For three of the six banks in my sample, I cannot reject the null that the bank performs sparse maximization at conventional sizes—for two of those banks, it is even impossible to reject this null at *any* size. For the remaining three banks, I can reject the null of sparse maximization at virtually any size: the p -values of the associated tests are very close to zero. However, it should be noted that the test of Romano, Shaikh, and Wolf (2014) rejects the null whenever only one the inequalities (10) is sufficiently violated. The test does not provide information on the extent of the violation. Therefore, I provide in the second column of Table 13 the fraction of moment inequalities that are violated in my estimated model. As can be seen, the violations are relatively infrequent. On the whole, if a bank pays attention to one state variable but not another, more often than not this is consistent with sparse maximization. The most common violation is that a bank is estimated to not pay attention to its own market share, but that it does pay attention to another bank's market share.

To check the robustness of the model to the fact that the Lasso does not select attention vectors that are fully compatible with sparse maximization, I re-estimate the model choosing the Lasso penalty parameter λ to maximize the p -values of the test of consistency with sparse maximization, instead of through cross-validation. In other words, I choose the penalty parameter to maximize the model's structural interpretation instead of the out-of-sample fit. Since the most common violation of sparse maximization is that banks do not pay attention to their own market share, I impose that banks must pay attention to their own market share. The other variables I include are the outcome of the multi-task Lasso procedure as in Section 6.1.2, except that I choose the penalty parameter λ to maximize the p -values of the null that the model is consistent with sparse maximization. Using this procedure, I get p -values of 1 for all banks, that is a model that is perfectly compatible with sparse maximization. Table 14 compares the results of this method with the results obtained cross-validation: they are very similar. Thus, it is not necessary to obtain a model that is perfectly

Table 14: Comparison of the estimated effect of the ban on history-based price discrimination for different methods of choosing the Lasso penalty parameter.

Effect of ban on:	Consumer surplus	Profits	Welfare
Using cross-validation	415	-279	125
Maximizing the model's structural interpretation	385	-248	137

Note: The values are in euro per year for a mortgage of €150,000. The model estimated using cross-validation is the model described in the main text. When maximizing the model's structural interpretation, one instead estimates the model various values of the Lasso penalty λ and chooses the one with the highest p -value on the specification check described in Appendix A.

consistent with sparse maximization to obtain reasonable results.

B Estimation details

B.1 Matching households over time

To match switching households over time, I employ the following algorithm. First, I discard all households that do not switch between the LLD in year t and year $t + 1$. I can exactly identify these households because every bank uses a unique scheme to identify households over time.⁶¹ For the remaining loans, I calculate the distance between all loans in the old and the new LLD. The distance between loan i in year t and loan s in year $t + 1$ is the norm between its loan and household characteristics, i.e.

$$d(i, j) = \|X_i - X_j\|,$$

where X_i are the standardized characteristics of loan i . As characteristics I take the birth year of the primary borrower, the payment type of the loan, the outstanding balance at the moment of switching and the maturity year of the loan. I then assign loan i as being loan j 's previous loan with probability

$$\frac{\exp\{d(i, j)\}}{\sum_k \exp\{d(k, j)\}}$$

where the sum in the denominator is over all loans in year t that I cannot match based on loan id in year $t + 1$. I further adjust the matching probabilities such that the aggregate probability of switching equals the switching probability I observe in the DHS and such that the correct proportion of loans in year $t + 1$ is not assigned any previous loan, i.e. is a new mortgage.

61. The only exception is ABN Amro, which switches its identifying scheme once. I match those mortgages using the same algorithm.

B.2 Supply side first order conditions

The derivatives of bank b 's profits for a product k with respect to its interest rate for old and new customers are

$$\begin{aligned}\frac{\partial \pi_b}{\partial r_{k0}}(r_b, s) &= \sum_{j \in \mathcal{J}_b} \left((r_{k0} - i_t - \gamma_j) \frac{\partial D_j}{\partial r_{k0}}(b, \psi, r_b, \hat{\rho}^b(s)) \right) + D_k(b, \psi, r_b, \hat{\rho}^b(s)) \\ \frac{\partial \pi_b}{\partial r_{k1}}(r_b, s) &= \sum_{j \in \mathcal{J}_b} \sum_{d \neq b} \left((r_{j1} - i_t - \gamma_j) \frac{\partial D_j}{\partial r_{k1}}(d, \psi, r_b, \hat{\rho}^b(s)) \right) + \sum_{d \neq b} D_k(d, \psi, r_b, \hat{\rho}^b(s)).\end{aligned}$$

