Firms have Partial Knowledge and they Partially Optimize: Evidence From A Reform

Avner Strulov-Shlain*

March 6, 2024

Abstract

Firms may try to maximize profits but fail for various reasons. I use a reform to argue that the only model of the firm consistent with the data is one with misoptimization due to insufficient knowledge of the world. A reform in Israel limited prices to end with 0 as the cent digit (e.g. 2.90 but not 2.99). Since consumers are left-digit biased, specifically causing demand to discretely fall by 5%-9% at round prices, optimal pricing implies bunching at just-below prices (e.g. 2.99 pre-reform or 2.90 post) and avoiding round prices (e.g. 3.00). In fact, before the reform, supermarkets set just-below prices for 45% of prices and rarely used round prices. If price-setting before the reform was driven by the correct model of demand, firms' response to the reform would have been to update immediately according to their beliefs and avoid round prices. However, pricing after the reform is inconsistent with their pre-reform revealed beliefs, setting 20% of clearly dominated prices for almost a year. Whether firms were optimizing a wrong model or making decisions in a model-free way, their knowledge had to be partial. Partial knowledge driven by incomplete learning can explain how firms behave suboptimally in a persistent way and challenges counterfactual exercises that rely on the assumption of model-based optimization. I suggest an approach of almost-optimization to calculate counterfactuals that are more likely to contain the truth.

*Strulov-Shlain: University of Chicago Booth School of Business, avner.strulov-shlain@chicagobooth.edu
I am thankful for Stefano DellaVigna for endless encouragement and support. Ned Augenblick, Zarek Brot-Goldberg, Ben Handel, and Dmitry Taubinsky provided many thoughtful comments. Christy Kang provided excellent research assistance. This research greatly improved thanks to discussions with Nick Barberis, Modibo Camara, Jean-Pierre Dube, Ben Enke, Johannes Hermle, Ella Honka, Ali Hortacsu, Alex Imas, Tesary Lin, Sanjog Misra, Sarah Moshary, and Muriel Niederle. All errors are my own. I am in debt to Itai Ater, Roni Rojkin (ICC), Saul Lach, and Merav Yiftach (CBS) for greatly assisting me with data collection and thus making this project viable.

1 Introduction

It is important to understand what firms know and how they learn. However, it is often hard to observe and separate between different explanations, since we only observe one equilibrium, and the data do not lend themselves to sharp tests. I take advantage of a reform that forced a change in the equilibrium of an entire market by narrowing down the set of prices firms can charge. This reform allows us to test different explanations of firm behavior in the long and short run, specifically what do firms know and how they learn. I find that the only consistent explanation for firms behavior before and after the reform is that they do not fully learn about demand and operate under constant partial knowledge.

My findings are consistent with a growing body of evidence that documents persistent mis-optimization of firms[1]. These findings go against the cornerstone assumption, and strongly held belief, of firm optimization in equilibrium. The assumption states that firms can be described as if they use a model to evaluate the best available actions, and commonly further assumes that firms choose optimally. This assumption holds from neoclassical economics (Muth (1961)) to behavioral IO (Heidhues and Kőszegi, 2018; Spiegler, 2011). In contrast, the mis-optimization empirical literature cited above states that firms may be wrong, and that in order to be consistently wrong they must have the wrong model, wrong beliefs, or that some frictions arise. These conditions are often doubted to hold since wrong models and beliefs should be corrected in the market, since firms operate in a high incentive, high-feedback, environment. This paper builds on the argument that sustained wrong model or wrong beliefs stem from limited learning that leads to *sustained partial knowledge*: If firms have a wrong prior model in mind and they do not try to refute it, suboptimal behavior will persist (Hanna et al., 2014; Gagnon-Bartsch et al., 2021; Schwartzstein, 2014); If firms are model-free decision makers (Arthur, 1991; Cross, 1973, operationalized with reinforcement learning, Sutton and Barto 2018) and do not explore the action space enough, wrong beliefs and suboptimal behavior will persist.

This paper uses data and a reform affecting supermarket pricing to empirically document a case where large firms hold sustained partial knowledge, which can be a source of long run mis-optimization. It then studies possible as-if explanations for firm behavior, of either a wrong model or model-free decision making based on wrong beliefs, and discusses their implications.

To test if firms' decision making is consistent with full knowledge optimization I use data from a reform that limited the set of admissible actions a firm could take. To illustrate my empirical strategy, consider the following thought experiment: there are three options with an objective ranking of $a \succ b \succ c$. The decision maker is taking action a in steady state, behaving consistently with the ranking. Then, a is removed from the choice set. What should the decision maker choose?

[1] See Bloom and Van Reenen 2007; Cho and Rust 2010; DellaVigna and Gentzkow 2019; Dube et al. 2017; Dubé and Misra 2023; Goldfarb and Xiao 2011, 2019; Hanna et al. 2014; Hortaçsu et al. 2019; Hortacsu et al. 2021; Huang 2021; List et al. 2023; Rao and Simonov 2019; Shapiro et al. 2021; Strulov-Shlain 2022

The next best option, b. Instead, imagine they initially choose c, and later on, in the new steady state switch to b. While we cannot reject that the decision maker knew that $a \succ b$ and that $a \succ c$, which was sufficient knowledge initially, we can conclude they did not know that $b \succ c$. Thus, they operated under partial knowledge in the old steady state. The empirical strategy in this paper tracks this thought experiment.

In 2013, the Israeli government banned pricing in nonexisting coins — the equivalents of the US Penny and Nickle — forcing companies to price at .X0-ending prices (Ater and Gerlitz, 2017). For example, 2.96 is banned, but 2.90 or 3.10 are allowed. Before the reform, the modal price-ending in Israeli supermarkets was 99 (e.g. 2.99) with $\sim$ 45% of prices, and there were almost no 00-ending prices. This price-ending distribution *before* the reform is consistent with at least partial maximization, specifically with a model of pricing to left-digit biased demand[2] — a demand structure driven by consumers with distorted price perception leading to discontinuous drops in demand at round prices.[3] When this is the demand structure, prices that absent a bias would have ended with low endings (e.g., 5.00 or 5.29) are better priced at the lower 99-ending price (4.99 in this case). When 99-ending—previously a profit-maximizing price ending for many costs and elasticities—is banned and the highest just-below price-ending is 90, still, pricing predictions are that prices should not end with 00.[4] Indeed, demand analysis using scanner data shows that changing a product's price from, for example, 4.99 to 5.00 versus from 4.99 to 4.90 was accompanied with excess demand loss of 5%-9%. In the notation of the above thought experiment, 99-ending prices are option a, 90-ending prices are b, and 00-ending prices are c, with $a \succ b \succ c$. Thus, the behavior of the companies before the reform was consistent with at least knowledge of $a \succ b$ and $a \succ c$.

At the time of the reform's implementation, 20% of prices ended with 00. This high share is inconsistent with the demand structure where $b \succ c$. In the years that followed, the proportion of 00-endings decreased again. For some chains, a discrete reduction in the share of 00-ending prices occurred about 8 months after the reform. The evolution of price-endings has been such that products whose price-ending changed from 99 to 90 stayed at 90, while those whose price changed to 00-ending quickly reverted to 90-ending prices. That is, we have evidence of moving from a to b and c, but then moving away from c to b.

The behavior of firms over the long term is also inconsistent with complete knowledge optimization. Beyond the stylistic three-choice example, chains' pricing behavior in the long run, either before the reform or long after, is suboptimal due to underuse of 99-ending prices (consistent with List et al., 2023; Strulov-Shlain, 2022). The magnitude of left-digit bias should have led to more 99-ending and 90-ending prices than observed. One way to explain this is that

[2] That supermarkets realize that they face left-digit biased demand, at least to some degree, was also found in supermarkets in the U.S. (Strulov-Shlain, 2022).

[3] For example, if 4.99 is perceived as much lower than 5.00, demand discontinuously drops at 5.00.

[4] There is indeed a tradeoff due to the reform: a larger revenue loss from pricing below the round price threshold (e.g., at 4.90 vs. 4.99), and the discontinuity is effectively "smoother." Yet, left-digit bias that is enough to justify many 99-ending prices also leads to prices being either 90-ending or higher, but not at 00.

chains used 99-ending prices before the reform because they understood the value of 99 versus 00; however, firms did not use *enough* 99-endings before the reform (or 90-endings after) possibly because they operated under partial knowledge and did not understand the full demand structure. Further, this partial knowledge is also consistent with the above mentioned excessive 00-endings after the reform.

How can partial knowledge be sustained in the long-run? The key is under-exploration. We can explain all patterns in the data in a few ways. For example, we can assume firms were maximizing against a wrong model of demand that is non-monotonic at 99-ending prices. i.e., downward sloping demand but with high demand spikes at 99-ending prices. This structure, although refuted in the data (as well as by List et al. (2023); Strulov-Shlain (2022)), was suggested for example by Stiving (2000) and Anderson and Simester (2003). Another alternative is that firms are model-free learners, and were able to learn the relative benefits of 99-endings to 00-endings but not of 90-endings to 00-endings, and only learned those after the reform. Either of these two explanations can only hold if they are coupled with under exploration. Had firms tried to refute the demand-with-spikes model, they would have; had they used more price endings, they would have learned the full structure of demand.

