Competition and Dynamic Pricing in Markets with Consumer Switching Costs*

Jurre H. Thiel†

January 21, 2020

Abstract

In markets with switching costs, prices increase as firms' market shares grow. I study the effect of entry on these price dynamics in the Dutch mortgage market. I exploit incumbents' multi-product nature to control for unobserved firm-level differences, while a difference-in-differences strategy controls for market-level shocks by exploiting that only some mortgage products faced increased competition. Consistent with a simple theoretical framework, banks with more locked-in customers decrease their interest rates less after entry than smaller incumbents. This can explain why policies that encourage competition often disappoint in markets with switching costs. I also discuss implications for competition policy.

JEL Classification: D43, L11, L13

1 Introduction

A central question in industrial organization concerns the relationship between competition and prices. In standard frictionless models, an increase in the number of firms in a market causes prices to decrease. However, in a market with consumer switching costs, firms employ dynamic pricing strategies as they trade off charging low prices to attract new customers (investment) and charging high prices to their existing, locked-in customers (harvesting).[1] A large theoretical literature (e.g. Beggs and Klemperer 1992), as well as empirical evidence (Carbo-Valverde, Hannan, and Rodriguez-Fernandez 2011; MacKay and Remer 2019), shows that as a result, a firm's price increases in its number of locked-in customers. Hence, as opposed to frictionless markets, the question in markets with switching costs is not only how an increase in the number of firms influences the average price, but potentially also how it affects the relationship between the amount of locked-in customers and prices.

*I am grateful to Marcus Asplund, Aljoscha Janssen, Andrew Rhodes and Xuan Zhang for useful comments and José Luis Moraga-González and Mauro Mastrogiacomo for their comments and guidance. This paper has also benefited from presentations at EARIE 2019 in Barcelona. 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.

†Copenhagen Business School. juth.eco@cbs.dk

1. This trade-off exists if price discrimination between old and new customers is not possible, an assumption I maintain throughout this paper. This assumption reflects my empirical setting, as I discuss in more detail in Section 3.

It is important to understand the effects of competition on these dynamic pricing strategies for the following reasons. First, policy makers often implement policies to encourage competition in markets with switching costs. A concrete example is the introduction of retail competition in electricity, opening up former local monopolies in electricity provision and allowing consumers to choose their provider. However, consumers tend not to switch away from the incumbent electricity provider due to switching costs (Hortaçsu, Madanizadeh, and Puller 2017). As a result, the benefits from increased competition derive primarily from incumbents' reactions. Hence, it is crucial to understand how incumbents respond to entry to evaluate the merits of policies that aim to increase competition. Similar issues arise in health insurance, where policy makers around the world have tried to foster competition (Gaynor, Ho, and Town 2015). However, recent empirical evidence shows that in these markets consumers display significant inertia (Handel 2013). The second importance is in competition policy. Competition authorities frequently cite consumer lock-in as a concern.[2] Yet, it is unclear how competition affects prices in such markets, which is important for merger analysis. A related issue is the question whether switching costs form a barrier to entry (see e.g. Farrell and Shapiro 1988), which has implications for a market's contestability.

This paper provides novel theoretical and empirical results on the effect of an increase in the number of firms on incumbents' pricing strategies. It does this by studying the effect of pension funds' entry into the Dutch mortgage market on incumbents' interest rates. In this market, households tend to fix their interest rates, most commonly for a period of ten years. After this fixed interest rate period ends, a household's current bank offers it to renew its mortgage, that is to fix its interest rate again. At this point a household can also switch to a different bank. However, switching is costly: the largest mortgage broker in the country reports that on average a switching household incurs about €3,700 in fees, in addition to the opportunity cost of time.[3] Since empirical evidence suggests that pension funds' entry into this market was spurred by changing regulations governing pension funds (Kim and Mastrogiacomo 2019), the timing of entry is plausibly exogenous.

My setting allows me to disentangle the effect of entry on the average price, i.e. the intercept of firms' policy functions, from the effect of entry on the relation between customer lock-in and price, i.e. the slope of firms' policy functions. The feature of the market that enables this distinction is that mortgage providers sell multiple mortgage products, differing in the way in which the principal is repaid. Since consumers typically do not switch between different repayment methods, even if they switch mortgage providers, there are in essence *multiple* mortgage markets in the Netherlands. Entrants sell only a subset of mortgage products offered by incumbents. Hence, some mortgage products face increased competition, while

2. For example, both in the US and the EU, the alleged difficulty of switching to a different browser was central to competition authorities' arguments in their cases against Microsoft (Edlin and Harris 2012). MacKay and Remer (2019) discuss some other recent cases in the US where switching costs were explicitly mentioned by competition authorities.

3. De Hypotheker. `https://www.hypotheker.nl/jouw-woonsituatie/hypotheek-oversluiten/`. Accessed Mar 11, 2019. The broker estimates notary costs of €750, taxation costs of €500, costs of removing the original mortgage from the official mortgage registry of €150, costs for government insurance of the mortgage of €2,000 and advisory costs of €2,990. This comes to a total of €6,390. These costs are however tax-deductible. Since the average marginal tax rate is 42%, this means monetary switching costs are around €3,700. In addition, switching takes time and effort. Thiel (2018) uses a structural model to estimate switching costs for this market and finds that they are substantial.

others do not, resulting in clearly defined treatment and control markets and a difference-in-differences strategy identifies the average effect of entry on interest rates. To estimate the effect on the slope of incumbents' policy functions, I estimate the interaction between an incumbent's response to entry and its market share among renewing customers. Here, I face the issue that incumbents' market shares are likely correlated with unobserved firm characteristics, such as marginal costs or quality, that also influence their interest rates.[4] I solve this issue by exploiting cross-sectional differences between different treated mortgage products sold by the same provider in the amount of locked-in customers. Under the assumption that mortgage products sold by the same provider face similar supply and demand shocks, differences in response to entry are then entirely due to differences in the amount of consumer lock-in. To the extent this assumption holds, I hence control for any persistent or transitory vertical differentiation or cost differences.

My empirical results imply that an increase in the number of firms amplifies incumbents' dynamic pricing incentives. In my preferred specification, an incumbent without any locked-in customers decreases its interest rate with 12 basis points, or 4.5%, relative to the control market. For every 10 percentage points that its market share amongst renewing customers increases, the decrease is 3 basis point less. Hence, incumbents with a small number of locked-in customers decrease their prices more than larger incumbents, reflecting the changing incentives firms have to attract new customers or to exploit their existing customers.

To interpret these results, as well as to elucidate the identification challenges of my empirical exercise, I provide comparative statics from a simple model to understand the effects of competition on prices in markets with switching costs. Surprisingly, this issue has not been studied in the theoretical literature, perhaps owing to the focus of this literature on duopolies. Building on the model of Somaini and Einav (2013), I show that an increase in the number of firms active in a market makes the relation between firms' past market shares and prices steeper. Hence, consistent with my empirical results, in the period immediately after a new firm enters the market, firms with smaller customer bases decrease their prices more than firms with larger customer bases. The intuition for this result is that competition strengthens *both* the investment and the harvesting motive and that therefore, an incumbent's reaction to entry depends on the amount of locked-in customers it has.

My results provide one explanation why policies to foster competition do not always live up to expectations. Continuing the example of retail competition in electricity markets, Defeuilley (2009) notes in a review of this policy around world that while it was expected to lower prices, "the expected results did not always materialize" (p. 377). The results in this article indicate that, on the contrary, incumbents not decreasing their prices is exactly what should have been expected given that retail electricity is a market characterized by switching costs: the larger a firm's market share, the less likely it is to reduce prices after an increase in competition. In the context of electricity markets, the incumbent is typically a former state mo-

4. Say firm i experiences a positive and permanent demand shock at time t. Then, because the shock causes expansion of demand at time t, firm i's number of locked-in customers will be relatively high at time $t+1$. At the same time, firm i will charge a higher price in periods $\tau \geq t$ than before. Hence, unobserved shocks to product quality create a positive correlation between past market shares and prices even absent switching costs. A similar argument shows that shocks that decrease marginal costs create a negative correlation between prices and market shares.

nopolist. Hence, it has a market share of 100% around the introduction of competition and this logic holds in the extreme. More broadly, my results suggest that there is negative selection of markets in which policies to encourage competition are introduced: policy makers are likely to see a larger need to stimulate competition in more concentrated markets, but insofar these markets are subject to switching costs, these are exactly the markets where such policies can be expected to have the least impact.

It should be stressed from the outset that my empirical results concern the effect of entry on incumbents' prices in the periods immediately following entry, and not in the steady state.[5] However, if these price changes are due to a rotation of incumbents' policy functions, as my theoretical results suggest, my results provide insights into the changing nature of competition in all periods following entry, up to and including the steady state. Moreover, markets with switching costs need not reach a steady state, but can also be characterized by price cycles (Janssen 2019) and empirical results show that it can take a long time to reach a new steady state.[6] Hence, if the social planner discounts future periods, the periods immediately after entry will form a large part of any welfare calculation and it is these periods that I observe in my data.

This paper contributes to a large theoretical and empirical literature on switching costs. The theoretical literature on switching costs (Klemperer 1987b; Farrell and Shapiro 1988; Beggs and Klemperer 1992; Padilla 1995; To 1996; Doganoglu 2010; Arie and Grieco 2014; Rhodes 2014; Fabra and García 2015; Cabral 2016; Pearcy 2016; Ruiz-Aliseda 2016) has mostly focused on the competitive effects of the *level* of switching costs for a given number of firms (typically two), but has not looked at the effects of competition in markets where switching costs are relevant. Theoretical papers that do explicitly consider entry are Klemperer (1987a, 1988), Farrell and Shapiro (1988), and Klemperer (1989). There are two important differences with this prior literature. The first is that the aforementioned papers all consider the case where the incumbent is a monopolist, while in my framework and empirical setting there is also competition before additional firms enter. My results suggest that the prior focus on monopolist incumbents has been restrictive, as smaller incumbents react very differently to entry than larger ones. A second difference is that this literature focuses primarily on the question whether switching costs should be seen as a barrier to entry. This paper however is more on the effects of entry per se. Nevertheless, my results do offer the new insight that pre-entry market concentration influences whether switching costs are a barrier to entry, as I discuss in Section 6. I am unaware of any empirical work that studies the effect of competition in markets with switching costs.

This paper also contributes to a rich empirical literature documenting the effect of competition on prices.[7] With the exception of Allen, Clark, and Houde (2014) and Lach and Moraga-

5. I do not take a stand whether the market is in a steady state in the periods immediately *before* entry.

6. For example, Hortaçsu, Madanizadeh, and Puller (2017) find that in retail electricity markets, the incumbent enjoys a significant advantage even four years after entry of new firms.

7. Bresnahan and Reiss (1991) look at the effect of competition on prices for a wide array of industries. In retail markets Basker (2005), Hausman and Leibtag (2007), Jia (2008), Basker and Noel (2009), and Ailawadi et al. (2010) focus specifically on the effects of Walmart's entry into a market on prices. The focus of Lira, Rivero, and Vergara (2007) is on the Chilean supermarket sector. Abe and Kawaguchi (2010) consider the Japanese supermarket sector. In the US, Courtemanche and Carden (2014) look at the effect of Costco and Sam's Club, while Hosken, Olson, and Smith (2018) study the US supermarket sector more broadly. Russell (Forthcoming) considers the effect of a reduction on competition, through horizontal mergers, on college tuition fees in the United States. The effect of competition on prices has also been studied in developing countries (Atkin, Faber, and Gonzalez-Navarro 2018; Bennett and Yin

González (2017), who study the effect of competition in markets with search costs, this literature has not explicitly considered the role of frictions or implicitly assumed that markets are frictionless. I add to this literature by studying the effect of competition on prices in markets with switching costs. My theoretical framework gives reasons to believe that the effect of competition might be different in such markets than in frictionless markets. Since switching costs have been widely documented (e.g Handel 2013; Honka 2014; Shcherbakov 2016; Raval and Rosenbaum, Forthcoming), it is important to extend the existing empirical literature to this setting.