(When banks can only set a single interest, the derivative of profits with respect to that interest is the sum of these two terms.) The market share of bank $d \neq b$, if it is in bank b 's sparse state representation, evolves as follows

$$\begin{aligned}\frac{\partial \Gamma_d}{\partial r_{k0}}(r_b, s) &\propto \sum_{j \in \mathcal{J}_d} \frac{\partial D_j}{\partial r_{k0}}(b, \psi, r_b, \hat{\rho}^b(s)) \\ \frac{\partial \Gamma_d}{\partial r_{k1}}(r_b, s) &\propto \sum_{j \in \mathcal{J}_d} \sum_{d' \neq b} \frac{\partial D_j}{\partial r_{k1}}(d', \psi, r_b, \hat{\rho}^b(s)).\end{aligned}$$

I ignore a constant that measures the size of tomorrow's market versus today's market since it will be subsumed by the Lagrange multipliers μ .

Substituting these expressions into (8) gives that bank b 's first order conditions in month t can be written as

$$\begin{aligned}\sum_{j \in \mathcal{J}_b} \left((r_{k0} - i_t) \frac{\partial D_j}{\partial r_{k0}}(b, \psi, r_b, \hat{\rho}^b(s)) \right) + D_k(b, \psi, r_b, \hat{\rho}^b(s)) &= \sum_{j \in \mathcal{J}_b} \gamma_j \frac{\partial D_j}{\partial r_{k0}}(b, \psi, r_b, \hat{\rho}^b(s)) + \mu_{bt}^T \frac{\partial \hat{\Gamma}}{\partial r_{k0}}, \\ \sum_{j \in \mathcal{J}_b} \sum_{d \neq b} \left((r_{j1} - i_t) \frac{\partial D_j}{\partial r_{k1}}(d, \psi, r_b, \hat{\rho}^b(s)) \right) + \sum_{d \neq b} D_k(d, \psi, r_b, \hat{\rho}^b(s)) &= \sum_{j \in \mathcal{J}_b} \sum_{d \neq b} \gamma_j \frac{\partial D_j}{\partial r_{k1}}(d, \psi, r_b, \hat{\rho}^b(s)) + \mu_{bt}^T \frac{\partial \hat{\Gamma}}{\partial r_{k1}}.\end{aligned}$$

The left side of these equations can be calculated given the estimated demand model and bank b 's policy functions. The right side is composed of quantities that can similarly be calculated and unknown parameters γ_j, μ_{bt} . The right hand side is linear in these parameters, therefore the supply side first order conditions give rise to a linear system of equations.

B.3 Calculating a bank's best response

To find the interest rates r that maximize (9), I approximate the value function by a complete product of Chebyshev polynomials of degree 4. I use twenty nodes per state variable. I then use the following algorithm to calculate a bank's best response:

1. Initialize $i = 1$ and $V_b(\sigma_0) = \sum_{t=0}^{\infty} \beta^t \pi(\hat{r}(\sigma_t), \sigma_t)$. Here, $\hat{r}(\cdot)$ are the estimated policy functions of all banks and $\sigma_{t+1} = \Gamma(\hat{r}(\sigma_t), \sigma_t)$. That is, I initialize bank b 's value function as the value function it would obtain when all banks would set interest rates according to their estimated policy functions forever.

2. For every node in the basis of the Chebyshev polynomials, $\hat{\sigma}_b$, calculate bank r_i to maximize

$$r_i = \arg \max_r \pi_b(r, \hat{r}_{-b}(\hat{\sigma}_b), \hat{\sigma}_b) + \beta V_i(\Gamma(r, \hat{r}_{-b}(\hat{\sigma}_b), \hat{\sigma}_b)).$$

3. Calculate bank b 's value function

$$V_{i+1}(\hat{\sigma}_b) = \pi_b(r_i, \hat{r}_{-b}(\hat{\sigma}_b), \hat{\sigma}_b) + \beta V_{i+1}(\Gamma(r_i, \hat{r}_{-b}(\hat{\sigma}_b), \hat{\sigma}_b)).$$

As when calculating the costs of inattention (Appendix B.4), this can be done using a contraction mapping.