Sustained suboptimal behavior has implications on counterfactual exercises in two ways. First, a researcher who is recovering unobservables from a model does so under the assumption of optimization (see, e.g., every structural paper with a counterfactual section); second, a researcher who is predicting outcomes when conditions change does so under the assumption of optimization. Suboptimal response and partial knowledge do not mean lack of optimization, but a failure to reach full optimization. Therefore, I do not recommend throwing the baby out with the bathwater and assume "anything goes", but rather test the sensitivity of countefactual conclusions to varying degrees of suboptimization. In the spirit of Afriat (1972) and Rabin (2013), I introduce a "fudge factor" $X \in [0, 1]$, which is how far from optimization we allow the decision maker to be. The idea is simple - the standard revealed preference argument is that if we observe a firm taking an action $p = p^*$ and we know that under optimization $p = \Pi(p; a) = f(a)$ is a function of latent parameters of interest a, then we infer that $a = f^{-1}(p^*)$. If we observe a firm taking an action with the goal of maximizing profits, but we allow it to fail to optimize up to X, then any a that satisfies $\Pi(p^*, a) \geq (1 - X) \cdot \Pi(f(a), a)$ is a possible value of the parameter of interest.

Conceptually, this factor is not new. It is closely related to Afriat's famous Critical Cost Efficiency Index (Afriat, 1972) and to subsequent work focusing mainly on individual decision makers (e.g., Choi et al. (2014); de Clippel and Rozen (2021); Halevy et al. (2018); Varian (1990)). These papers propose to calculate such a factor in order to measure how far a decision maker is from maximizing utility, or offer guidelines on how to choose the best descriptive model of behavior by minimizing this factor. I deviate from this literature in two ways. Superficially, I focus on firms maximizing profits, rather than consumers maximizing utility. Substantively, I offer to use this factor even when the factor cannot be estimated from the data. The existing literature thinks of settings where data reject all models and how to reconcile that gap; instead, I consider an exercise where the model is assumed and asking what are the implications of deviations from optimization.

As an illustrative example, consider a profit-maximizing monopolist selling one good. An $X = 0.05$ means that the firm acts with at most 5% of unexploited profits. Therefore, if we see a monopolist setting a price of \$2.00 and estimate a price elasticity of -2, then instead of determining that the marginal cost is exactly \$1 (as it is under optimization), we get a range of costs $[0.78, 1.22]$. For each cost in this range, a price of \$2 provides at least 95% of the profit as with otherwise optimal prices. Therefore, allowing for a 5% friction in optimization leads up to 22% deviation in inferred costs, and markups of 39%-61% instead of exactly 50%. The proposed use of the fudge is as a layer of protection against model misspecification while preserving the assumption that firms *strive* to maximize the objective function the researcher has in mind.

A series of recent papers find large values for such fudge factors. Bloom and Van Reenen (2007, 2010) document persistent and common productivity gaps between firms and attribute these gaps to management practices. Across industries (such as tech companies, national retail chains, large airlines, car rental companies, international brands), the literature documents a loss of profits of dozens of percent in specific actions that translate into several percent of the overall profits (Blake et al., 2015; Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Dubé and Misra, 2023; Hortacsu et al., 2021; List et al., 2023; Shapiro et al., 2021; Strulov-Shlain, 2022; Zhang and Misra, 2022). These are selected examples, but many of the quoted papers do not seek to show misoptimization and yet find it. In general, I believe it is justified to acknowledge the possibility that X is positive, and thus to test the sensitivity of conclusions to various values of X.

The main contribution of this paper is to provide empirical evidence of sustained partial knowledge by large sophisticated firms, even in the presence of a learning process. Close papers studying individuals are Schwartzstein (2014); Hanna et al. (2014); Gagnon-Bartsch et al. (2021) showing theoretically and empirically that in a multi-attribute decision making environment, people can make persistent mistakes if they ignore some optimization-relevant dimensions. Anderson et al. (2021); Bloom et al. (2013) show in field experiments in India and Uganda that management and marketing training increased profits and growth, and attribute the pre-intervention suboptimal behavior to information gaps. Theoretically, Gagnon-Bartsch et al. (2021) propose a framework that argues "why misconceptions may persist: we channel our attention to data we deem relevant under our model, whereas it is the data we deem irrelevant that is most likely to contradict our model." In this paper's setting the framework of Gagnon-Bartsch et al. means that to sustain misunderstanding of the role of left-digit bias, the chains need to deem price endings as irrelevant for demand. However, it is not that firms thought that price endings do not matter, as they have demonstrated by pricing at 99-endings before the reform. That is, firms thought that price endings could be relevant, yet not to the extent that they would take the time to fully understand if and how much so.

Second, this paper the studies what positive model of a firm can generate these behaviors. It provides empirical evidence that sheds light on how firms make decisions, suggesting that firms can be described as model-free decision makers with partial knowledge, a formalization that can generalize to many settings. In contrast, to be described as standard model-based optimizers, they must have the wrong model of demand, which may take many forms. Why model-free? The theoretical idea that firms are model-free learners was suggested before (Cross, 1973) and there is ample support in the neuroscience and psychology literature that reinforcement learning is the way humans learn (e.g., Daw et al., 2011; Thorndike, 1911). Reinforcement learning had been studied, mainly in the lab, for individuals (Camerer, 2011), and since decision makers in firms are often people, it seems likely to extend this idea to firms. In addition, reinforcement learning is a useful way to solve complicated problems and a central tool of AI (Dayan, 1992; Sutton and Barto, 2018; Watkins and Dayan, 1992); we have evidence that when these methods are explicitly employed by firms, they affect market outcomes (Assad et al., 2020, 2021; Calvano et al., 2020); and these ideas seem useful in explaining phenomena in financial trading (Barberis and Jin (2021)). This paper provides suggestive evidence that model-free learning occurs at the firm level and leads to long-term model-free decision making. Therefore, the implications of that description of the firm should be further studied.

This paper offers a unique empirical opportunity that adds to the literature on firms' learning, and suggests considering model-free learning and processes that lead to partial knowledge (see review by Aguirregabiria and Jeon (2020)). Unlike in other settings, firms did not have to learn a new demand structure, most notably in the papers by Doraszelski et al. (2018) and Huang et al. (2018). In Doraszelski et al. (2018) electricity providers' ultimate bidding behavior is consistent with optimization, and their learning is studied as learning about the parameters of the correct model. Further, there is no sense in which learning type (model-based or model-free) can be identified, because we cannot know firms' initial beliefs. Yet, the authors find it hard to explain firms' pricing in the first 1.5 years of data, presumably because behavior is inconsistent with any stable beliefs about model parameters. Huang et al. (2018) also remain in the realm of model-based optimization, studying retailers setting prices for liquor in Washington. However, the article shows that companies initially put a large weight on recent outcomes, a finding that is also consistent with reinforcement learning. This paper proposes to use the lens of model-free learning in such exercises.

Finally, the paper provides another example for a case in which firms do not optimize, both in the long-run and in the short-run, and contributes an explanation for how that can happen. Apart of the papers by Schwartzstein and coauthors (Schwartzstein (2014); Hanna et al. (2014); Gagnon-Bartsch et al. (2021), the literature explores other explanations. When firms are found to be mis-optimizing it is explained by adding frictions and heuristics (DellaVigna and Gentzkow (2019); Hortacsu et al. (2021); Huang (2021); Strulov-Shlain (2022)), organizational economics arguments of principal-agent incentives misalignment (Hortacsu et al. (2021); Shapiro et al. (2021)), management practices (Anderson et al., 2021; Bloom and Van Reenen, 2007, 2010; Bloom et al., 2013), attention costs (Goldfarb and Xiao (2011, 2019)), or wrong beliefs (Aguirregabiria and Jeon (2020); Hortaçsu and Puller (2008); Strulov-Shlain (2022); Xie (2018)). This paper further provides evidence for partial knowledge as a common, and here necessary, component, with under-exploration as a mechanism that sustains it. If firms operate under partial knowledge, what their actions reveal should be taken with care. As such, I suggest a different approach to counterfactual calculations that allows persistent sub-optimization.

Partial knowledge implies that, even if firms converge to nearly optimal behavior in the long run, they are sensitive to small changes in the environment. Therefore, counterfactual exercises that assume perfect model-based optimization, such as merger implications, inferring markups, or studying the effects of regulation, should be done with this alternative description of the firm in mind and test its sensitivity to a fudge factor.

The paper proceeds as follows. Section 2 describes the setting, the data, and left-digit biased demand and its effects on pricing. Section 3 then shows how prices evolve and the demand response. Section 4 suggests two interpretations of the data, through the lenses of model-based and model-free learning. Section 5 introduces the fudge factor, discusses its justification, magnitude, and provides use cases. Section 6 concludes with a discussion of the implications of these findings for future research.

2 Setting and Data

2.1 Reform details

On October 17, 2013, the Israeli Ministry of Economics announced that starting January 1, 2014, it will start to enforce an existing law banning pricing in non-existing coins.[5] Israel abolished the 1-Agora and 5-Agorot coins (the Shekel consists of 100 Agorot) in 1991 and 2008, respectively.[6] Yet prior to 2014, the majority of prices ended in non 10-Agorot ("dime"), most commonly with 9- and 5-endings. Therefore, most prices had to change due to the reform. This policy came about because consumers felt that they were being tricked by non-existing prices (e.g., a store posts a price of 4.99 but the actual price is 5.00; prices are VAT inclusive) as they pay the total shopping basket price when paying with a card or the price rounded to the nearest dime when paying cash.

2.2 Data

The effects of the reform are manifested through the dynamics of price setting and their effects on demand. Therefore, I use data on prices and purchases in supermarkets. I employ three

[5] except for "continuous" products—gas, water, and electricity

[6] https://web.archive.org/web/20201127080757/https://www.boi.org.il/en/Currency/CurrentCurrencySeries/Pages/Default.aspx#Top

data sources: two of them are shelf prices before and after the policy changes, and the third one is scanner data containing revenue and quantities. In all datasets, the unit of observation is a product in a store in a period. I first describe the pricing data and then the scanner data.