Finally, the results in this paper provide empirical support for the existence of the investment-harvesting trade-off. The existence of this trade-off is a central result of the theoretical literature, as well as a key assumption of recent structural models of pricing in markets with switching costs (e.g. Cosguner, Chan, and Seetharaman 2016; Janssen 2018; Thiel 2018; MacKay and Remer 2019). There is some prior empirical support for the existence of this trade-off, based on testing the implication that in markets with more unattached customers, prices should be lower. Hannan and Adams (2011) test this implication for US bank deposits, while Carbo-Valverde, Hannan, and Rodriguez-Fernandez (2011) do so for Spanish bank deposits.[8] The latter paper additionally tests whether banks with larger market shares indeed offer lower deposits rates (which is the equivalent to setting higher prices in their setting). When doing so, they use an instrumental variables strategy to control for unobserved firm characteristics that influence both deposit rates and market shares. Relative to this prior literature, I make two contributions. First, I test a different implication of models of pricing with switching costs, namely that dynamic pricing incentives become amplified when the number of firms increases. Second, and perhaps more importantly, a significant advantage of my setting is the possibility to control for unobserved firm characteristics. Like Carbo-Valverde, Hannan, and Rodriguez-Fernandez (2011) and Hannan and Adams (2011), I find empirical support for firms employing dynamic pricing strategies in markets with switching costs.

This paper proceeds as follows. Section 2 lays out my theoretical framework. Section 3 gives more background on the Dutch mortgage market and the recent entry into this market. Section 4 describes my data set. In Section 5 I develop my empirical strategy and test the theory from Section 2. Section 6 discusses the implications of my results for economic and competition policy. Section 7 concludes.

2 The effects of entry in a simple model of switching costs

To frame my empirical results, I employ the model of Somaini and Einav (2013). This model is relatively rich compared to much of the theoretical literature on switching costs while still being highly tractable. Because I focus on the effects of competition, it is particularly attractive that the model of Somaini and Einav (2013) allows for an arbitrary number of firms. Most of

2019; Busso and Galiani 2019).

8. Stango (2002) also studies pricing in markets with switching costs, but contrary to this paper and the cited papers, does so in a setting where firms can charge different prices to new and existing customers.

the existing theoretical literature deals with model of duopoly.[9]

2.1 Base model

I now briefly describe the setup of Somaini and Einav (2013). I do not alter or extend their model but will just point out the implications of entry in their setting. Therefore, I highlight only the model's most relevant features and refer the reader to the original paper for more details.

There are N firms in the market, each selling a horizontally differentiated product. Differentiation is modeled spatially: the products are located at the vertices of the $(N-1)$-dimensional simplex. For example, for $N = 3$, the products are located at the corners of an equilateral triangle. This model of product differentiation extends the Hotelling (1929) linear city model to higher dimensions and is isomorphic to the spokes model of Chen and Riordan (2007).

Consumers live for two periods. At the beginning of every period, the mass of old consumers is normalized to one. In addition, in every period g new, unattached consumers are born. Note that g can be smaller than, larger than or equal to one, so that the size of the market can grow, shrink or stay constant over time.

Consumers are located uniformly on the $\frac{N(N-1)}{2}$ edges of the simplex. In the first period, a consumer is assigned a location on one of the edges of the simplex with uniform probability. By assumption, a consumer only considers the two firms located on his own edge. To purchase a certain product, a consumer pays the product's price and incurs a linear transportation cost proportional to the distance he has to travel along the simplex' edge. The market is assumed to be covered, so all consumers make a purchase.

In the second period of a consumer's life, his position is redrawn in the following manner. He receives a new, random position on the same edge of the simplex as he inhabited in the first period. The consumer therefore considers the same two products as in the first period. He however faces a switching cost s if he purchases from a different firm than before. Therefore, in the second period the consumer purchases the product that minimizes the sum of the product's price, transportation cost and (if applicable) switching costs.

Consumers are forward-looking and have correct beliefs on firms' pricing strategies in equilibrium. Hence, in the first period, the consumer purchases the product that minimizes his total discounted costs. Consumers employ a discount rate $\delta_c \in [0, 1)$. The model thus nests the case of myopic consumers for $\delta_c = 0$.

Firms are infinitely-lived and maximize their discounted payoffs. A firm faces two sources of demand. The first is from old consumers. Because a firm's demand is less elastic when it has more locked-in consumers, the demand from old consumers depends on past market shares $(x_i)_{i=1,\dots,N}$, where $\sum_{i=1}^{N} x_i = 1$. In addition, a firm sells to the g newly-born, unattached consumers. Firms discount future profits with a discount rate r_f. It is assumed that firms' effective discount rate $\delta_f \equiv r_f g$ is smaller than one. I normalize marginal costs to zero for now to highlight the main intuition from the model. I discuss the effects of asymmetric marginal costs briefly below, and in more detail in Appendix B.

9. Two exceptions are Arie and Grieco (2014) and Pearcy (2016), who also allow for an arbitrary number of firms. However, their models are significantly less tractable than that of Somaini and Einav (2013).

Even though this model is rich in various dimensions——it contains product differentiation, an arbitrary number of firms and an infinite horizon—it has a surprisingly simple Markov Perfect Equilibrium (MPE):[10]

Proposition 1 (Somaini and Einav (2013))**.** *There exists a Markov Perfect Equilibrium in which every firm has the policy function*

$$p_i^* = \alpha + \beta x_i, \tag{1}$$

where $\alpha, \beta > 0$ are constants that depend on the model's parameters and where x_i is the firm's market share amongst old consumers.

Somaini and Einav (2013) moreover show that this is the unique MPE in parallel linear strategies.[11] Not only is a firm's optimal strategy linear, it also depends on just its own past market share rather than the full vector $(x_i)_{i=1,\dots,N}$. This is particularly useful when studying competition, because it means that the effective dimension of the game does not depend on the number of firms in the market. However, because β is the solution to a fourth-degree polynomial (Somaini and Einav 2013, p. 975), it is not possible to derive useful closed-form solutions for α and β.

Before looking at the effects of the number of firms in this model, it is instructive to study the structure of this policy function. The ratio $\frac{\beta}{\alpha}$ essentially measures the strength of dynamic pricing incentives. If this ratio is large, firms with many locked-in customers charge significantly higher prices than firms with lower past market shares. Indeed, the ratio of firm i's price and firm j's price can be written as

$$\frac{p_i}{p_j} = \frac{1 + \frac{\beta}{\alpha} x_i}{1 + \frac{\beta}{\alpha} x_j}.$$

When $x_i > x_j$, this ratio is increasing in $\frac{\beta}{\alpha}$. As $\frac{\beta}{\alpha}$ goes to zero, the firms charge the same price even though i has the larger market share. As $\frac{\beta}{\alpha}$ increases, the price difference between firms with smaller and larger past market shares increases. That is, for larger firms the harvesting motive becomes relatively more important when $\frac{\beta}{\alpha}$ increases, while for smaller firms the investment motive increases in relative importance. Hence, $\frac{\beta}{\alpha}$ can be seen as a summary statistic that measures the degree to which firms engage in dynamic pricing.

I now study the effects of competition on firms' equilibrium strategies. First I consider the effect of the number of firms on firms' policy functions. Because the policy functions have the same functional form regardless of the number of firms, any effect of entry operates through the coefficients α and β. Since the solutions for α and β are very complex, it is difficult to analytically derive comparative statics with respect to the number of firms N. Therefore, I proceed with a numerical investigation. I calculate firms' optimal policy functions on a grid of 110,000 different parameter constellations (Table 1). Because I take eleven different values for N, this means that I investigate the effect of competition for 10,000 different environments.[12]

10. In the formulation in Somaini and Einav (2013) the slope α also differs from firm to firm. Since I assume that marginal costs are the same across firms, the policy functions reduce to the form shown here (see Somaini and Einav 2013, eq. 30).

11. A linear strategy profile is parallel if the coefficient on x_i is the same for all firms.

12. I thank Somaini and Einav (2013) for making their code for calculating MPE's publicly available.

Table 1: Parameter values for which the model is simulated

	Min value	Max value	Number of points
Switching cost (s)	.01	.99	10
Discount rate consumers (δ_c)	.01	.99	10
Discount rate firms (δ_f)	.01	.99	10
Growth rate market (g)	.025	2.475	10
Number of firms (N)	3, 4, 5, 6, 7, 8, 9, 10, 20, 50, 100		

Note: The unique Markov Perfect Equilibrium in parallel linear strategies is computed for the Cartesian product of the parameter ranges (110,000 equilibria in total). Intermediate points are taken equidistant between the endpoints.

For each environment, the following proposition holds true.

Proposition 2. *As the number of firms N increases, the intercept of firms' policy functions α strictly decreases and the slope of firms' policy functions β strictly increases.*

From this proposition it follows immediately that after entry $\frac{\beta}{\alpha}$ increases. In other words, competition increases dynamic pricing incentives.

Given the predicted effect of N on the coefficients α and β, it is possible to derive comparative statics on incumbents' prices. I do so in the following way. I take the established market shares of the incumbent firms $(x_i)_{i=1,\ldots,N}$ as given and look at the entry of firm $N+1$. An entrant does not have locked-in customers, i.e. $x_{N+1} = 0$. This means that I look at the effects of an increase in competition in the period right after entry has taken place. This reflects my empirical setting, in which I observe the same period.

The main prediction on incumbents' prices is the existence of a *market share-entry response gradient*. By observing the policy function (1), it is obvious that if α decreases and β increases there exists a cutoff market share x' such that firms increase their price if their market share is below x' and increase their price if their market share is above x'.

Proposition 3. *There exists a cut-off market share, such that incumbent firms with a customer base above this share increase their price after entry and incumbent firms with a customer base below this share decrease their price after entry.*

The intuition for this result is as follows. An entrant has no market share and hence no harvesting motive. As a result, it will price aggressively. An incumbent firm must then lower its price to increase its future market share, which decreases its margin on old consumers. Thus, entry increases the opportunity cost of investing in market share. For firms with large established customer bases, this opportunity cost may be so large that it may then become worthwhile to give up on attracting new customers and instead "feed off" existing consumers. For firms with low market shares this strategy is not an option: they have to decrease their price to compete with the aggressive entrant.

Because the effects on β are multiplied by the firms' market shares, firms with market shares further away from the cut-off market share change their price more after entry.

Proposition 4. *Among incumbent firms that increase their prices after entry, a firm with a larger customer base increases its price more. Among incumbent firms that decrease their prices after entry, a firm with a smaller customer base decreases its price more.*

Hence, the main prediction of the model is the existence of a market share-entry response gradient: a correlation between previous market shares and post-entry price responses. Denoting with Δx the post-entry change in quantity x, the effect of entry on firms i's price is

$$\Delta p_i = \Delta\alpha + \Delta\beta \cdot x_i \tag{2}$$

and the simulations predict $\Delta\alpha < 0$ and $\Delta\beta > 0$. It is these predictions that I take the data below.