4. Terminate if

$$\sup \left| \frac{V_{i+1} - V_i}{1 + V_{i+1}} \right| < \epsilon.$$

Otherwise, increase i by one and go to step 2.

B.4 Calculating the cost of inattention

Here, I describe how to calculate the cost of inattention required to calculate the critical value of the moment inequalities (10). This quantity depends on an estimated model: since the choice of λ influences the demand estimates (through the control function) as well as the marginal cost estimates (through the sparsity structure of banks' states), both the demand and supply model have to be calculated for each value of λ . Given the model, perform the following steps to calculate $r_{\sigma_i}^T v_{rr} r_{\sigma_i}$ for a given bank b

1. **Calculate the objective value function.** To calculate the test, the objective value function is required. The objective value function measures the actual profits a bank makes given its actions. That is,

$$V_b(\sigma) = \pi_b(\rho_b(\sigma), \rho_{-b}(\sigma), \sigma) + \beta \mathbb{E}_{i'} [V_b(\Gamma(\sigma, \rho_b \sigma, \rho_{-b}(\sigma)), i')].$$

Given the policy functions $\rho(\cdot)$, the value functions be calculated using the following mapping T :

$$V_b = T(V_b) = \pi_b(\rho_b(\sigma), \rho_{-b}(\sigma), \sigma) + \beta \mathbb{E}_{i'} [V_b(\Gamma(\sigma, \rho_b \sigma, \rho_{-b}(\sigma)), i')].$$

Gabaix (2017, Lemma 3.6) implies this is a monotone contraction, so calculating this is easy.

To implement this, I approximate $V_b(\sigma)$ by a complete product of Chebyshev polynomials of degree four on a grid consisting of the Cartesian product of twenty Chebyshev nodes per state variable. The flow profits $\pi_b(\cdot)$ can be calculated using the estimated demand, marginal costs and policy functions.

Denote

$$v_b(r_b, r_{-b}, \sigma) = \pi_b(r_b, r_{-b}, \sigma) + \beta V_b(\Gamma(r, r_{-b}, \sigma)).$$

2. **Calculate the default action.** Recall that the default action contains the interest rates a bank

sets in the default state. In the default state, the bank assumes every state variable equals its long-run average in every period. The default action of bank b is

$$r_b^d = \arg \max_r v_b(r, \rho_{-b}(\sigma^d), \sigma^d).$$

The actions of b 's competitors can be calculated using the estimated policy functions.

3. **Calculate r_{σ_i} and u_{rr} .** r_{σ_i} is the derivative of b 's optimal interest rate with respect to state variable σ_i . Since the first order condition is

$$(r - c) \frac{\partial D}{\partial r} + D + \beta \frac{\partial V(\Gamma)}{\partial \sigma} \frac{\partial \Gamma}{\partial r} = 0,$$

the implicit function theorem gives that

$$\frac{\partial r}{\partial \sigma} = (r - c) \frac{\partial^2 D}{\partial r \partial \sigma} + \frac{\partial D}{\partial \sigma} + \beta \frac{\partial V(\Gamma)}{\partial \sigma} \frac{\partial^2 \Gamma}{\partial r \partial \sigma} + \beta \frac{\partial \Gamma}{\partial r} \frac{\partial^2 V(\Gamma)}{\partial \sigma^2} \frac{\partial \Gamma}{\partial r}.$$

(I suppress arguments and the subscript b for legibility.) Simply differentiating the definition of $v(\cdot)$ above twice gives

$$v_{rr} = \frac{\partial^2 r}{\partial r^2} + \beta \frac{\partial \Gamma}{\partial r} \frac{\partial^2 V(\Gamma)}{\partial \sigma^2} \frac{\partial \Gamma}{\partial r} + \beta \frac{\partial V(\Gamma)}{\partial \sigma} \frac{\partial^2 \Gamma}{\partial r^2}.$$

Note that both r_{σ_i} and v_{rr} are evaluated at the default action.