In the pricing data, the main variables are the shelf price as displayed in a store, and identifiers of the product, the store, and the chain. The main pricing data was collected by a consumer advocacy group, "Israel Consumer Council" (ICC). Its employees went to supermarket stores with a varying list of products every 2–4 weeks between early 2013 and early 2015 and wrote the prices of a set of predetermined products. A "product" can be thought of as a UPC—of unique producer, weight, and size. ICC collects and posts the data on their website to inform consumers which of the nearby stores sells the cheapest bundle in a given week. Therefore, stores and chains are identified by their name and address. The other pricing data, from the Central Bureau of Statistics (CBS), is similar, but a "product" in the data is a product type, for example, a carton of omega3-enriched brown medium eggs (any brand). Data are a balanced panel, but stores identities are anonymized without chain identifiers, and prices are sampled monthly. More details are provided in Appendix Section A.

The pricing data are not a balanced panel due to missing observations and changes in the basket of goods sampled. Therefore, I focus on a consistent subset of products and impute missing prices as described in Appendix Section A. For example, for the main analysis, I require that the price of a product in a store be sampled at least once in each reform period.

Table 1 shows the differences in the number of products, stores, and prices between the samples. The main ICC sample consists of 21 products in 143 stores, and the main CBS sample consists of 171 product-types in 99 stores. The share of 99-ending prices is 45%–47% in the pre-period.

Finally, a third data set is a mix of daily and weekly scanner data, in which the main variables are the number of units sold and the revenue of each UPC in a store. Like in the ICC data, each store and chain has a unique identifier, though anonymous. Data are provided by StoreNext Ltd, a market research firm that collects transaction level data from store registers and aggregates it. The unit of observation is a product (UPC) in a store in a week.[7] From an initial sample of 10 UPCs from a balanced set of 295 supermarket stores over 3 years (2013--2015), 6 products can be used for the analysis.[8] I merge these data with another sample of StoreNext's *weekly* data for another 13 products from 287 national supermarket stores covering 2013 to 2015.[9]

The scanner data are not ideal for precise price measurements due to the averaging of possibly different prices paid in a time period (see Einav et al., 2010; Strulov-Shlain, 2022). StoreNext's data

[7]Because of the mix of daily and weekly readings, I aggregate the data at the weekly level.

[8]One of the UPCs has a single price throughout the sample (canned corn); another one is a product that was not available before the reform (400g Hummus). Two others are cottage cheeses that follow unusual pricing behavior (see Hendel et al., 2017).

[9]While the set of stores is likely largely overlapping, I cannot match stores or chains between these two sources since store identifiers are unique to each data set.

is the only research-available scanner data from Israel but has some excess measurement errors, requiring careful data cleaning (elaborated in Appendix A).

The main scanner dataset, described in column "StoreNext Main" in Table 1, includes 185k observations. Compared to the price data, there are fewer 99-ending prices (23% vs. 43%) and more 00-ending prices after the reform (15% vs. 11%). In some analyses, I also use a "long price spells" sample. A same-price spell is when the average price of the product in a store is identical over consecutive weeks. In these cases I restrict the sample to be same-price spells of at least 2 weeks, aiming to capture non-promotional product-weeks.

Together, these data allow one to explore the price evolution of products in supermarket stores at the aggregate, use the panel structure of the data, and examine the effects of price endings on demand.

2.3 Pricing under left-digit bias

Many prices end with 99 before the reform and with 90 after the reform, likely as a response to left-digit biased demand. This section describes the concept of left-digit bias and what it implies for pricing.

Left-digit bias is the tendency of consumers to put excess weight on the leftmost digits of a price. Under left-digit bias, the perceived difference between prices is bigger than actual if they have different left-most digits and smaller than actual if they share the left digits. In turn, demand will slope downward with discontinuous drops when the left digit changes. The key implication of such a demand structure, i.e. downward sloping with discontinuities, is that prices to the right of the discontinuities should be avoided.

The literature broadly supports this explanation, finding left-digit bias in supermarket pricing and other settings. Early literature on the reasons why firms use psychological prices reached mixed conclusions about their profitability (Basu, 1997; Gedenk and Sattler, 1999; Ginzberg, 1936; see Stiving and Winer, 1997 for a brief review) and suggested various explanations for their existence (e.g., Basu, 1997, 2006; Bizer and Schindler, 2005; Anderson and Simester, 2003). Some experimental and empirical analyses were conducted to see whether and how 9 and 99 price endings affect demand in practice, and found mixed results as well (Anderson and Simester, 2003; Bizer and Schindler, 2005; Sokolova et al., 2020; Thomas and Morwitz, 2005). Strulov-Shlain (2022) provides large scale evidence from 25 US supermarket chains and thousands of products supporting high levels of left-digit biased demand. Left-digit bias was previously and recently established in non-price settings (e.g., Lacetera et al. (2012); Busse et al. (2013); Li and Qiu (2023)) and recently extended to other demand settings by Hilger (2018) and List et al. (2023), providing further support for its effects on demand.

To study the implications on pricing requires incorporating left-digit bias into demand and solving for optimal pricing. To solve for optimal pricing, I am using the model from Strulov-Shlain (2022), in which biased consumers meet monopolistic firms. In the model, a price is perceived in a distorted way, as a mix with weight $1-\theta$ on the true price, and θ on the price with a fixed focal price ending Δ.[10] It causes the perceived price to change discontinuously when the left-most digit changes, and also be less sensitive to price changes that do not involve left-most digit changes. For example, with left-digit bias $\theta = 0.2$, a 1 cent difference between 4.99 and 5.00 is perceived as a 20.8 cents gap, while the difference between 5.00 and 5.01 is only a 0.8 cent difference. Formally, if the true price is p, the perceived price is

$$\hat{p} = \hat{p}\left(p;\theta\right) = \left(1-\theta\right)p + \theta\left(\lfloor p \rfloor + \Delta\right) \tag{1}$$

where $\lfloor p \rfloor$ is the floor of the price. Demand is a function of perceived prices, while consumers pay the true price when purchasing an item. Assume a demand curve with constant elasticity. The demand curve is then

$$D\left(p;\theta\right) = A\hat{p}^{\epsilon} = A\left(\left(1-\theta\right)p + \theta\left(\lfloor p \rfloor + \Delta\right)\right)^{\epsilon} \tag{2}$$

which means that gross profits are of the form

$$\Pi\left(p;\theta\right) = D\left(p;\theta\right)\left(p-c\right)$$

where the assumption of a unit cost c is an approximation for supermarket pricing. It can be shown that prices under this demand structure bunch just below round numbers (denote those just below prices with q), with ranges of missing prices with low price-endings, i.e., excess mass below round prices is mainly drawn from higher prices. As shown by Strulov-Shlain (2022), the predicted ranges of missing prices are pinned down by the parameters of the problem.

Given the pricing function, we can focus on a key outcome—the *Next-Lowest Price*—which defines the range of missing prices given the parameters of the problem. The *Next-Lowest Price* P is the lowest price that should be charged that is greater than a just-below price q. The just-below price, q, ends with 99 before the reform and with 90 after. The Next-Lowest Price satisfies two conditions – it maximizes profits for some cost c, and for that cost profits are equal at P and at q. It can be shown that it is the solution of the following implicit equation:

$$P + \frac{(1-\theta)P + \theta\left(\lfloor P \rfloor + \Delta\right)}{\epsilon\left(1-\theta\right)} - \frac{\left((1-\theta)P + \theta\left(\lfloor P \rfloor + \Delta\right)\right)^{\epsilon} P - \left((1-\theta)q + \theta\left(\lfloor q \rfloor + \Delta\right)\right)^{\epsilon} q}{\left((1-\theta)P + \theta\left(\lfloor P \rfloor + \Delta\right)\right)^{\epsilon} - \left((1-\theta)q + \theta\left(\lfloor q \rfloor + \Delta\right)\right)^{\epsilon}} = 0 \tag{3}$$

[10] Δ is in a sense a demand level-shifter, and bears almost no influence on pricing.

Proposition 2 in Strulov-Shlain (2022) shows that given an elasticity $\epsilon < -1$, larger left-digit bias θ increases P. Therefore, since for $\theta = 0$ demand is smooth and $P = \lfloor P \rfloor$, there is some threshold $\theta_0(\epsilon)$ such that for $\theta > \theta_0(\epsilon)$, $P > \lfloor P \rfloor$. Meaning, there is some threshold $\theta(\epsilon)$ above which there should be no 00-ending prices. Since prices are discrete, I will ask what the corresponding θ_0 is for $P = \lfloor P \rfloor + 0.01$, or, in other words, for the next lowest price 1 cent above the round price. While there is no analytical solution for θ_0, solving for it numerically shows that for an extreme elasticity of -7 and $\lfloor P \rfloor = 2$, θ_0 is a mere 0.00075 (θ_0 will be even lower for less elastic demand and higher prices[11]). In contrast, typical left-digit bias estimates are between 0.1 and 0.5, and typical elasticities in these settings are about -1 to -4 (Butters et al., 2019; DellaVigna and Gentzkow, 2019; Hausman et al., 1994; Hitsch et al., 2017; Nevo, 2001). This calibration exercise implies that for any reasonable levels of left-digit bias and price elasticity, *if firms optimize*, we should observe no 00-ending pricing when 99-ending is available.

However, the reform made the trade-off stronger: The post-reform just-below price is 90-ending rather than 99-ending, and the next price above 00-ending is 10-ending rather than 01-ending. The former makes just-below prices provide lower revenue per item, and the latter makes an upward correction more constrained. Due to these two forces, the threshold levels of θ_0 are higher in the post-reform regime. I find θ_0 in the post-reform setting by asking for what parameters would $P > \lfloor P \rfloor + 0.1$, given that the just-below prices are now 90-ending. Although these two forces do have a combined strong effect on the minimal bias leading to no 00-ending prices, the levels are still very low. The threshold levels are depicted in Figure 1. Indeed, while before the reform with elasticity of -4 and $P_1 = 2$, the minimal bias is close to zero, after the reform the minimal bias becomes $\theta_0^{Post} = 0.02$, which is still an order of magnitude lower than levels found in such data, but not outside the scope of what firms might consider the bias to be.