2.2 Asymmetric competition

So far, I have assumed that although firms are horizontally differentiated, they are symmetric. When firms additionally differ in marginal costs or, equivalently, are vertically differentiated, a more complex picture emerges.[13] The reason is that a firm's marginal costs also determine how it reacts to an increase in competition, with, all other things equal, a firm with lower marginal cost responding more strongly. Nevertheless, in Appendix B I show that the main insight from the preceding discussion, that competition amplifies dynamic pricing incentives, holds, if i) the entrant's marginal cost are not above the industry average, and ii) there is not too much dispersion in marginal costs. Under these conditions, incumbents' responses to entry look something like Figure 1. When comparing two firms with the same marginal costs (in an otherwise asymmetric industry), the story is exactly as before: there is a cut-off market share at which incumbents do not respond to entry. Moreover, the further an incumbent's past market share is away from the cut-off, the more it will increase or decrease its price. Similarly, for a given market share, firms tend to increase their prices if they have high marginal costs and decrease them if their marginal costs are low.

The results from Appendix B imply that marginal costs and vertical differentiation are potential confounders when estimating the market share-entry response gradient. For example, firms with lower marginal costs tend to set lower prices for a given past market share. Since they tend price lower, they also have higher market shares. Hence, if all firms had the same customer base, there would be a negative correlation between prices and market shares. As a result, ignoring marginal costs will bias downwards the estimate of the market share-entry response gradient.

Since marginal costs and the amount of vertical differentiation are not observed in the data, it is a major empirical challenge to disentangle what part of a firm's entry response is due to its marginal costs and what part is due the amount of locked-in customers it has.[14] This is in particular true for my setting, as Thiel (2018) shows by structurally estimating a model of this market that there are large marginal cost differences between banks.[15] As I explain in more

13. The reason that vertical differentiation is equivalent to marginal cost differences follows from consumers' quasi-linear preferences. As a result, an increase in consumers' willingness to pay for a particular product is equivalent to a similarly sized decrease in the product's marginal costs. This is a standard result in theoretical industrial organization.

14. Accounting marginal cost data are available for my empirical application, the Dutch mortgage market, but only for the largest banks in the market. Since the goal of the empirical exercise is to look at how banks of different sizes respond to entry, this information is fairly useless. The amount of vertical differentiation is never directly observable in the data.

15. For computational reasons, Thiel (2018) also only estimates the marginal costs for the largest banks in the mar-

Figure 1: Incumbents' price responses to entry

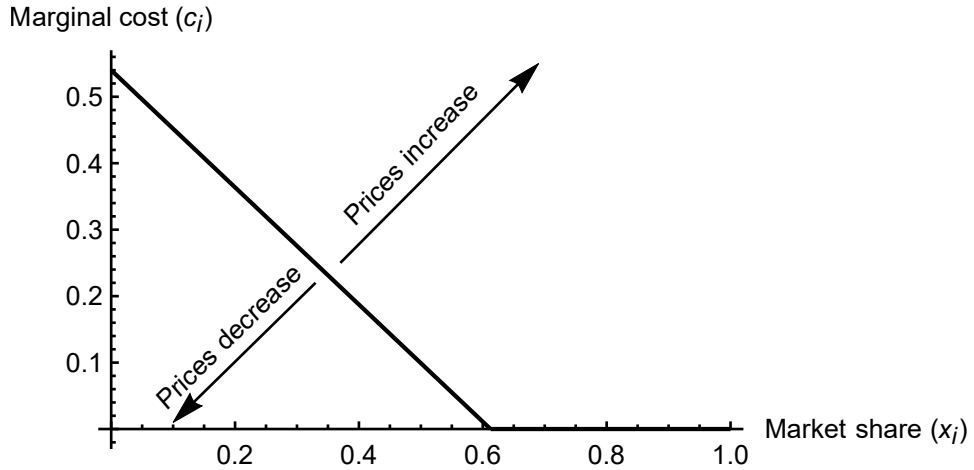

Note: The graph displays regions for which incumbents' prices increase and decrease as a function of an incumbent's established market share and marginal cost. The graph is plotted for an increase of $N = 3$ to $N = 4$ firms, switching costs $s = .1$, consumers' discount rate $\delta_c = .95$, firms' effective discount rate $\delta_f = .95$ and a constant market size over time ($g = 1$). The average marginal cost in the industry is $\bar{c} = .3$ and the entrant has marginal cost .3 as well, so that the average marginal cost does not change after entry.

Table 2: Entry in the Dutch mortgage market in 2014

Firm	Month of entry	Notes
Hypotrust	August 2014	Originally a mortgage broker; started to sell white-label mortgages financed by pension funds.
Tellius	September 2014	Financed by pension funds and insurance companies.
MUNT	November 2014	Financed by pension funds.
Post-entry market share of entrants		5.2%

Note: The table shows all entrants into the Dutch mortgage market in the year 2014. The month of entry is the origination date of the first mortgage in the Loan Level Data by the Dutch Central Bank. The post-entry market share of entrants is the entrants' share of newly originated mortgages in the period Aug 2014–Jul 2016.

detail below, I tackle this issue by comparing the interest rates a bank sets in a given month for different mortgage products, thereby differencing out unobserved firm characteristics.

3 The Dutch mortgage market

I test the predictions from Section 2 in the Dutch mortgage market. The Dutch mortgage market is an attractive testing ground for the theory for the following reasons. First, the Dutch mortgage market features significant switching costs. Second, as in the theoretical model, there is no price discrimination between new as to renewing customers. Third, in recent years there has been quite some entry. Fourth, this entry is caused by regulatory changes in other markets so that the timing of entry can be seen as plausibly exogenous. I now discuss these points in more detail.

In the Dutch mortgage market, most households fix their interest rates, typically for a period of ten years. However, the total duration of the mortgage is normally longer than

ket.

that—most mortgages are paid off in thirty years. When the fixed interest rate period ends, a consumer's current bank is legally required to offer a new fixed interest rate. At this point the consumer can also switch to a different bank without penalty. This does not mean that switching is free however, as various fees must be paid, for example reappraisal and notarial fees. The largest mortgage broker in the country estimates that monetary switching costs are around €3,700.[16] In addition, switching takes time and effort. Thiel (2018) estimates the total switching costs in a structural model of this market and finds that, even though there is significant heterogeneity in switching costs, switching costs are so high that the average household essentially never switches.

A second attractive feature of the institutional context is that it is illegal to engage in price discrimination based on purchasing history. That is, firms in this market are not allowed to charge lower interest rates to new than to renewing customers. This ban was instituted in 2013 to increase the competitiveness of the market. Thiel (2018) shows that while there was a small amount of non-compliance with this ban, banks indeed charge almost the same interest rate to new and renewing customers. While in many markets with switching costs, one might worry about (unobserved) price discrimination between new and renewing customers, this can therefore be ruled out here.

I study entry into the Dutch mortgage market around August 2014. Historically, the Dutch mortgage market has been reasonably concentrated. Three large banks—Rabobank, ABN AMRO and ING—dominate the market, all having market shares of over 20%.[17] Since the start of the financial crisis, there had been no entry and some smaller foreign players even exited the market. This all changed when in 2014 various new players entered. The source of entry was arguably quite surprising. Rather than (foreign) banks entering the market, Dutch pension funds set up new vehicles to sell mortgages. Table 2 displays the entrants in this period. Pension funds typically combined through so-called "servicers" that handle marketing and administration of the mortgages. Two mortgages sold by the same entrant can thus be financed by a different institution. However, this process is transparent to the consumer. The entrants were relatively successful, obtaining a combined market share of 5.2% in the period August 2014-July 2016.

Pension funds seem to have entered the Dutch market because of changes in the regulations governing their balance sheets. The new Financial Assessment Framework ("Financieel Toetsingskader" or FTK in Dutch) caused a flight to safe assets. Since repayment rates in the Dutch mortgage market are high, regulators view investments in the Dutch mortgage market as relatively safe. Kim and Mastrogiacomo (2019) show that pension funds that were affected more by the new regulations were more likely to enter the Dutch mortgage market. The Dutch central bank has also attributed the increased activity of non-bank players during this period to changes in regulations outside the mortgage market (De Nederlandsche Bank 2016). For this reason, I argue that the timing of entry can be seen as plausibly exogenous.

A defining feature of the Dutch mortgage market concerns the tax deductibility of mortgage interest payments. For this purpose, the tax code distinguishes two types of mortgages.

16. See Footnote 3.

17. For the remainder of this article, I anonymize bank names. This is a precondition for using my main data source, the LLD. I explain the LLD in more detail in Section 4.

The first consists of amortizing mortgages. These are mortgages where the balance is paid off during the duration of the mortgage. This category includes annuities and linear mortgages. Interest payments on amortizing mortgages are always deductible from income taxes. The second type of mortgage is non-amortizing mortgages. For these mortgages, households only make interest payments——the balance is paid off as a lump sum on the end date of the mortgage. Often, these mortgages are sold together with savings or investments products, the proceeds of which are used to pay off the balance. I call mortgages that are sold together with another product *complex*, while I denote all other (including amortizing) mortgages as *simple*. Interest payments on non-amortizing mortgages are only tax-deductible for households that already had a non-amortizing mortgage before 2013. As a result, first-time buyers only purchase amortizing mortgages. The fact that these are tax-deductible makes them categorically more attractive than non-amortizing mortgages. This fact will be crucial in separating the treatment and control markets, as I explain below.

4 Data

The main data source of this article is the Loan Level Data (LLD) of the Dutch Central Bank (DNB). The LLD are a micro-level data set on most mortgages in the Dutch market. It contains 75%–80% of the market, with only some smaller and foreign players missing (Mastrogiacomo and Van der Molen 2015). The LLD are reported quarterly, but this study is based on the reports from the fourth wave only so that I in effect have yearly waves. Participating institutions report their full portfolio. Because the entrants are relatively small, they do not routinely report their data to the LLD. However, in the middle of 2016, DNB required *all* institutions in the market to give a one-time snapshot of the outstanding mortgages in their portfolio. Since the entry I focus on happens in 2014, this means that I have information on the most relevant part of the entrants' portfolio: in the first years of their existence.

Note that since institutions report the *stock* of outstanding mortgages in 2016, and virtually all mortgages are active for at least a couple of years, I observe all sales in the years directly prior to 2016. While I do observe prior periods in my data, I start my sample in 2013. The reason is that at the beginning of 2013 many reforms took place in the Dutch mortgage market.[18] Therefore, the pre-2013 market might be fundamentally different than the post-2013 market.

To make mortgages as comparable as possible, I restrict the sample to mortgages insured by the Dutch government under the so-called NHG scheme. Under Dutch law, households are personally liable for any remaining debt if they default on their mortgage. For mortgages with NHG protection, this default risk is insured by the government for the most common causes of default, such as divorce, disability or unemployment. As a result, mortgages with NHG protection are seen as significantly lower risks than other mortgages. Hence, banks tend to price all mortgages with NHG protection equally, increasing comparability. NHG protection is only available if the value of the underlying property is below the mean property value in the country. Since banks offer significant interest rate discounts on mortgages with this pro-

18. For example, the tax deductibility of non-amortizing mortgages changed to the regime described in Section 3, the ban on history-based price discrimination described in Section 3 took effect in 2013 and there were changes to broker regulations as well.

Table 3: Market shares (%) of payment types

	Incumbents	Entrants	Total
Simple loans			
Annuity	36.68	70.06	38.17
Linear	3.62	7.46	3.79
Fixed installments	0.09	0.00	0.09
Bullet	33.65	22.48	33.16
Complex loans			
Savings	16.52	0.00	15.79
Life	6.64	0.00	6.34
Investment	1.66	0.00	1.59
Other	1.13	0.00	1.08
Total	100.00	100.00	100.00

Note: The table displays the distribution of payment types across loans, for loans with a fixed interest period starting in the period Jan 2013–Jul 2016. Entrants are defined as institutions selling their first mortgage in the period Aug–Nov 2014. The remaining institutions are labeled incumbents. A loan is complex if it sold as a bundle with another product. Otherwise, it is simple.

tection, there is no reason to worry about selection bias: it virtually always pays for qualifying households to acquire the insurance.