Figure 1 shows threshold values of $\theta_0(\epsilon)$, above which there should be no 00-ending prices, before and after the reform. The dark lines are the pre-reform thresholds, showing that 00-ending prices should not be observed for almost any value. The light curves are for a post-reform setting, showing that the threshold indeed increases substantially, but not enough to justify 00-ending pricing given the estimated levels of θ in the data. The figure also overlays estimates from Strulov-Shlain (2022) showing that US retailers underestimate the bias, and as such, are getting closer to, but still above, the thresholds.

Therefore, the literature shows that there is economically significant left-digit bias, and the model predicts that with these levels of bias there should be bunching at 99- and 90-ending prices, and no 00-ending prices.

[11]Proposition 2 in Strulov-Shlain (2022) shows that less elastic demand increases $P - \lfloor P \rfloor$, and Corollary 1 shows that $P - \lfloor P \rfloor$ is increasing in $\lfloor P \rfloor$. Meaning, less elastic demand and higher prices will lead to an increase in $P - \lfloor P \rfloor$.

3 Price Patterns

Given that the model of pricing to left-digit biased consumers predicts no 00-ending prices, we expect that prediction to hold, and can learn from it failing. The actual patterns in the data inform us regarding how firms make decisions (as model-based or model-free learners), and I explain how in the next section, Section 4. This section shows the main patterns without interpretation. First, I describe aggregated and chain-level price-endings shares before, during, and after the reform. Second, I analyze price-endings paths at the product-store level. Finally, I estimate the impacts on demand from increasing a 99-ending price to 00-ending versus lowering it to 90-ending.

3.1 Price-endings shares

Figure 2 shows the shares of prices that end with 99, 90, and 00 by period. The shaded area represents the period between the announcement and enactment of the policy change.

Before the reform, the share of products that end with 99 is high and stable while the shares of 00 and 90 are very low. In both pricing data sets, the share of products that end with 99 before the policy announcement was about 45% and kept stable. In contrast, the shares of 00 and 90 were low, with 1.2% for 00 and 1.7% for 90.[12]

Immediately after the policy is announced but before it is enacted, firms start to change prices. Chains had 2.5 months to update prices, and 79% of products had a price change in December 2013. However, chains did not wait for the policy to bind to update prices and started changing prices away from 99-endings promptly.

As the policy starts to bind in January 2014, there are about 40% of prices that end with 90 and 20% that end with 00. Yet, a gradual and consistent downward trend in the share of 00-endings starts soon after.

A year after the reform, the share of 00-endings is lower, at about 5% in ICC data and 8% in CBS data; the share of 90-ending prices is potentially higher than the share of 99-ending prices before the reform. A sub-sample of StoreNext's data extends to 2019, showing the share of 00-endings consistently declining to 3% of prices by 2018–2019 (from roughly 20% in 2014–2015 in the same sample).

Between-chains heterogeneity While the decline in share of 00-endings is gradual on the aggregate, for some chains there is a step-wise update downward. Figure 3 shows the price-endings shares after the policy change for the biggest 6 chains. For most chains, the share of 00-endings behaves in an almost step-wise fashion, where for "Mega in the City" and "Rami Levy" (the second and

[12]The second and third most frequent price endings were 49 (6.6%) and 89 (3%).

third most largest chains in the market) the shares go to almost zero. "Shufersal" (the largest chain) with its two subsidiaries "Deal" (discount stores) and "Shelli" (urban stores) goes down to 10% of 00 price endings in mid 2014. "Wine Pavilion" and "Shufersal Deal" seem to also have fewer 00-endings by 2015, but the data ends too soon to see that convincingly.

The share of 00-endings prices therefore had been low before the reform, increased until its enactment, and since then gradually declined. In contrast, 99-endings, which were the just-below prices before the reform, were quickly and persistently replaced with 90-endings prices after the reform. But were some products priced at 90-endings and others at 00, or were many products priced at both price endings?

3.2 Price-endings paths

We now move from aggregated price patterns to consider the price-ending paths at the product-store level. The price-ending paths can inform us better about the learning processes used by firms. Different leaning processes can generate these aggregate patterns. I elaborate on these tests in Section 4, and the key differentiation between model-based and model-free learning is the contrast between testing multiple actions in model-based learning versus settling on "good" outcomes in model-free learning.

To see those patterns, consider the price path of a product in a store that was initially priced at some 99-ending. To simplify the problem, we consider three potential price-endings: 90, 00, or Other. That is, if a price p_{ist} of product i in store s at period t ends with 90, the current price-ending state is $pe_{ist} = 90$. The states history $h_{ist} = (pe_{is1}, pe_{is2}, ..., pe_{ist})$ is then categorized into price-ending path histories, as described by the automaton (state-machine) in Figure A-2.

The results of the exercise are shown in Figure 4, showing the shares of the different price-ending paths as different shades. The green represents 99, the blue shades 90 (or Other after 90), and the red shades 00 (or Other after 00).[13]

The main patterns are the following. First, products that changed their prices into 90 are unlikely to change, and specifically not into 00. To see that, compare the large and stable dark blue "90, never 00" to the small orange "00 after 90." Second, in contrast, products with a 00-ending price are unlikely to stay at 00, and highly likely to change to 90, e.g., compare the shrinking dark red "00, never 90" to the blue "90 after 00." Indeed, products that changed into 00 initially are 15 times more likely to end with 90 than 00 at the last period of data. [14] Third, there is a small set of products with prices ending with both 00 after 90, and 90 after 00, and those tend to eventually end at 90 (e.g., 99 → 90 → 00 → 90). This is evident in the growing light blue "90 after switching"

[13]The green are the products that ended with 99. Almost all of these changed at one point or another to either 90 or 00. Rarely is there a product that was priced at 99-ending and changed its price to anything else—in the figure, if changed to Other remains green at the post period.

[14]Products that changed into 00 initially are 3.5 times more likely to end with Other (not 90, not 00) than 00 at the last period.

versus small light orange “0 after switching.” At the last period, switching products are 8.4 times more likely to end with 90 than with 00.

Therefore, the price-ending analysis shows that product-stores that try a 90-ending are likely to just stay at 90, while 00-ending prices are unstable. There is some “experimentation” of switching between both, but that is relatively rare, and also tends to quickly resolve in 90-ending.

3.3 Demand response

In the previous sections we saw that there were many 00-ending prices early after the reform. A natural question to ask is if these prices had any impact on revenues. The model predicts that they are detrimental to demand and should be avoided, and the price updating behavior suggests that this is plausibly the case.

The ideal experiment to learn the effect of price-endings on demand is if price-endings were randomized. In a supermarket setting, it means to randomize prices within a product-store across different periods.[15] A second-best is an experiment in which comparable products' prices are randomized between different stores (the approach taken by Chetty et al. (2009)). The reform allows a quasi-experimental version of these ideal experiments. Consider the same product, priced the same way in two stores of the same chain before the reform. For example, it is priced at 5.99 a few times in 2013. For the former design, imagine that the products' price is switching between 5.90 and 6.00 within the same store. For the latter, imagine that after the reform a product is priced at 6.00 in one store and 5.90 in another. The goal of the following analysis is to leverage the reform to estimate these effects.[16]

I use the scanner data from StoreNext for the analysis., and estimate the following specification using instrumental variables for prices:

$$\log Q_{ist} = \alpha^{90}\kappa^{90}_{ist} + \alpha^{00}\kappa^{00}_{ist} + \alpha^{99}\kappa^{99}_{ist} + \beta_{ic(s)} \log P_{ist} + X'\beta + \epsilon_{ist} \tag{4}$$

where i is product, s is store, t is date (week), and c is chain. Q is the number of units sold and P is the average price. For each product-store in 2013 I define the modal 99-ending price $p^{Modal99}_{is}$. The αs are the price-endings fixed effects, such that κ^{99}_{ist} equals 1 if the price in week t equals $p^{Modal99}_{is}$ in 2013 (0 otherwise), κ^{90}_{ist} and κ^{00}_{ist} equal 1 if the price is 10-Agorot lower (ends with 90) or 1-Agora higher (ends with 00) than $p^{Modal99}_{is}$, respectively, in 2014 (0 otherwise). X includes various controls—a product-by-store fixed effect (to capture baseline store-product demand), a proxy for a product-store on-sale dummy (to capture transitory sales effects[17]), same-price spell-length decile

[15] If prices could have been randomized at the consumer level, that would also been ideal, but that is hard to implement in a brick-and-mortar store. In online marketplaces List et al. (2023) use an RD design using the smooth distribution of prices which creates quasi-random distribution around 00-ending prices, while Dubé and Misra (2023) randomize prices in a field experiment.

[16] Strulov-Shlain (2022) tests for these drops in demand using larger and higher quality data, but lacks the random variation.

[17] The dummy variable equals 1 if the price is lower by more than 3% than any price within 3 weeks (before or after).

by product fixed effect (to capture transitory sales effects), a month-of-year by product fixed effect (to capture seasonality), and year-by-product fixed effects.

I instrument for the product-chain elasticity with the leave-out average price of the product in stores of the same chain. The idea behind these instruments (often referred to as Hausman IVs after Hausman (1996)), is that the price within the chain in other stores captures shifts in costs but not in local demand shocks.