Table 3 shows the distribution of payment types in the Dutch mortgage market. Many different mortgages are available. Annuity and linear mortgages, which are typical in many other countries, have an overall market share of 42%. Bullet mortgages, where no payments towards the principal are made during the mortgage, have a market share of 33%. Often, bullets are combined with a savings, life or investment product—these types of mortgages have a combined market share of 24%.

The distribution of payment types differs in several respects between incumbents and entrants. First, entrants do not sell savings, life or investment mortgages. Second, entrants sell more annuities and fewer bullets. This is because bullets are non-amortizing mortgages. As explained above, non-amortizing mortgages are not attractive to first-time buyers for tax reasons and are hence only purchased by renewers and switchers. Since entrants do not have any renewers, it is expected that they sell relatively fewer bullets.

In the LLD, I observe contractually agreed instead of posted interest rates. While mortgage pricing to some extent takes place at the individual level, or depends on particular features of a mortgage that I do not observe in my data, the bulk of variation in observed differences in interest rates is between rather than within mortgage products. I define a *mortgage product* by its three main characteristics: the bank that sells the mortgage, the duration of the fixed interest rate duration (5, 10, 15 or 20 years) and the method by which the mortgage is repaid. Conditioning on mortgage product explains about 82% of variation in contractually agreed interest rates. Hence, even though there is some residual price dispersion, it is sensible to talk about the interest rate of a particular mortgage product.

5 The effects of entry on interest rates

Section 2 predicts the existence of a market share-entry response gradient in markets with switching costs. In this section, I test for this gradient using the 2014 entry of mortgage sellers funded by pension funds described in Section 3.

5.1 Main empirical specification

To estimate the causal effect of pension funds' entry on incumbent banks' interest, I use the identification approach illustrated in Figure 2. The main idea is to combine two separate sources of identification: differential entry into different sub-markets, and within-bank differences in the size of its customer base for the different mortgage products it sells.

First, I use a difference-in-differences approach to control for contemporaneous shocks to the mortgage market. Here, I exploit the fact that the entrants do not sell all types of mortgages that are available in the market. As Table 3 shows, the entrants do not sell savings, life or investment mortgages. Hence, it is possible to split the market into two segments: a market *affected* by entry comprised of simple mortgages, and a market *unaffected* by entry, comprised of complex mortgages.[19] Under the assumption that incumbents face similar demand and cost shocks in both markets, the unaffected market serves as a control for the affected market.

A difference-in-differences approach identifies a homogeneous (or average) treatment effect, but the main implication from the theoretical model is that incumbents with different amounts of locked-in customers react differently to entry. This section also shows that a simple treatment heterogeneity analysis can give biased estimates of the market share-entry response gradient, as unobserved firm characteristics are likely correlated with the number of locked-in customers a bank has. To alleviate this concern, I condition on Bank $\times$ Fixed Rate Duration $\times$ Month fixed effects. Hence, I exploit within-bank variation in the share of locked-in customers of its mortgage products to estimate the heterogeneous effects of entry. Intuitively, I compare the interest rate a particular bank sets for a 10-year annuity mortgage with the interest rate the same bank sets for a 10-year bullet mortgage. Since both annuity and bullet mortgages are treated, I can exploit cross-sectional as well as time variation in the share of locked-in customers a bank has *across* these products to identify the slope of the market share-entry response gradient. The identifying assumption is that the unobservable shocks a bank faces are similar across the different mortgage products it sells, so that absent any differences in the share of locked-in customers, the change in interest rates of different treated products sold by the same bank should be the same.

Combining these two identification arguments leads to the regression model:

$$Interest_{ibdmt} = \gamma_1 Size_{bmt} + \gamma_2 Post_t \times Treat_m + \gamma_3 Post_t \times Treat_t \times Size_{bmt} + \tau_{bdt} + \tau_m + X'_{ibdmt}\delta + \varepsilon_{ibdmt} \quad (3)$$

Here, i indexes the consumer, bdm the mortgage product he purchases and t the start month of the fixed interest rate period. In particular, b is the bank that sells the mortgage, d indicates the

19. I ignore fixed installments and other types of mortgages since they have negligible market shares.

Figure 2: Graphical depiction of the identification strategy of this paper.

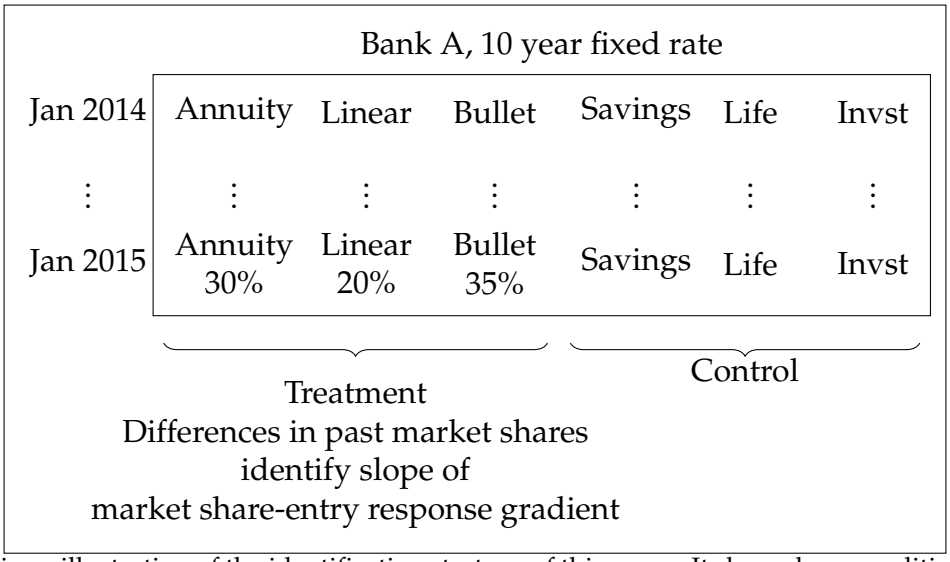

Note: The picture is an illustration of the identification strategy of this paper. It shows how conditioning on the mortgages of a single bank (in this case 10 year-mortgages sold by "Bank A") controls for unobserved bank characteristics and how a difference-in-differences approach controls for contemporaneous shocks to the mortgage market. Past market shares are for purpose of illustration only.

duration of the fixed interest rate period and m indicates the type of mortgage (as in Table 3). $Post_t$ is a dummy indicating whether entry has taken place.[20] $Treat_m$ is a dummy indicating whether product market m is affected by entry or not. $Size_{bmt}$ indicates the size of the bank in the type of market in which mortgage bm is sold. For example, if m indicates the annuity market, $Size_{bmt}$ measures the past sales of bank b across consumers with annuity mortgages that are up for renewal at time t. For $Size_{bmt}$ I take various cutoffs that divide the market into small and large banks, as well as a bank's past market share amongst renewing customers as a continuous variable. τ_t contains month fixed effects. X_{ibdmt} contains additional controls which might affect the interest rate a consumer pays.[21] In this sense, the regression can be interpreted as a generalized hedonic pricing regression. Such an approach is useful since it allows me to control for the fact that mortgage pricing to an extent takes place at the individual level, as well as for potential compositional changes between the treatment and control groups.

Note that two dummies that typically appear in a difference-in-differences specification are subsumed by other fixed effects: a $Post_t$ dummy is subsumed by the τ_{bdt} dummies and a $Treat_m$ dummy is subsumed by the τ_m dummies. I cluster standard errors at the firm level.

The main coefficient of interest is γ_3, which measures the strength of the market share-price response gradient. In terms of the model from Section 2, γ_3 measures how much the slope of firms' policy functions, β, changes after entry. According to Proposition 2, it is expected that $\gamma_3 > 0$. γ_2 measures the interest rate decrease in the treatment market of an incumbent without locked-in customers relative to the control market. Hence it is comparable to the change in the slope of firms' policy functions, α, and theory predicts that $\gamma_2 < 0$.

20. I take the first date of entry among all entrants, i.e. August 2014.

21. X_{ibdmt} always includes the loan-to-value ratio of the mortgage. As a robustness check, I will also include further variables as detailed in Table 5.

Figure 3: Interest rates over time, by treatment and control market

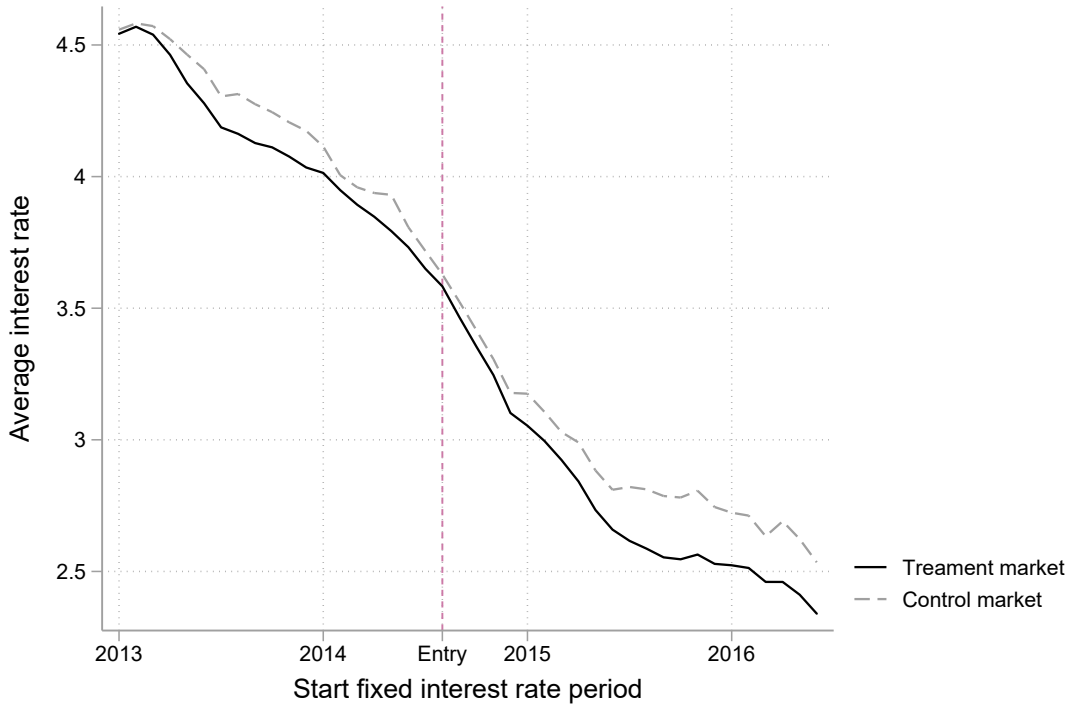

Note: The figure shows average interest rates of the incumbents by start of the fixed interest rate period for government-backed loans with a 10-year fixed interest rate period.

5.2 Identification

As discussed above, my approach identifies the differential causal effect of entry on incumbents' interest rates for incumbents of different sizes under the following assumptions: i) a parallel trends assumption ii) a common shock assumption, i.e. mortgages sold by the same incumbent face similar unobserved demand and supply shocks. I now discuss these assumptions in more detail, as well as potential challenges to them.

5.2.1 Parallel trend assumption

My approach identifies a causal effect only if absent entry, incumbents' interest rates in the treatment and control markets would have followed similar trends. This is a plausible assumption, because incumbents face the same cost of funding for mortgages sold in the treatment and control markets. In the same way, demand shocks, for example due to a marketing campaign, are likely to affect mortgages in both markets similarly. Figure 7 plots average interest rates for government-backed loans with a 10-year fixed interest rate duration.[22] The pre-treatment trends are similar: this supports the argument that the treatment and control markets are affected by similar shocks. In Section 5.3.1, I formalize this argument by testing for "leads" and "lags" of the treatment. This exercise also indicates the similarity of the treatment and control markets.

A potential challenge to the parallel trends assumption is that there are systematic differ-

22. I plot a subsample of all loans to make the sample more homogeneous. Government-backed 10-year loans are the most popular category in the market.

Figure 4: Post-entry correlation between interest rates of annuity and bullet mortgages, by bank

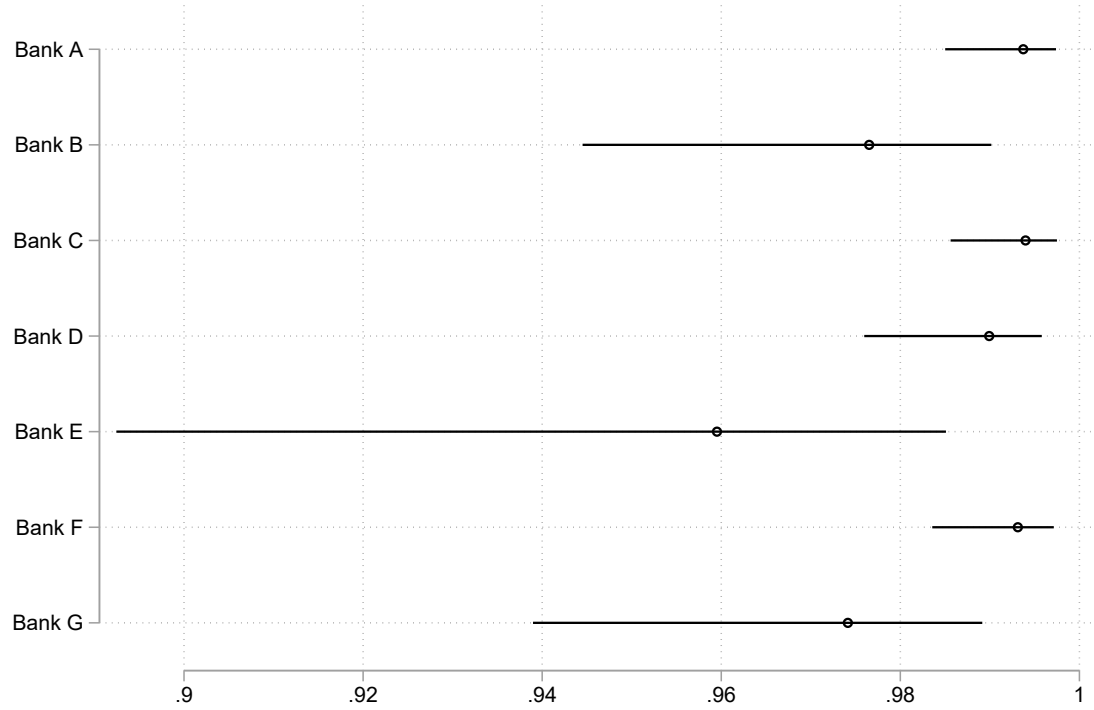

Note: The graph shows the Pearson correlation between the average monthly interest rate on 10-year annuity and bullet mortgages, by bank. The bars indicate 95% confidence intervals, as calculated by the Stata corrci package (Cox 2008). The correlations are only displayed for incumbent banks that have a market share of at least 1%. Bank names have been anonymized to comply with regulations on using the LLD for research purposes.

ences between borrowers in the treatment and control markets (Table A.1). Households in the treatment market are on the whole older and richer. They also have smaller loan-to-value ratios and their houses are worth more. Such systematic differences are no problem for difference-in-differences designs *per se*, as long as the parallel trend assumption holds. However, systematic differences between the control and treatment market raise the possibility of there being differential changes in customer composition between the treatment and control market. Such differential changes in composition might render the parallel trends assumption invalid. For example, if younger households on average pay higher interest rates than older households and the number of young households in the treatment market increases relative to the control market, I would overestimate the effect of entry on incumbents' interest rates. Here, the hedonic pricing approach introduced in Section 5.1 helps. Because I can control for some observable factors that might influence the interest rates that households pay, I can (at least partially) also control for differential changes in composition

5.2.2 Common shock assumption

To identify the relationship between past market shares and response to entry, I assume that different mortgage products sold by the same incumbent are affected by similar unobserved demand and cost shocks. If this is the case, then differences in how an incumbent changes the interest rates of two different mortgage products are fully attributable to differences in past market shares.

Table 4: Switching between treatment and control markets

	2013	2014	2015	2016	Total
Proportion switchers	.049 (.006, .165)	.016 (.000, .088)	.017 (.000, .091)	.068 (.030, .130)	.043 (.023, .074)
N	41	61	59	117	278
p-val. no diff. in switching[1]	.545				

95% exact confidence intervals are in parentheses

[1] p-value of Fisher exact test with the null hypothesis that the proportion of switchers is the same in 2013/2014 as in 2015/2016.

Note: The table displays the proportion of households that switches between a mortgage in the control market and a mortgage in the treatment market in the DNB Household Survey. The baseline population is households present in two subsequent waves of the survey, who i) have at least one mortgage in both waves, ii) have at least one mortgage with a fixed interest period starting in the displayed year. A household with one mortgage is a switcher if the mortgage is from a different market in the first wave compared to second wave. A household with more than one mortgage is considered a switcher if it has only mortgages from one market in the first wave and only mortgages from the other market in the second wave.

To test the plausibility of this assumption, I calculate the correlation between monthly average interest rates for annuity and bullet mortgages *in the post-entry period* (Figure 4). These types of mortgages are the two most popular in the treatment market (Table 3). The idea is that if, for a given incumbent, demand and cost shocks for these mortgages are similar, the interest rates of these two mortgage types should display significant co-movement. The restriction to the post-entry period is important since the observed shock of pension fund entry might change this correlation.

Figure 4 shows that the interest rates on annuity and bullet mortgages display significant correlation, with point estimates over .95 for every incumbent bank. I interpret this as support for the common shock assumption, for if annuity and bullet mortgages would face different shocks, interest rates should not move together so closely over time.

5.2.3 Separation of treatment and control markets

My empirical strategy is valid only if the treatment and control markets can be seen as separate markets. If not, the entry would also affect the control market. That the treatment and control markets can be seen as separate is not immediately obvious, because in principle simple and complex mortgages are substitutes. However, institutional details as well as statistical evidence indicate that to a large extent the treatment and control market can be seen as separate.

The institutional argument is based on the tax treatment of mortgages. As detailed in Section 3, from 2013 onwards, first time buyers' interest payments are only tax deductible for amortizing (annuity and linear) mortgages. As a result, first time buyers are only active in the treatment market. The tax advantages of amortizing mortgages are so large that they never substitute towards a complex mortgage. Renewing households can only deduct interest payments on non-amortizing mortgages if they had such a mortgage prior to 2013. Therefore, renewing households similarly never switch from the treatment to the control market.

However, it is possible that renewing households switch from the control to the treatment market. In the LLD the possibility to follow switching households over time is limited. There-

fore, I look at the DNB Household Survey to assess how large the fraction of switchers is. The DNB Household Survey is a yearly survey of a random sample of the Dutch population on their finances. Therefore, even though its sample size is much smaller than that of the LLD, it still allows me to obtain a representative estimate of the extent of renewing households' switching. Moreover, it has the advantage that households can be followed from one wave of the survey to the next. Table 4 shows that the extent of switching is modest. On average, 4.3% of households that could potentially switch, switch from the control to the treatment market. Moreover, I find no statistical evidence that the rate of switching changed after entry had taken place, as one would expect if the entry had affected the control market.

A second potential challenge to separation is that banks might have cost functions without constant returns to scale. In this case, a bank's sales in the treatment market could affect the costs it faces in the control market. Even if there were no demand-side substitution, this would still imply that banks would change their interest rates in the control market after entry. I do not have information on the structure of banks' cost functions. However, Figure A.1 shows that most incumbents' sales are relatively constant over time. Hence, even if incumbents' costs feature non-constant returns to scale, it is unlikely that this had a large impact on mortgage pricing.

For these reasons, I argue that the treatment and control market can be seen as largely separate. However, the separation is not perfect. It is unfortunately not obvious what direction a potential bias due to imperfect separation introduces. On the one hand, prices of different firms are strategic complements. Hence, if banks do adjust their interest rates in the control market as a result of the entry into the treatment market, they will do so in the same direction. Therefore, I would underestimate the magnitude of the effect of entry. On the other hand, banks in this market are multiproduct firms. Hence, if consumers substitute between the control and treatment market, banks will internalize the business-stealing effect of a change in interest rates in the treatment market on profits in the control market. This means that they set higher interest rates in the treatment market compared to the case where there is perfect separation and that I underestimate the effect of entry.

5.2.4 Selection and anticipation

In difference-in-differences designs, selection into the treatment is typically a large concern. In this application, however, selection is unlikely to be an issue. This is for two reasons. First, any selection is on the part of entrants, while I only consider outcomes for the incumbents. This means selection is only an issue if entrants select on *incumbents'* post-entry outcomes. While this is certainly not impossible—for example, firms might be more likely to enter in markets where post-entry competition is less fierce—it does limit the amount of selection that is possible. Second, the fact that the entrants sell only simple and no complex mortgages is likely driven by regulation rather than by selection. First, because first-time buyers no longer receive tax benefits for complex mortgages, the control market is slowly dying out. This makes it less attractive to enter this market than the treatment market. Second, a reason why pension funds entered the mortgage market is because changes in pension regulations made investing in mortgages more attractive. The reason for this is that mortgage debt is seen as relatively safe (Kim and Mastrogiacomo 2019). However, complex mortgages are less safe than simple

Table 5: Main results: The effect of entry on incumbents' interest rates

	(1)	(2)	(3)	(4)
Loan-to-value ratio	0.116*	0.112*	0.112*	0.161**
	(0.0537)	(0.0526)	(0.0525)	(0.0501)
Market share[1]	0.176	0.141**	0.0246	-0.0552
	(0.107)	(0.0584)	(0.0507)	(0.0514)
Post × Treated	-0.0602	-0.121***	-0.131**	-0.160**
	(0.0382)	(0.0184)	(0.0422)	(0.0573)
Post × Treated × High market share[2]		0.108**		
		(0.0476)		
Post × Treated × Market share[1]			0.304**	0.268**
			(0.0951)	(0.0955)
Constant	3.466***	3.484***	3.543***	3.319**
	(0.0897)	(0.0693)	(0.0558)	(1.243)
Bank × fixed rate duration × month effects	Yes	Yes	Yes	Yes
Payment type effects	Yes	Yes	Yes	Yes
Additional controls[3]	No	No	No	Yes
Observations	557495	557495	557495	401960
R^2	0.803	0.803	0.803	0.815

Standard errors in parentheses

Standard errors are clustered at the firm level

* $p < .10$, ** $p < .05$, *** $p < .001$

[1] Market share amongst renewing consumers with the same mortgage type in the relevant month. Mortgage types are defined as in Table 3.

[2] A bank has a high market share for a certain mortgage if it has a market share of at least 20% amongst renewing customers with a mortgage of that type. Mortgage types are defined as in Table 3.

[3] Birth year and income of primary borrower; property location fixed effects; original loan balance; mortgage purpose fixed effects.

mortgages. Therefore, it was less attractive for pension funds to enter into the control market. Both these factors are unlikely to be correlated with incumbents' post-entry responses, so that selection is not a large issue here.

A second potential issue is that incumbents might have anticipated pension funds' entry into the Dutch mortgage market. The issue would be that, anticipating fiercer competition for locked-in customers in the future, incumbents would already invest less in market share before the ban. That is, under anticipation effects one would expect the interest rates in the treatment market to increase pre-entry, relative to the control market. Figure 3 shows that this does not happen.