The meaning of this exercise is that the α coefficients capture the unexplained demand that is driven by the price-ending, controlling for the overall price elasticity. The most interesting comparison is between α^{90} and α^{00}. This comparison shows the difference in demand for "the same" product if it were priced at a 90-ending price versus a 10-Agorot higher 00-ending price in 2014, when price endings were more likely set at random. To a lesser degree, the comparison between a 99-ending price in 2013 and the 1-Agora larger 2014 00-ending price is also informative, but time is confounded with price-endings.[18]

Table 2 shows the results of estimating Equation 4. Overall, the concurrent excess difference in demand between a 90-ending price and a 10-Agorot higher 00-ending price, when the price was the nearest 99-ending before the reform, is estimated to be between 5%--9%. Meaning, *beyond* the price sensitivity of a 10-Agorot increase, demand is about 5%–9% lower. The difference in demand between a 99-ending price in 2013 and the 1-Agora higher 00-ending price in 2014, which is harder to interpret, is -1%–4%. Column (2) is closest to the first experiment, as it compares the demand at a 90-ending and a 00-ending for a product priced at, say, 5.99 in a store and then priced at 6.00 and 5.90 in the same store. The difference in demand is 8.2% with a p-value of 0.025. Columns (3) and (4) resonate the second approach, comparing demand for the same product priced, for example, at 5.99 in two stores and then priced at 6.00 in one store and at 5.90 in another. The difference in demand is 5.8% (p-value 0.061) in column 3, and 8.6% (p-value 0.0045) in column 4.[19]

With an average elasticity of about -2.5, and average price of 11, and 5%-8% drop corresponds to left-digit bias of $\hat{\theta} \approx 0.22 - 0.36$ (since $\theta \approx \frac{\%\text{drop}}{\text{elasticity}} \cdot \text{price}$). This estimate is depicted as X (Israeli demand) in Figure 1.

The model and previous literature predicted no 00-ending prices, and the demand analysis suggests that 00-ending pricing was a losing pricing strategy.

[18]For example, if many of the prices tend to end with 99 in 2013, the demand slope will be closer to the 99-ending prices due to the year fixed effects. Yet, removing the year fixed-effect will confound the price-ending effects with other year effects.

[19]Column 4 focuses on prices that were identical for at least two consecutive weeks to reduce price measurement error.

4 Model-based and Model-free Learning

4.1 Framework

Now we interpret the findings in Section 3 assuming that the true demand is left-digit biased and that firms, if using a model, use that correct structure[20]. The premise of this exercise is that all patterns must be taken into account while striving for the simplest explanation. The main three patterns are: (1) high shares of 99-endings before the reform and of 90-endings after it, low shares of 00-endings before the reform and long after; (2) high shares of 00-endings initially with negative demand impact; and (3) shares of 00-endings decreased consistently (and abruptly in some chains), and once a product-store was priced at 90-ending it was unlikely to be priced at 00-ending later.

First, low shares of 00 and high shares of 99/90 before and long after the reform imply acknowledging of left-digit bias. Bunching at just-below prices and avoiding low-ending prices are manifestations of correct pricing facing left-digit biased demand. Meaning that in the long run companies realize that just-below prices are better than 00-ending prices, even if they do not realize the extent of it (Strulov-Shlain, 2022).

Second, a high initial share of 00 and negative demand consequences imply ignorance about the value of 90 versus 00. Meaning, while firms priced as if they were aware of the relative benefit of 99-ending versus 00-ending, that knowledge did *not* extend to predict the relative value of 90-endings. In other words, although pre-reform pricing behavior is consistent with an as-if model of optimal pricing facing left-digit bias, it fails to predict behavior post-reform.

Third, downward trends of 00, chains' abrupt changes, and price ending paths are consistent with learning. It seems that over time chains realized that 00-endings are worse than 90-endings. The abrupt changes are consistent with "lightbulb" moments of sudden realization, and can rule out a slow adjustment process driven by frictions. Finally, the price-ending paths, in which 90-endings are an absorbing state while 00-ending are not—and further 00-endings are unlikely to follow 90-endings while the converse is not true—are in line with model-free reinforcement learning procedures.

Taken together, these patterns can inform us about how firms learn and therefore make decisions.

Model-based decision making is the consensual tool for modeling economic agents. It consists of two components—a function that translates parameters of a problem and actions into outcomes, and a maximization operator choosing the action delivering the best outcome. Importantly, in order to know what the best action is, the decision maker needs to learn about the parameters of

[20]That is, I do not entertain the possibility of a completely wrong model since it has no predictive power – if any model is possible all behavior is consistent. Indeed, any structural work in economics assumes some model. At the extreme, one can imagine that firms have a completely nonparametric model where each price has some associated demand without any structure related to it. That extreme model can be viewed as the "model-free" framework.

the target function, which leads to model-based learning. The goal of learning in a model-based world means solving the uncertainty about the parameters of the problem. For example, when a monopolist sets a price to maximize profits, it needs to know the parameters of the problem, which are the price elasticities and marginal costs of the products it sells (and potentially left-digit bias).[21] That is, model-based learning updates the expected value of every possible action.

Compare the above to model-free decision making, which is a private case and result of reinforcement learning (Sutton and Barto (2018)). The model-free reinforcement learning system is comprised of an environment of actions and states, and three components: a reward signal – the feedback from taking an action; a value function – the benefits of an action given the state and predicted behavior and evolution of the environment from taking an action (or the optimal action); and a policy – defining which action to take based on beliefs and state. Similarly to model-based learning, in model-free learning, the agent has some beliefs about the value of different actions given the state of the world. The beliefs are updated based on the reward signals, and a policy will choose the highest-value action. However, in contrast to model-based learning, resolving uncertainty means only learning about the value of a particular action taken. For example, assume that a company expects to make \$10 in profits if the price is \$1 and \$12 if the price is \$2. If it sets the price at \$2 and makes \$8, the company will update the value of pricing at \$2, but will not change its expectations about the value of pricing at \$1. That is, model-free learning informs the value of actions directly. Indeed, inferring from one action to another is not model-free, it requires some form of a "model", one that allows us to interpolate and extrapolate.

In theory, full learning can be achieved by using either model-free or model-based learning, making behavior indistinguishable in the long-run. For example, Q-learning, a leading model-free reinforcement learning rule, converges to the correct values if every action is tried in every state infinitely often and if new estimates, from trials of actions, are blended with previous ones using a slow enough exponentially weighted average (Dayan, 1992; Watkins and Dayan, 1992). Even during the learning process, distinguishing the processes requires some clever manipulations (Daw et al., 2011; Kurdi et al., 2019).

This paper explores a different identification strategy. If learning has subsided but "full learning" had not been achieved, the behavior can differ in the following way: If there is no full learning, a change in the action set alone may require learning in the model-free environment, can spur experimentation in a wrong model-based environment, but will not lead to experimenatation or learning in the correct model-based world.[22] To illustrate, consider the above pricing example and imagine that costs doubled, thus requiring higher prices that were not tested before. The model-based monopolist will have a set of beliefs about the demand function shape and therefore will know the best price, while the model-free learners will rely on an uninformed prior and

[21] Recent empirical examples are reviewed in Aguirregabiria and Jeon, 2020.

[22] A change in the action set may require learning, depending on the exact algorithm used. In machine learning, model-free learning algorithms that are "generalized" (e.g., Rummery and Niranjan, 1994), for example, use neural networks to extend value predictions to actions that had not been tested. However, I think of the monopolist problem as closer to the discrete case.

have to start learning again. Since 90-endings were rarely charged, the reform provides such an experiment, thus allowing one to separate between the methods.

4.2 Partial learning

Since the model-based framework is in consensus, I describe here an attempt to reconcile the data patterns with firms' behavior. This exercise is similar in spirit to Doraszelski et al. (2018); Huang et al. (2018). I assume that firms know that demand is left-digit biased, but allow them to be wrong about the level of left-digit bias. If firms think that bias is small (and demand is very elastic), then pricing at 00-ending can be perceived as optimal even though it is accompanied by demand losses. Wrong beliefs will not be able to generate excess shares of 00-endings, but even ignoring the excess share, can they support any 00-endings?

In principle, underestimated left-digit bias might support 99-endings versus 00, but not 90 versus 00, because of the stronger trade-off. However, as Figure 1 shows, these beliefs about left-digit bias and elasticity must be extreme. Yet, we can estimate beliefs about left-digit bias before the reform, using the same algorithm as in Strulov-Shlain (2022), and see if they allow for 00-endings after the reform—at least in principle.

The algorithm uses the empirical distribution of prices as moments and predicts the price distribution according to the model. The excess shares in 99-endings and the distribution of lower-ending prices identify the parameters. In Appendix Section C I describe the algorithm in detail.[23]

Although the perceived bias in the left digits is much smaller than the demand-estimated level, it is still higher than the thresholds. I estimate that the perceived left-digit bias ranges between 0.02-0.04 while the threshold values of left-digit bias that can support 00-ending prices are less than 0.01. In other words, supermarkets price as if they believe that drops in demand from a 1c increase from 99-ending to 00-ending is equivalent to a 3-5 cent price increase, while a belief of a 2c-equivalent increase is enough to avoid 00-endings. Importantly, the demand-side estimated left-digit bias is 0.22-0.36 (e.g, the impact of a 1c increase from 99-ending to 00 is like a 23-37 cent equivalent price change). The finding of underestimation of bias is common in the US market as well (Strulov-Shlain, 2022), and possibly not surprising by itself. However, even an underestimation cannot support 00-endings after the reform serves the main case against full knowledge of the shape of demand. That is, estimated beliefs before the reform are inconsistent with pricing behavior after the reform under the null of knowing the shape of demand.

The finding that after the reform, the 00 ending rarely follows the 90 ending at the product-store level suggests, though not conclusive by itself, inconsistency with model-based learning as well. If firms are trying to learn the left-digit bias parameter, they need to randomize price endings.

[23] Appendix section C is copied from Strulov-Shlain (2022).