5.3 Main results

I now estimate the effects of entry on incumbents' interest rates. Column (1) in Table 5 shows the estimated regression model (3), but assuming a homogeneous treatment effect. The esti-

Figure 5: Treatment heterogeneity by past market share

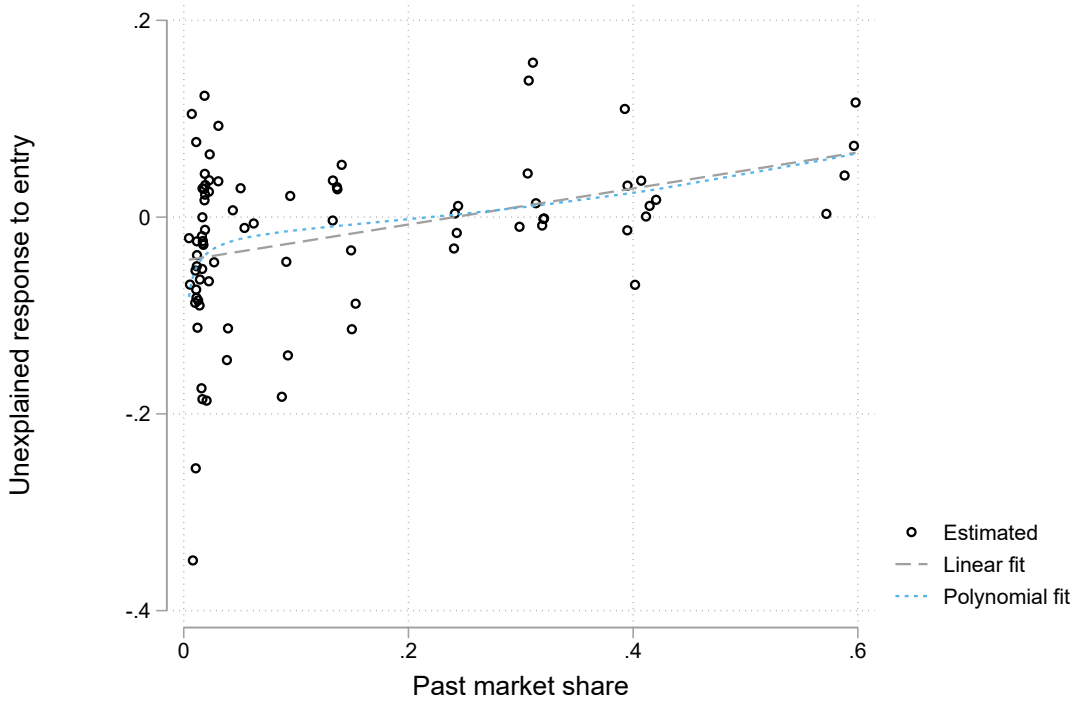

Note: The graph plots unexplained entry responses versus past market shares, by mortage product. The y-axis shows the average residuals from regression (1) in Table 5, that is a difference-in-difference model assuming a homogeneous tretment effect of entry. On the x-axis the average past market share in the post-entry period is displayed. Only treated mortgage products are included.

mates show that, on average, incumbents decreased the interest rates on mortgages sold in the treatment market relative to mortgages in the control market. The average decrease is about 6 basis points. Given the post-entry average interest rate in the Dutch mortgage market of 2.9%, this is a plausible decrease of about 2%.

Before estimating the full model, I provide two supplementary pieces of evidence for the existence of a market share-entry response gradient. First, Figure 5 plots the average residuals from the regression in column (1) of Table 5 by past market share. It shows that for mortgage products with larger past market shares, a homogeneous treatment effect model tends to under-estimate the effect of entry (and vice versa for mortgage products with smaller market shares). This points to entry responses differing by past market share. Second, Figure 6 shows average interest rate trends for small and large incumbents.[23] It shows that smaller banks, on average, seem to have decreased their rates more post-entry than larger banks. Hence, the residuals from a homogeneous treatment effect regression as well as high-level graphical evidence seem to support the theory from Section 2.

Columns (2)-(4) of Table 5 show estimates of regression (3), all allowing for treatment effect heterogeneity. Column (2) treats an incumbent being "large" as binary: an incumbent is large in a particular market segment if it has a market share of at least 20% in that segment.[24] If

23. It does not matter whether I classify a bank as big or small based on the treatment, control or full sample.

24. The 20% cut-off is based on the fact that, for the whole market, three banks have market shares that are significantly higher than 20%, while the remaining banks have market shares under 20%. Hence, there is a "gap" in the market share distribution around 20%, making it a natural cut-off point.

Figure 6: The effect of entry for small and large incumbents

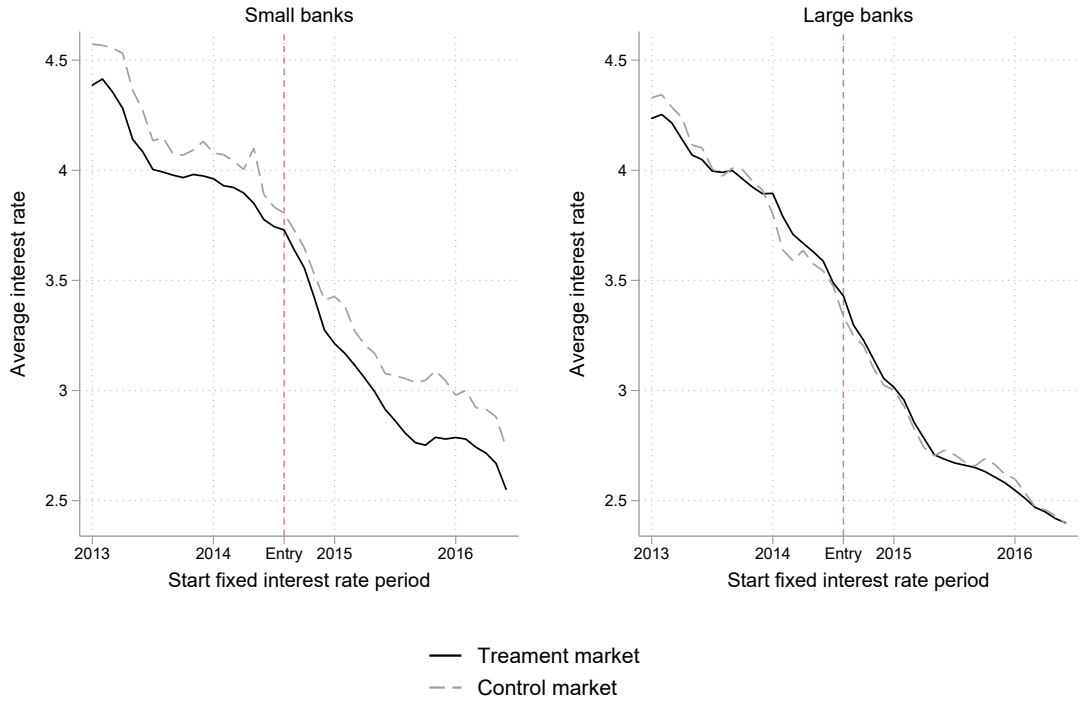

Note: The figure shows average interest rates by start of the fixed interest rate period for government-backed loans with a 10-year fixed interest rate period by size of the incumbent. An incumbent is considered large if its past market share amongst renewing customers is at least 20% on average in the post-treatment period.

an incumbent is small in a particular segment, it decreases its interest rate by 12 basis points, or about 4%. Large incumbents however barely decrease their interest rate post-entry: the average effect of entry is $-12 + 11 = -1$ basis point, or .3%. Table A.2 in the Appendix re-estimates this model using different cut-offs for an incumbent being large. Except when the cut-off is quite large (40%), the results are qualitatively the same: large incumbents respond much less to entry than small incumbents.

Column (3) interacts the entry dummy with a mortgage product's past market share, taken as a continuous variable. Hence, this specification is the closer to the theoretical model of Section 2, wherein entry responses are also linear in past market shares (see equation (2)). The estimate show that, if a particular mortgage product has a negligible market share, its interest rate decreases by 12 basis points, or 4.5%, on average. For every 10 percentage points increase in past market share, a mortgage product's interest rate decreases by 30 basis points less, or by about 1 percentage point less. This implies that if a mortgage product has a past market share of at least 43%, its interest rate increases instead of decreases. Mortgage products with such large market shares are rare, but do exist in my sample. Hence, my results indicate some prices increasing after entry. Column (4) adds additional individual-level controls. Noting that these variables are not available for my full sample, my results are quite robust to the inclusions of these additional controls.

Figure 7: Estimated treatment effects over time

(a) Homogeneous treatment effect

(b) Heterogeneous treatment effect: treatment on 0% market share bank

Note: The figure shows the treatment effect with associated 95% confidence interval over time. The per-month treatment effect is the regression coefficient on a month $\times$ treatment group dummy in a hedonic pricing regression. Confidence intervals are based on standard errors clustered at the bank level.

5.3.1 Treatment effects over time

To test the plausibility of the parallel trends assumption, Figure 7 displays treatment trends over time. I do so for two cases. The left panel displays a estimated homogeneous treatment effect over time, which is analogous to column (1) in Table 5. The right panel shows the estimated treatment effect for the heterogeneous case, analogous to column (3) in Table 5. Since the heterogeneous treatment effect depends on past market share, . I have chosen to plot the effect of entry on an incumbent without any locked-in customers. This is simply because for such an incumbent, the evidence is the starkest, so that the graph is easiest to interpret. Obviously, for incumbents with larger market shares, the post-entry treatment effect line would lie above the plotted line—for incumbents with sufficiently large market shares, the estimated treatment effect is not statistically different from zero.

For both models, consistent with the graphical evidence discussed before, "leads" of the treatment are not significantly different from zero. Figure 7 moreover shows that the effect of the treatment increases over time. In other words, it seems to have taken some time for the entry to have full effect on the market.

5.3.2 The importance of within-bank variation

To establish the importance of controlling for unobserved bank-level shocks, I rerun the analysis using additive instead of multiplicative fixed effects. That is, instead including Bank $\times$ Fixed rate duration $\times$ Month effects, I include Bank, Fixed rate duration and Month fixed effects separately. While such a specification does control for systematic interest rates differences between, for example, different incumben tbanks, it does not absorb the full shock a particular bank incurs at a particular time. Hence, the causal identification of this model rests only on the parallel trends assumption.

Table A.3 in the Appendix shows that compared to the main specification, the estimated

market share-entry response gradient is smaller. As explained in Section 2, this is consistent with the presence of unobserved bank-level marginal cost shocks. The results are also less stable, for example to including additional controls, than for the main specification, suggesting that the additive fixed effects specification indeed suffers from omitted variable bias. Hence, using multiplicative fixed effects to control for unobserved bank-level shocks is crucial to get a good estimate of the market share-entry response gradient.

6 Policy implications

I now discuss the policy implications of my results, for two broad policy areas: first, for policies that try to encourage competition; second, for competition policy.

6.1 Retail competition

Governments often try to encourage competition in (former) public services to lower consumers' costs. One prominent example of such a policy is the liberalization of electricity markets around the globe.[25]. Another important area is health insurance, where multiple countries have tried to encourage competition.[26] However, such industries are typically characterized by large switching costs.[27] The results in this paper imply that the expected gains from fostering competition can be limited. Indeed, the more dominant an incumbent is, the less it can be expected to respond to competition, so that in markets where an increase in competition seems to be the most useful, the less it actually accomplishes. Perversely, it is exactly the firms with the largest customer bases, where hence the largest gain could be had, that respond the least. Although I find limited empirical support for this in the current setting, the theory even allows for incumbents to increase their prices after an increase in competition. Hence, fostering competition in markets with switching costs might also raise important distributional concerns, where unaffiliated consumers gain at the expense of consumers locked-in at a large incumbent.

These results, however, do not imply that encouraging competition in markets with switching is always without benefit. First, unaffiliated consumers will almost certainly benefit from increased competition as firms without an established customer base, i.e. entrants, will tend to set low prices. Moreover, in the long run the benefits of competition are likely larger than the short run. For example, in the steady state of Somaini and Einav (2013), an increase in the number of firms is unambiguously good for consumers. Rather, since it might take a long time to reach a steady state, my results imply that policies that foster competition cannot be expected to immediately deliver in markets with switching costs.

25. Such policies have been introduced in all of the European Union, in various US states, as well as in Australia and New Zealand (Defeuilley 2009).

26. Gaynor, Ho, and Town (2015) contains a good overview of recent reforms with the explicit aim of fostering competition.

27. To continue the examples, Hortaçsu, Madanizadeh, and Puller (2017) provides evidence of switching costs in retail electricity and Handel (2013) in health insurance.

6.2 Competition policy

My results have two different implications for competition policy. The first is for merger analysis. In a market with switching costs, a horizontal merger has two effects: i) a decrease in the number of firms; ii) the combination of the merging firms' customer bases. Since a merged firm has a larger market share than its constituents, it will, holding the number of firms constant, set a higher price. However, a decrease in the number of firms at the same time *reduces* dynamic pricing incentives. In other words, after the merger firms have smaller incentives to "harvest" locked-in customers. This is likely to dampen the price increase of the merging firms. At the same time, the dynamic pricing incentives of the non-merging firms are also reduced: hence, while large non-merging firms might even decrease their price, small firms increase their price after a merger as the investment motive becomes less strong. It is outside the scope of this paper to give a full account of the effects of horizontal mergers in markets with switching costs, but the results at least suggest them to be non-trivial.

A second implication concerns the interpretation of switching costs as a barrier to entry, which can for example be important to determine the contestability of a market. As Klemperer (1987a) already noted, it is not clear whether switching costs should be seen as a barrier to entry, as "in addition to locking in customers to buying from the incumbent, high switching costs may in effect lock in the incumbent to selling only to repeat customers" (p. 100). Hence, he concludes that "entry may be more deterred by either very high or very low switching costs than by switching costs in a middle rang" (p. 100). My results imply that not only the level of switching costs, but also the degree of concentration determines to what extent switching costs should be seen as a barrier to entry. If, as in Klemperer (1987a), the market is served by a monopolist it is indeed quite likely that it will not compete aggressively with a potential entrant and switching costs do not form a significant barrier to entry. However, if for the same level of switching costs, there is additionally a fringe firm active, this fringe firm will compete aggressively with a potential entrant. An entrant then faces significant competition for unattached customers from the fringe firm, while attached customers are locked-in due to switching costs. In such a scenario, it is quite likely that switching costs form a significant barrier to entry. This effect does not occur in the existing literature (Klemperer 1987a; Farrell and Shapiro 1988), which restricts itself to pre-entry monopolies.

7 Conclusion

This paper has theoretically and empirically studied the effect of competition in markets with switching costs. On the theoretical front, I study the model of Somaini and Einav (2013). By simulating equilibria of their model for a large number of parameter constellations, I show that, contrary to common intuition for frictionless markets, competition does not necessarily lead to price decreases. Rather, competition increases dynamic pricing incentives. Firms for which the incentive to invest in market share is already strong, decrease their price post-entry. On the other hand, firms for which the incentive to "harvest" existing locked-in customers is strong, increase their price post-entry.

I test the derived implications by studying entry into the Dutch mortgage market. Because

the entrants only sell a subset of all available mortgages in the market, entry splits the market into a treatment and a control market. I employ a difference-in-differences framework, to study the effect of entry on incumbents' interest rates. Consistent with theory, I find that smaller banks decrease their interest rates more than larger banks.

The results in this paper suggest that, when a market features switching costs, incumbents will not react much to increased competition and that policies to encourage competition can hence be ineffective. This result hold only in the periods after entry has taken place—it is possible that in a new steady state competition nevertheless significantly benefits consumers. It would be interesting to explicitly account for the transition to a new steady state in further work.

This paper also sheds new light on an old literature that discusses whether switching costs are a barrier to entry. My results suggest that this literature's attention to pre-entry monopolies has been restrictive and that pre-entry concentration can play a large role. This is another avenue for future work: the theoretical framework developed in this paper for example does not allow for entry deterrence. Hence, it would be interesting to see what a model of entry deterrence with pre-entry competition would look like.

References

Abe, Naohito, and Daiji Kawaguchi. 2010. "Incumbent's price response to new entry: the case of Japanese supermarkets." *Journal of the Japanese and International Economies* 24 (2): 196–212.

Ailawadi, Kusum L, Jie Zhang, Aradhna Krishna, and Michael W Kruger. 2010. "When Wal-Mart enters: How incumbent retailers react and how this affects their sales outcomes." *Journal of Marketing Research* 47 (4): 577–593.

Allen, Jason, Robert Clark, and Jean-François Houde. 2014. "The effect of mergers in search markets: Evidence from the Canadian mortgage industry." *American Economic Review* 104 (10): 3365–96.

Arie, Guy, and Paul LE Grieco. 2014. "Who pays for switching costs?" *Quantitative Marketing and Economics* 12 (4): 379–419.

Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro. 2018. "Retail globalization and household welfare: Evidence from Mexico." *Journal of Political Economy* 126 (1): 1–73.

Basker, Emek. 2005. "Selling a cheaper mousetrap: Wal-Mart's effect on retail prices." *Journal of Urban Economics* 58 (2): 203–229.

Basker, Emek, and Michael Noel. 2009. "The evolving food chain: competitive effects of Wal-Mart's entry into the supermarket industry." *Journal of Economics & Management Strategy* 18 (4): 977–1009.

Beggs, Alan, and Paul Klemperer. 1992. "Multi-period competition with switching costs." *Econometrica* 60 (3): 651–666.

Bennett, Daniel, and Wesley Yin. 2019. "The Market for High-Quality Medicine: Retail Chain Entry and Drug Quality in India." *Review of Economics and Statistics* 101 (1): 76–90.

Bresnahan, Timothy F, and Peter C Reiss. 1991. "Entry and competition in concentrated markets." *Journal of Political Economy* 99 (5): 977–1009.

Busso, Matias, and Sebastian Galiani. 2019. "The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment." *American Economic Journal: Applied Economics* 11 (1): 33–56.

Cabral, Luis. 2016. "Dynamic pricing in customer markets with switching costs." *Review of Economic Dynamics* 20:43–62.

Carbo-Valverde, Santiago, Timothy H Hannan, and Francisco Rodriguez-Fernandez. 2011. "Exploiting old customers and attracting new ones: The case of bank deposit pricing." *European Economic Review* 55 (7): 903–915.

Chen, Yongmin, and Michael H Riordan. 2007. "Price and variety in the spokes model." *Economic Journal* 117 (522): 897–921.

Cosguner, Koray, Tat Y Chan, and PB Seetharaman. 2016. "Dynamic pricing in a distribution channel in the presence of switching costs." *Management Science* 64 (3): 1212–1229.

Courtemanche, Charles, and Art Carden. 2014. "Competing with Costco and Sam's Club: Warehouse club entry and grocery prices." *Southern Economic Journal* 80 (3): 565–585.

Cox, N. J. 2008. "Speaking Stata: Correlation with confidence, or Fisher's z revisited." *Stata Journal* 8 (3): 413–439.

De Nederlandsche Bank. 2016. *Loan markets in motion: Larger role of pension funds and insurers boosts financial stability.* Technical report. Accessed May 9, 2019.

Defeuilley, Christophe. 2009. "Retail competition in electricity markets." *Energy Policy* 37 (2): 377–386.

Doganoglu, Toker. 2010. "Switching costs, experience goods and dynamic price competition." *Quantitative Marketing and Economics* 8 (2): 167–205.

Edlin, Aaron S, and Robert G Harris. 2012. "The role of switching costs in antitrust analysis: a comparison of Microsoft and Google." *Yale JL & Tech.* 15:169.

Fabra, Natalia, and Alfredo García. 2015. "Market structure and the competitive effects of switching costs." *Economics Letters* 126:150–155.

Farrell, Joseph, and Carl Shapiro. 1988. "Dynamic competition with switching costs." *RAND Journal of Economics* 19 (1): 123–137.

Gaynor, Martin, Kate Ho, and Robert J Town. 2015. "The industrial organization of health-care markets." *Journal of Economic Literature* 53 (2): 235–84.

Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review* 103 (7): 2643–82.

Hannan, Timothy H, and Robert M Adams. 2011. "Consumer switching costs and firm pricing: Evidence from bank pricing of deposit accounts." *Journal of Industrial Economics* 59 (2): 296–320.

Hausman, Jerry, and Ephraim Leibtag. 2007. "Consumer benefits from increased competition in shopping outlets: Measuring the effect of Wal-Mart." *Journal of Applied Econometrics* 22 (7): 1157–1177.

Honka, Elisabeth. 2014. "Quantifying search and switching costs in the US auto insurance industry." *RAND Journal of Economics* 45 (4): 847–884.

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L Puller. 2017. "Power to choose? An analysis of consumer inertia in the residential electricity market." *American Economic Journal: Economic Policy* 9 (4): 192–226.

Hosken, Daniel S, Luke M Olson, and Loren K Smith. 2018. "Do retail mergers affect competition? Evidence from grocery retailing." *Journal of Economics & Management Strategy* 27 (1): 3–22.

Hotelling, Harold. 1929. "Stability in Competition." *Economic Journal* 39 (153): 41–57.

Janssen, Aljoscha. 2018. "Switching Costs, Quality Misconceptions and Behavioral Pricing in the Pharmaceutical Market." *Mimeo.*

———. 2019. "Price Dynamics of Swedish Pharmaceuticals." *Mimeo.*

Jia, Panle. 2008. "What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry." *Econometrica* 76 (6): 1263–1316.

Kim, Yeorim, and Mauro Mastrogiacomo. 2019. "The effect of the Dutch Financial Assessment Framework on the mortgage investments of Pension funds and their borrowers." *Mimeo.*

Klemperer, Paul. 1987a. "Entry deterrence in markets with consumer switching costs." *Economic Journal* 97:99–117.

———. 1987b. "Markets with consumer switching costs." *Quarterly Journal of Economics* 102 (2): 375–394.

Klemperer, Paul. 1988. "Welfare effects of entry into markets with switching costs." *Journal of Industrial Economics* 37 (2): 159–165.

———. 1989. "Price wars caused by switching costs." *Review of Economic Studies* 56 (3): 405–420.

Lach, Saul, and José L Moraga-González. 2017. "Asymmetric price effects of competition." *Journal of Industrial Economics* 65 (4): 767–803.

Lira, Loreto, Rosario Rivero, and Rodrigo Vergara. 2007. "Entry and prices: Evidence from the supermarket sector." *Review of Industrial Organization* 31 (4): 237–260.

MacKay, Alexander, and Marc Remer. 2019. "An Empirical Model of Consumer Affiliation and Dynamic Price Competition." *Mimeo.*

Mastrogiacomo, Mauro, and Remco Van der Molen. 2015. "Dutch mortgages in the DNB loan level data." *DNB Occasional Studies* 13 (4).

Padilla, A Jorge. 1995. "Revisiting dynamic duopoly with consumer switching costs." *Journal of Economic Theory* 67 (2): 520–530.

Pearcy, Jason. 2016. "Bargains followed by bargains: when switching costs make markets more competitive." *Journal of Economics & Management Strategy* 25 (4): 826–851.

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. 2014. "Re-examining the effects of switching costs." *Economic Theory* 57 (1): 161–194.