The natural way to do so is to change prices into 90- and 00-endings in a methodological way. Yet, that the vast majority of products stay at 90-endings even if they cannot be compared to 00-endings suggests reliance on recent feedback. This evidence is weaker, since there are multiple ways to experiment (e.g., randomize prices between stores and not within product-store).

Finally, an *excess* of 00-ending prices and 90-ending prices at the same time is in contrast to any belief of left-digit bias. Excess 00-endings require beliefs about the shape of demand that mixes positive and negative left-digit bias. In contrast, pre-reform pricing is only consistent with beliefs of positive left-digit bias. Meaning, pre-reform model-based beliefs are inconsistent not only in magnitude but also in sign.

Taken together, the evidence is inconsistent with full learning that had occurred over the years leading to the reform. Therefore, long-run behavior also implies partial knowledge.

5 Fudge Factor

The findings show a gap in optimization, associated with prolonged partial knowledge. If researchers try to analyze companies' behavior under the wrong assumption of full knowledge and optimization, it would lead us to draw wrong conclusions. However, companies may not optimize for a multitude of reasons. For example, one reason is partial knowledge and lack of experimentation (as in this paper and see also supporting evidence in Berman and Israeli (2022); Bloom and Van Reenen (2010); Huang (2021); Koning et al. (2022); Zhang and Misra (2022)); another is organizational frictions and agency problems (DellaVigna and Gentzkow, 2019; Hortacsu et al., 2021; Rao and Simonov, 2019; Shapiro et al., 2021); yet another is psychological biases of managers (Malmendier and Tate, 2005, 2015); and many others, of course, exist. At the same time, we wish to maintain the hypothesis that companies strive to optimize, even if they fail to do so perfectly.

I suggest a method to capture partial optimization and allow testing of the sensitivity of the results to misoptimization. A *Fudge factor* $X \in [0, 1]$ captures the percentage of potential profits that the firm forgoes. Given X, we ask what unobservables are consistent, under the model, with the observed behavior such that the observed behavior brings at least $1 - X$ of the optimum. Formally, assume that the objective function is some $\Pi(p, a)$, where a is a state variable and p is an action the firm can take to maximize Π. Assume that a is the unknown parameter of interest. The standard assumption is that an observed p is a function of a such that $p = p^*(a) = \arg\max \Pi(p, a)$, which then allows finding $a = \Pi^{-1}(p^*, a)$. What I propose is to assume that p is not the maxmizer but a partial-maximizer, and find all a such that $\Pi(p, a) \geq (1 - X)\,\Pi(p^*(a), a)$. That is, assume that p is consistent with partial optimization such that the profits under a and p are at least $1 - X$ the profits under the same a and an optimal $p^*(a)$.

For example, assume that a monopolist faces a profit function $D(p)(p-c)-F$ and the parameter of interest is c. Therefore,

$$D(p)(p-c)-F \geq (1-X)\left(D(p^*(c))(p^*(c)-c)-F\right)$$

Although in general this inequality is not analytically solvable, it provides an implicit function that can be solved numerically. Further, if Π is convex the solution will be a segment of costs.[24]

The fudge factor is similar to the Critical Cost Efficiency Index (CCEI) pioneered by Afriat (1972) and extended and elaborated by others (Choi et al., 2014; de Clippel and Rozen, 2021; Halevy et al., 2018; Varian, 1990). Afriat's idea was to think about differences in productivity between firms, given a known production function. The CCEI then measures how far, in terms of production efficiency, a given firm is from the frontier. Similarly, he introduced an index for violations of GARP (the General Axiom of Revealed Preference) to measure how "far" are a set of choices from being able to be represented by a utility function. The follow up work then took this individual choices index and expanded it to measures in terms of budget constraints and other money metrics. For example, Halevy et al. (2018) treat the index as a loss metric that allows to do model selection based on minimizing this loss. Specifically, to find the model that best describes participants' behavior, we seek to find the model that minimizes out-of-sample loss, as measured by this index. Another similarity is with the computer science and operation research literature, which often deals with complicated optimization problems. In these problems, which cannot be fully solved, the solution concept is often an approximation. A proposed algorithm is then judged on the basis of how well it approximates the full solution or by minimizing the maximal regret. In one way or another, firms often face complicated problems and it is almost necessary that they only approximate their solutions rather than fully optimize them. Unlike this literature, I assume that the correct model is known, and ask either what are the implications of a known X on the range of recovered parameters, and how sensitive are conlusions to varying levels of X.

As noted above, different reasons may lead to $X > 0$. For example, a problem that is too complicated to solve routinely will lead to approximations (Camara, 2021; Hortacsu et al., 2021); partial knowledge would make it unlikely to find the optimal action (this paper or Bloom et al. (2013); Dubé and Misra (2023); Hitsch et al. (2017); Zhang and Misra (2022)); organizational frictions in which a principal only monitors an agent if they deliver catastrophic results and otherwise trusts them; low belief in the importance of an attribute and lack of competition on specific margins would likely lead a constrained firm to not explore and try to optimize leading to potentially large Xs (Gagnon-Bartsch et al., 2021). Therefore, the idea of X is to capture a whole host of explanations,

[24]For example, if $F = 0$ and $D(p) = A - Bp$, then $p^*(c) = \frac{A+Bc}{2B}$ and we can find an analytical, albeit unappealing, solution for the range of costs c that will rationalize an observed p given known demand parameters A and B and a fudge factor X:

$$c \in \left[\frac{2Bp - 2(A-Bp)\sqrt{X} - A(1+X)}{B(1-X)}, \frac{2Bp + 2(A-Bp)\sqrt{X} - A(1+X)}{B(1-X)}\right]$$

and its value will be determined by the specific setting. For example, we might believe that in more competitive environments, or for standard, high feedback, important, frequent decisions X will be smaller.

How big is X? In most cases researchers assume optimization because it is impossible to identify deviations from optimality. Yet, some examples should move our priors away from $X = 0$. For example, Blake et al. (2015) find ROI of -63% for eBay Search Enginge Marketing expenditures, and Shapiro et al. (2021) study well established CPG manufacturers and find that the median ROI from TV advertising on a given week is -88% and the overall ROI of advertising is -57%. Given typical shares of advertising expenditures, these lead to a loss of a few percents of the overall profit (if these were the only mistakes these companies were making); Dubé and Misra (2023) find that ZipRecruiter's pricing before their pricing experiment led to 35% lower profits compared to uniform pricing; Zhang and Misra (2022) find that before randomization coupons sent by a food delivery platform led to 58% lower profits compared to the best uniform promotion depth; DellaVigna and Gentzkow (2019) estimate a profit loss of approximately 10% from uniform pricing across supermarket stores; Hortacsu et al. (2021) find that organizational structure and pricing heuristics lead to 13% revenue loss compared to a unified optimization benchmark; Cho and Rust (2010) demonstrate the rental companies forgo 20%-30% of profits with their fleet management practices. Finally, two closely related papers show the importance of left-digit bias in pricing: Strulov-Shlain (2022) finds that the partial response to left-digit bias causes large US supermarket chains to forgo 3%-5% of gross profits; and List et al. (2023) show that Lyft was completely ignoring the importance of left-digit bias and could have made $160M more in profits. Although the lost profits are not the main point of some of these papers, these are a selected sample. In general, across companies and time, it seems likely that X is much closer to 0 than these papers find. However, together these papers document, for a myriad of reasons, instances where X was much higher than 0 for many sophisticated companies and for sustained periods of time. Therefore, acknowledging the possibility of a positive X in any given analysis seems justified.

For example, when studying left-digit bias we care about the magnitude of the bias as it is manifested in demand. However, estimating it from demand requires high-quality data on multiple prices and quantities, which is often unavailable. Evidence shows that firms do not respond enough to left-digit bias. With the fudge factor, even if we only observe pricing behavior, we can learn the magnitude of left-digit bias as a function of X. Some papers, such as List et al. (2023); Strulov-Shlain (2022), study left-digit bias among Lyft and 25 national US supermarket chains, respectively, can estimate the bias directly on the demand side and compare it to companies' behavior. However, if we consider the current paper, estimating the left-digit bias of demand is only suggestive due to data limitations. We can still ask how large the actual left-digit bias of demand can be, such that the observed firm behavior will lead to at most X% loss.[25]

[25]work in progress

Another use case is to consider the smallest X that will revert counterfactual conclusions. For example, assume that one studies changes in markups over time, based on estimating price elasticities and assuming optimization, and finds an increase in markups due to lower elasticities. A more sophisticated version of the toy example above can allow one to determine that if X is greater than some threshold, constant markups cannot be rejected. If this threshold is 1%, 10%, or 30%, it will tell us how sensitive the finding is to the full optimization assumption.

6 Discussion

I find that following a pricing reform, supermarket chains acted in a way inconsistent with full knowledge and optimization. The host of evidence on the dynamics of price setting before and after the reform, and our priors and evidence on demand response, suggest that firms were facing left-digit biased demand, yet operated under partial optimization. This paper joins other papers in documenting sub-optimal behavior and suggests that sustained partial knowledge leads to this gap in behavior.

How do firms preserve partial knowledge? They can do so, as in Hanna et al. (2014); Gagnon-Bartsch et al. (2021); Schwartzstein (2014), if they do not think that actions matter and never try to learn if they do. In this paper, I document that they have some awareness to the influence of price endings on demand, but they clearly are not trying to figure out the exact form of this influence.