Ruiz-Aliseda, Francisco. 2016. "When do switching costs make markets more or less competitive?" *International Journal of Industrial Organization* 47:121–151.

Russell, Lauren. Forthcoming. "Price Effects of Non-Profit College and University Mergers." *Review of Economics and Statistics.*

Shcherbakov, Oleksandr. 2016. "Measuring consumer switching costs in the television industry." *RAND Journal of Economics* 47 (2): 366–393.

Somaini, Paulo, and Liran Einav. 2013. "A model of market power in customer markets." *Journal of Industrial Economics* 61 (4): 938–986.

Stango, Victor. 2002. "Pricing with consumer switching costs: Evidence from the credit card market." *Journal of Industrial Economics* 50 (4): 475–492.

Thiel, Jurre H. 2018. "Price discrimination, switching costs and welfare: Evidence from the Dutch mortgage market." *Mimeo.*

To, Theodore. 1996. "Multi-period competition with switching costs: An overlapping generations formulation." *Journal of Industrial Economics* 44 (1): 81–87.

A Additional tables and figures

Table A.1: Borrower and loan characteristics for the treatment and control market

	Control market	Treatment market	Total
Age primary borrower (years)	43	40	41
	(27.27)	(98.74)	(85.88)
Household income (€)	50341	50139	50184
	(22563.3)	(44237.2)	(40410.6)
Loan-to-value ratio	87	87	87
	(31.46)	(24.02)	(26.06)
Property valuation (€)	162422	174786	171721
	(79576.1)	(80737.3)	(80627.9)

Note: The table displays average borrower and loan characteristics for the treatment and the control market, for the whole sample. Standard deviations are in parentheses.

Table A.2: The effect of entry on incumbents' interest rates, for different cut-off definitions of "large bank"

	(1) 10% cut-off	(2) 20% cut-off	(3) 30% cut-off	(4) 40% cut-off
Market share[1]	0.164*	0.141**	0.0492	0.163
	(0.0846)	(0.0584)	(0.0566)	(0.168)
Loan-to-value ratio	0.114*	0.112*	0.113*	0.116*
	(0.0537)	(0.0526)	(0.0525)	(0.0539)
Post × Treated	-0.128**	-0.121***	-0.0887**	-0.0613
	(0.0309)	(0.0184)	(0.0386)	(0.0426)
Post × Treated × High market share[2]	0.0903	0.108**	0.0723	0.00698
	(0.0566)	(0.0476)	(0.0407)	(0.0435)
Constant	3.473***	3.484***	3.528***	3.472***
	(0.0822)	(0.0693)	(0.0479)	(0.108)
Bank × fixed rate duration × month effects	Yes	Yes	Yes	Yes
Payment type effects	Yes	Yes	Yes	Yes
Observations	557495	557495	557495	557495
R^2	0.803	0.803	0.803	0.803

Standard errors in parentheses

Standard errors are clustered at the firm level

* $p < .10$, ** $p < .05$, *** $p < .001$

[1] Market share amongst renewing consumers with the same mortgage type in the relevant month. Mortgage types are defined as in Table 3.

[2] A bank has a high market share for a certain mortgage if it has a market share of at least equal to the percentage is in the column header amongst renewing customers with a mortgage of that type. Mortgage types are defined as in Table 3.

Table A.3: The effect of entry on incumbents' interest rates, without exploiting within-firm variation

	(1)	(2)	(3)	(4)
Loan-to-value ratio	0.124*	0.125*	0.124*	0.176**
	(0.0605)	(0.0608)	(0.0602)	(0.0566)
Market share[1]	0.223**	0.224**	0.213*	0.238*
	(0.0662)	(0.0675)	(0.108)	(0.115)
Post × Treated	-0.0954**	-0.0915**	-0.103**	-0.118
	(0.0382)	(0.0364)	(0.0450)	(0.0614)
Post × Treated × High market share[2]		-0.00721		
		(0.0639)		
Post × Treated × Market share[1]			0.0312	-0.156
			(0.170)	(0.156)
Constant	3.532***	3.533***	3.534***	2.528
	(0.103)	(0.0992)	(0.105)	(1.395)
Bank effects	Yes	Yes	Yes	Yes
Fixed rate duration effects	Yes	Yes	Yes	Yes
Month effects	Yes	Yes	Yes	Yes
Payment type effects	Yes	Yes	Yes	Yes
Additional controls[3]	No	No	No	Yes
Observations	557495	557495	557495	401960
R^2	0.730	0.730	0.730	0.743

Standard errors in parentheses

Standard errors are clustered at the firm level

* $p < .10$, ** $p < .05$, *** $p < .001$

[1] Market share amongst renewing consumers with the same mortgage type in the relevant month. Mortgage types are defined as in Table 3.

[2] A bank has a high market share for a certain mortgage if it has a market share of at least 20% amongst renewing customers with a mortgage of that type. Mortgage types are defined as in Table 3.

[3] Birth year and income of primary borrower; property location fixed effects; original loan balance; mortgage purpose fixed effects.

Figure A.1: Coefficient of variation of monthly mortgage sales by incumbent

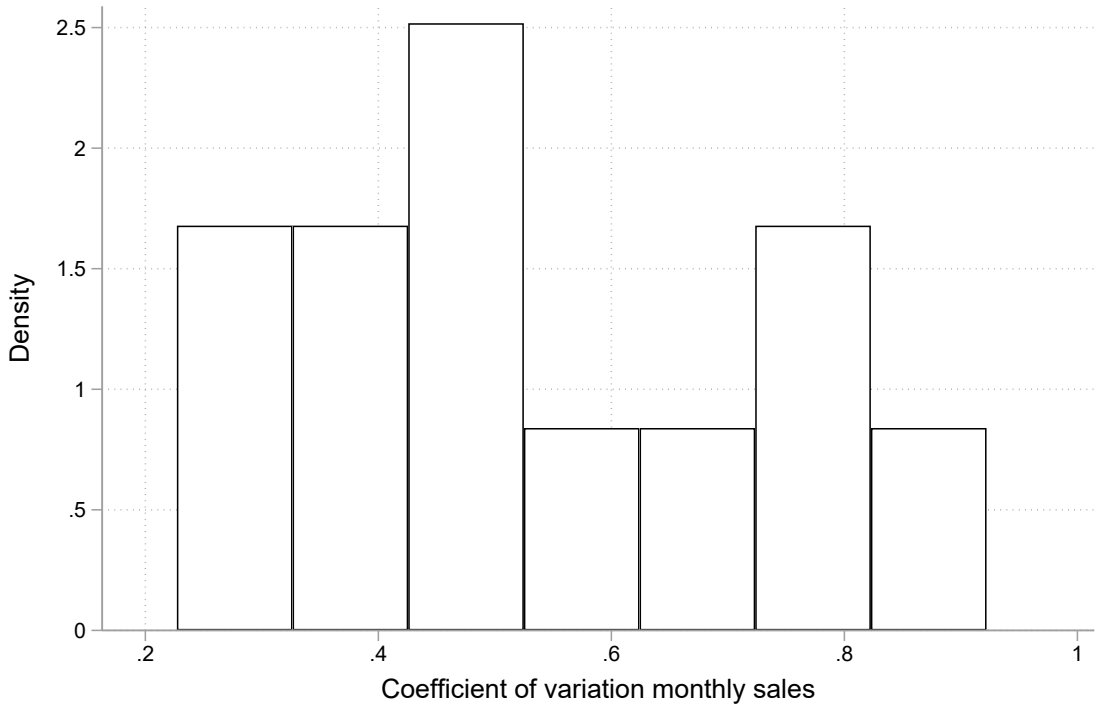

Note: The figure shows the coefficient variation (standard deviation divided by mean) of mortgage sales by incumbents. Sales are measured in total volume of originated or renewed mortgages in euro's. Incumbents with market shares smaller than 1% are not shown.

B The effects of entry in a model with asymmetric costs

In this appendix, I introduce asymmetric costs into the model analyzed in Section 2.[28] Somaini and Einav (2013, Theorem 1) show that there still exists a unique Markov Perfect Equilibrium in parallel strategies. In this case, the price of firm i satisfies

$$p_i = \bar{c} + \mu_0 + \mu_1(c_i - \bar{c}) + \beta x_i,$$

where x_i is the previous market share of firm i, c_i its constant marginal cost and $\bar{c} = N^{-1}\sum_{i=1}^{N} c_i$ the average marginal cost in the industry. μ_0, μ_1 and β are positive parameters that depend on the primitives of the model (but not directly on marginal costs). Denote $\alpha_i \equiv \bar{c} + \mu_0 + \mu_1(c_i - \bar{c})$ to write

$$p_i = \alpha_i + \beta x_i. \tag{B.1}$$

The difference with (1) is that the intercept is now firm-specific. Hence, the strength of dynamic pricing incentives β/α_i is now firm-specific, with firms with lower marginal costs having a stronger harvesting motive. Simulations on the grid in Table 1 show how the parameters in firms' policy functions change as N increases.

Proposition B.1. *As the number of firms N increases, μ_0 decreases, μ_1 increases and β increases.*

There are unambiguous predictions for the coefficient β, but the effect of entry on α_i is ambiguous:

$$\alpha_i = \underbrace{\bar{c}}_{?} + \underbrace{\mu_0}_{-} + \underbrace{\mu_1}_{+}\Big(\underbrace{c_i}_{=} - \underbrace{\bar{c}}_{?}\Big). \tag{B.2}$$

The effect of entry on $\bar{c}$ is ambiguous because the entrant can have higher or lower marginal costs than the pre-entry average. However, it seems reasonable to assume that the marginal costs of the entrant is no higher than $\bar{c}$, for it is unlikely that a firm with relatively high marginal costs finds it profitable to enter into an established market. Under this assumption $\bar{c}$ decreases and the comparative statics of α_i depend on the following elements:

$$\alpha_i = \underbrace{\bar{c}}_{-} + \underbrace{\mu_0}_{-} + \underbrace{\mu_1}_{+}\underbrace{(c_i - \bar{c})}_{+}.$$

Hence, for firms with below-average marginal costs, α_i decreases after entry and dynamic pricing incentives β/α_i increase. The same is true for firms for which the difference $c_i - \bar{c}$ is not too large, since then the first two terms dominate the last. Therefore, the prediction made in the main text that entry amplifies dynamic pricing incentives holds under the conditions that i) the entrant is at least as efficient as the industry average, ii) there is not too much dispersion in marginal costs.

Inspection of firms' policy function (B.1) shows that, for a fixed marginal cost c_i, firms with larger established market shares decrease their prices less or increase their prices more—the equivalent of Proposition 4. However, there does not necessarily exist a cut-off market share such that firms with established market shares below this cut-off decrease their prices

28. This is equivalent to the existence of vertical differentiation, in addition to the horizontal differentiation that is already present in the symmetric model.

and firms with established market shares above this cut-off increase their prices, even for a fixed c_i. The reason is that in the asymmetric model the intercept α_i need not decrease post-entry. Hence, it is possible that for some c_i firms increase their price no matter their established market share. Figure 1 shows an example of this: here, firms with sufficiently high marginal costs always increase their prices post-entry. However, when comparing firms with the same marginal costs, the overall picture is the same as in Propositions 3 and 4: the further a firm's established market share is away from the cut-off (if it exists), the stronger it reacts to entry.

What about the responses of firms that *differ* in their marginal cost? As equation (B.2) shows, the higher marginal costs c_i, the smaller the decrease of α post-entry. Hence, $\frac{\beta}{\alpha}$ increases more for firms with lower marginal costs. In other words, for firms with lower marginal costs dynamic pricing incentives are amplified more. Figure 1 plots the price response to entry for firms with different marginal costs and market shares. The lower marginal costs are, the less dynamic pricing incentives are amplified and hence the higher a firm's past market share must be to increase its prices.