Critiques There are several limitations to this study. First and foremost, it is merely a single example. One objection might be that this is not a representative case because price endings do not have a big effect on pricing. In other words, it is an unimportant example. Indeed, the view that left-digit bias is not important is common, including by retailers themselves. I interviewed several retail executives in the US and Israel, and they all think it is more of a norm than they make a calculated pricing decision. However, this view is misguided since under-appreciation of left-digit bias has large effects on profits (List et al., 2023; Strulov-Shlain, 2022). Furthermore, the fact that it might be viewed as unimportant supports the partial knowledge hypothesis. The most extreme example of perceived unimportance is in List et al. (2023), in which a (very sophisticated) company is not aware of the existence of the bias, completely ignores it in its pricing decisions, and loses money for it (following the research, the firm had changed its pricing to be mostly 99-ending). Here, firms are aware to some degree, but only partially.

A related objection is that this example is not a representative case because it is an outlier, and firms often do optimize. However, optimality is often assumed, not tested. A growing body of work shows cases in which companies persistently act in a sub-optimal manner (partial list - Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Goldfarb and Xiao, 2011; Hanna et al., 2014; Hortaçsu et al., 2019; Hortacsu et al., 2021; Huang, 2021; Shapiro et al., 2021; Strulov-Shlain, 2022).

Furthermore, partial learning can help answer the common question raised in those papers of "how firms are not optimizing."

Another limitation is that other models can be created to explain the data. For example, if a firm has beliefs about left-digit bias that are state-dependent, then any pricing behavior is possible—if at periods of change firms revert to a null of no bias (similar to Enke and Graeber (2019)), then the observed behavior is consistent with such a model. However, if we expand the model to agree with the data whenever the data do not agree with the model, then the notion of a "model" loses its bite.

While some readers might find these conclusions hard to accept, others might respond that firms' misoptimizing is obvious, but model-based is useful.[26] I concur with both claims. Therefore, to preserve the assumption of attempted optimization, I suggest adding a "fudge factor" to analyses in order to test the sensitivity of assumptions to partial optimization.

Implications and avenues for future research If firms are partial optimizers, it means that they may act in a persistent way in steady state, without knowing in advance how to respond to new regulations or changes in the market place. A change in the marketplace, such as the one in this paper, might jolt companies out of their steady state and into a learning mode (Gagnon-Bartsch et al., 2021). These changes occur a lot: New establishments open and close, rivals enter and leave, tastes shift, and costs fluctuate. Treating firms as all-knowing is bound to miss many interesting dynamics and potentially lead to wrong conclusions. For example, when analyzing the effects of a change in the market (say a new regulation or a new competitor entry), researchers should incorporate a learning period, which might be lengthy.

In the long run, firms may behave suboptimally consistently (e.g., Ely (2011) analyzes how path dependence may lead to suboptimal steady state behavior; Thrun and Schwartz (1993) show when reinforcement learning algorithms might fail). This is especially true due to partial knowledge, and further investigation of this avenue as a probable explanation can help solve cases in which it is unclear how firms consistently deviate from optimization.

To conclude, there is ample room for future research. Empirically, more examples and tests of partial knowledge should be explored; the effects of market forces on firms' suboptimal behavior is a topic of current work. Methodologically, the usefulness of the fudge factor can be studied and challenged. Theoretically, what are the implications of incorporating partial knowledge and optimization into problems of mechanism design, on how regulation might affect markets, or on welfare?

[26]In an episode of the Freakonomics podcast discussing Shapiro et al. (2021), the economist Steve Levitt says "any economist who tells you that firms are profit-maximizing has never worked with firms. That's a simple model we use when we teach beginning economics because it's easy to solve mathematically."

References

Afriat, S N, "Efficiency Estimation of Production Functions," *Int. Econ. Rev.*, 1972, *13* (3), 568–598.

Aguirregabiria, Victor and Jihye Jeon, "Firms' Beliefs and Learning: Models, Identification, and Empirical Evidence," *Review of Industrial Organization*, March 2020, *56* (2), 203–235.

Anderson, Eric T. and Duncan I. Simester, "Effects of $9 Price Endings on Retail Sales: Evidence from Field Experiments," *Quantitative Marketing and Economics*, 2003, *1* (1), 93–110.

Anderson, S J, P Chintagunta, F Germann, and others, "Do marketers matter for entrepreneurs? Evidence from a field experiment in Uganda," *Journal of*, 2021.

Arthur, W Brian, "Designing Economic Agents that Act like Human Agents: A Behavioral Approach to Bounded Rationality," *Am. Econ. Rev.*, 1991, *81* (2), 353–359.

Assad, S, R Clark, D Ershov, and L Xu, "Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market," 2020.

Assad, Stephanie, Emilio Calvano, Giacomo Calzolari, Robert Clark, Vincenzo Denicolò, Daniel Ershov, Justin Johnson, Sergio Pastorello, Andrew Rhodes, Lei Xu, and Others, "Autonomous algorithmic collusion: Economic research and policy implications," 2021.

Ater, I and O Gerlitz, "Round prices and price rigidity: Evidence from outlawing odd prices," *J. Econ. Behav. Organ.*, 2017.

Ater, Itai and Oren Rigbi, "Price Transparency, Media and Informative Advertising," 2019.

Barberis, Nichloas and Lawrence Jin, "Model-free and Model-based Learning as Joint Drivers of Investor Behavior," 2021.

Basu, Kaushik, "Why are so many goods priced to end in nine? And why this practice hurts the producers," *Economics Letters*, 1997, *54* (1), 41–44.

— , "Consumer Cognition and Pricing in the Nines in Oligopolistic Markets," *Journal of Economics & Management Strategy*, 2006, *15* (1), 125–141.

Berman, R and A Israeli, "The value of descriptive analytics: Evidence from online retailers," *Marketing Science*, 2022.

Bizer, George Y. and Robert M. Schindler, "Direct Evidence of Ending-Digit Drop-Off in Price Information Processing," *Psychology and Marketing*, 2005, *22* (10), 771–783.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Consumer heterogeneity and paid search effectiveness: A large-scale field experiment," *Econometrica*, January 2015, *83* (1), 155–174.

Bloom, N and J Van Reenen, "Measuring and explaining management practices across firms and countries," *Q. J. Econ.*, 2007.

— **and** — , "Why do management practices differ across firms and countries?," *J. Econ. Perspect.*, 2010.

— **, B Eifert, A Mahajan, and others**, "Does management matter? Evidence from India," *The*

Quarterly journal, 2013.

Busse, Meghan R, Nicola Lacetera, Devin G Pope, Jorge Silva-Risso, and Justin R Sydnor, "Estimating the effect of salience in wholesale and retail car markets," *American Economic Review*, 2013, *103* (3), 575–79.

Butters, R, Daniel W Sacks, and Boyoung Seo, "Why Don't Retail Prices Vary Seasonally with Demand?," *Kelley School of Business Research Paper*, 2019.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello, "Artificial Intelligence, Algorithmic Pricing, and Collusion," *Am. Econ. Rev.*, October 2020, *110* (10), 3267–3297.

Camara, M K, "Computationally Tractable Choice," *mimeo*, 2021.

Camerer, Colin F, *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, September 2011.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Cho, Sungjin and John Rust, "The flat rental puzzle," *The Review of Economic Studies*, 2010, *77* (2), 560–594.

Choi, S, S Kariv, W Müller, and D Silverman, "Who is (more) rational?," *Am. Econ. Rev.*, 2014.

Cross, John G, "A stochastic learning model of economic behavior," *Q. J. Econ.*, May 1973, *87* (2), 239.

Daw, Nathaniel D, Samuel J Gershman, Ben Seymour, Peter Dayan, and Raymond J Dolan, "Model-based influences on humans' choices and striatal prediction errors," *Neuron*, March 2011, *69* (6), 1204–1215.

Dayan, Peter, "The convergence of TD(?) for general ?," *Mach. Learn.*, May 1992, *8* (3-4), 341–362.

de Clippel, G and K Rozen, "Relaxed Optimization: How Close Is a Consumer to Satisfying First-Order Conditions?," *Rev. Econ. Stat.*, 2021.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform Pricing in US Retail Chains," *The Quarterly Journal of Economics*, 06 2019.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes, "Just Starting Out: Learning and Equilibrium in a New Market," *Am. Econ. Rev.*, 2018, *108* (3), 565–615.

Dube, Arindrajit, Alan Manning, and Suresh Naidu, "Monopsony and Employer Mis-optimization Account for Round Number Bunching in the Wage Distribution," *Unpublished manuscript*, 2017.

Dubé, Jean-Pierre and Sanjog Misra, "Personalized Pricing and Consumer Welfare," *J. Polit. Econ.*, January 2023, *131* (1), 131–189.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo, "Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?," *QME*, 2010, *8* (2), 207–239.

Ely, Jeffrey C, "Kludged," *American Economic Journal: Microeconomics*, 2011, *3* (3), 210–31.

Enke, Benjamin and Thomas Graeber, "Cognitive Uncertainty," Technical Report, National Bureau of Economic Research 2019.

Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein, "Channeled Attention and Stable Errors," 2021.

Gedenk, Karen and Henrik Sattler, "The impact of price thresholds on profit contribution should retailers set 9-ending prices?," *Journal of Retailing*, 1999, *75* (1), 33–57.

Ginzberg, Eli, "Customary prices," *The American Economic Review*, 1936, *26* (2), 296–296.

Goldfarb, Avi and Mo Xiao, "Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets," *American Economic Review*, 2011, *101* (7), 3130–61.

_ **and** _ , "Transitory shocks, limited attention, and a firm s decision to exit," 2019.

Halevy, Yoram, Dotan Persitz, and Lanny Zrill, "Parametric Recoverability of Preferences," *J. Polit. Econ.*, August 2018, *126* (4), 1558–1593.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein, "Learning through noticing: Theory and evidence from a field experiment," *The Quarterly Journal of Economics*, 2014, *129* (3), 1311–1353.

Hausman, Jerry A, "Valuation of new goods under perfect and imperfect competition," in "The economics of new goods," University of Chicago Press, 1996, pp. 207–248.

Hausman, Jerry, Gregory Leonard, and J Douglas Zona, "Competitive analysis with differenciated products," *Annales d'Economie et de Statistique*, 1994, pp. 159–180.

Heidhues, Paul and Botond Kőszegi, "Chapter 6 - Behavioral Industrial Organization," in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications 1*, Vol. 1 of *Handbook of Behavioral Economics: Applications and Foundations 1*, North-Holland, 2018, pp. 517 – 612.

Hendel, Igal, Saul Lach, and Yossi Spiegel, "Consumers' activism: the cottage cheese boycott," *The RAND Journal of Economics*, 2017, *48* (4), 972–1003.

Hilger, Nathaniel, "Heuristic Thinking in the Market for Online Subscriptions," *Available at SSRN 3296698*, 2018.

Hitsch, Günter J, Ali Hortacsu, and Xiliang Lin, "Prices and promotions in us retail markets: Evidence from big data," *Chicago Booth Research Paper*, 2017, (17-18).

Hortaçsu, Ali and Steven L Puller, "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market," *Rand J. Econ.*, March 2008, *39* (1), 86–114.

Hortaçsu, Ali, Fernando Luco, Steven L Puller, and Dongni Zhu, "Does strategic ability affect efficiency? Evidence from electricity markets," *American Economic Review*, 2019, *109* (12), 4302–42.

Hortacsu, Ali, Olivia Nata, Hayden Parsley, Timothy Schwieg, and Kevin Williams, "How do Pricing Algorithms Affect Allocative Efficiency? Evidence from a Large U.S. Airline," 2021.

Huang, Yufeng, "Pricing Frictions and Platform Remedies: The Case of Airbnb," June 2021.

__ , **Paul B Ellickson, and Mitchell J Lovett**, "Learning to set prices in the washington state liquor market," *Available at SSRN 3267701*, 2018.

Koning, Rembrand, Sharique Hasan, and Aaron Chatterji, "Experimentation and Start-up Performance: Evidence from A/B Testing," *Manage. Sci.*, September 2022, *68* (9), 6434–6453.

Kurdi, Benedek, Samuel J Gershman, and Mahzarin R Banaji, "Model-free and model-based learning processes in the updating of explicit and implicit evaluations," *Proc. Natl. Acad. Sci. U. S. A.*, March 2019, *116* (13), 6035–6044.

Lacetera, Nicola, Devin G. Pope, and Justin R. Sydnor, "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, 2012, *102* (5), 2206–2236.

Li, H and X Qiu, "Heuristics in Self-Evaluation: Evidence from the Centralized College Admission System in China," *Rev. Econ. Stat.*, 2023.

List, John A, Ian Muir, Devin Pope, and Gregory Sun, "Left-Digit Bias at Lyft," *the Review of Economic Studies*, 2023, p. rdad014.

Malmendier, Ulrike and Geoffrey Tate, "CEO overconfidence and corporate investment," *J. Finance*, December 2005, *60* (6), 2661–2700.

__ **and** __ , "Behavioral CEOs: The Role of Managerial Overconfidence," *J. Econ. Perspect.*, November 2015, *29* (4), 37–60.

Muth, John F, "Rational Expectations and the Theory of Price Movements," *Econometrica*, 1961, *29* (3), 315–335.

Nevo, Aviv, "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 2001, *69* (2), 307–342.

Rabin, Matthew, "An Approach to Incorporating Psychology into Economics," *Am. Econ. Rev.*, May 2013, *103* (3), 617–622.

Rao, J M and A Simonov, "Firms' reactions to public information on business practices: The case of search advertising," *Quantitative Marketing and Economics*, 2019.

Rummery, Gavin A and Mahesan Niranjan, *On-line Q-learning using connectionist systems*, Vol. 37, Citeseer, 1994.

Schwartzstein, Joshua, "Selective attention and learning," *J. Eur. Econ. Assoc.*, 2014, *12* (6), 1423–1452.

Shapiro, Bradley T, Günter J Hitsch, and Anna E Tuchman, "TV advertising effectiveness and profitability: Generalizable results from 288 brands," *Econometrica*, 2021, *89* (4), 1855–1879.

Sokolova, Tatiana, Satheesh Seenivasan, and Manoj Thomas, "The Left-Digit Bias: When and Why Are Consumers Penny Wise and Pound Foolish?," *J. Mark. Res.*, August 2020, *57* (4), 771–788.

Spiegler, Ran, *Bounded rationality and industrial organization*, OUP USA, 2011.

Stiving, Mark, "Price-Endings When Prices Signal Quality," *Management Science*, 2000, *46* (12), 1617–1629.

— **and Russell S. Winer**, “An Empirical Analysis of Price Endings with Scanner Data,” *Journal of Consumer Research*, 1997, *24* (1), 57–67.

Strulov-Shlain, Avner, “More Than a Penny’s Worth: Left-Digit Bias and Firm Pricing,” *The Review of Economic Studies*, 12 2022, p. rdac082.

Sutton, Richard S and Andrew G Barto, *Reinforcement Learning, second edition: An Introduction*, MIT Press, November 2018.

Thomas, Manoj and Vicki Morwitz, “Penny Wise and Pound Foolish: The Left-Digit Effect in Price Cognition,” *Journal of Consumer Research*, 2005, *32* (1), 54–64.

Thorndike, Edward L, “Animal intelligence; experimental studies,” 1911.

Thrun, Sebastian and Anton Schwartz, “Issues in using function approximation for reinforcement learning,” in “Proceedings of the Fourth Connectionist Models Summer School” books.google.com 1993, pp. 255–263.

Varian, Hal R, “Goodness-of-fit in optimizing models,” *Journal of Econometrics*, 1990, *46* (1-2), 125–140.

Watkins, Cjch and Peter Dayan, “Q-learning,” *Mach. Learn.*, 1992.

Xie, Erhao, “Inference in games without Nash equilibrium: An application to restaurants’ competition in opening hours,” Technical Report 2018-60, Bank of Canada Staff Working Paper 2018.

Zhang, Walter W and Sanjog Misra, “Coarse Personalization,” April 2022.

Table 1: Summary Statistics of Prices Data

	ICC		CBS		StoreNext	
Variable	Original	Main	Original	Main	Original	Main
No. of Products	67	21	233	171	19	17
No. of Stores	659	143	155	99	"580"	"573"
No. of Chains	38	20	–	–	"23"	"23"
Sampling Periods	41	40	48	48	156	156
First Observation	02/2013	03/2013	01/2012	01/2012	01/2013	01/2013
Last Observation	02/2015	02/2015	12/2015	12/2015	12/2015	12/2015
Mean Weekly Units Sold	–	–	–	–	56.69	66.09
Mean Price	12.99	6.95	10.48	10.04	12.36	11.06
Price SD	1.59	0.89	1.88	1.3	6.2	4.56
Pre 99 Share	0.38	0.45	0.36	0.47	0.13	0.23
Post 90 Share	0.52	0.59	0.39	0.48	0.29	0.39
Post 00 Share	0.12	0.10	0.09	0.12	0.15	0.15
Observations	472,572	116,995	292,471	156,336	846,392	185,775

The table shows summary statistics of the three datasets used in the paper, in their raw and cleaned form. ICC, Israeli Consumer Council, collected a panel of prices at the item-store-biweek level. CBS, Israel's Central Bureau of Statistics, collected prices for the consumer price index at the "product type"-store-month level. StoreNext, is a market research firm, collecting scanner data from retailers including quantities and revenues at the store-week-item level. These data are a merge of two datasets, and thus the exact number of stores is bounded between 295 and 580, but is more likely to be much closer to 295, and the number of chains closer to 12 than to 23.

Table 2: Price-ending effects

	Quantity (log-units)			
	(1)	(2)	(3)	(4)
α^{90}	−0.029	0.006	−0.037	−0.018
	(0.024)	(0.033)	(0.024)	(0.023)
α^{00}	−0.082***	−0.077**	−0.095***	−0.104***
	(0.028)	(0.032)	(0.027)	(0.029)
α^{99}	−0.058***	−0.091***	−0.062***	−0.066***
	(0.014)	(0.025)	(0.014)	(0.014)
90 minus 00	0.053	0.082	0.058	0.086
[p-value]	[0.083]	[0.025]	[0.061]	[0.0045]
99 minus 00	0.024	-0.014	0.034	0.038
[p-value]	[0.44]	[0.68]	[0.23]	[0.2]
Leave-out price IV	Yes	Yes	Yes	Yes
Spell-length restrictions				2+
99 prevalence 2013	1+	5+	5+	5+
90 or 00 post prevalence		5+ for both	5+ for either	5+ for either
N	176,996	176,996	176,996	149,784
R^2	0.859	0.859	0.859	0.872

The table shows the results from estimating equation 4 under different restrictions and samples. In all columns, a product-store is kept only if there is at least one occurrence of a 99-ending price in 2013. Further, in columns (2)--(4), $\kappa^{99}_{ist} = 1$ only if the price ends with 99 at period t and it ends in 99 for at least 5 weeks in 2013 for product i in store s. Sample restrictions in column (4) keep observations only if prices are the same for at least 2 consecutive weeks. In column (2) $\kappa^{90} = 1$ or $\kappa^{00} = 1$ if the price ends in 90 or 00 in 2014, at least 5 times *each* for the same product i in the same store s. Meaning, this is estimating the effect on items that switched price endings within the same store over time. In columns (3)--(4) $\kappa^{90} = 1$ or $\kappa^{00} = 1$ if the price ends in 90 or 00 in 2014, at least 5 times for *either* 90 or 00 in the same store. Meaning, this is comparing price-ending effects between "99 to 00" stores and "99 to 90" stores.

Figure 1: Minimal levels of left-digit bias $\theta_0(\epsilon)$ above which 00-ending is suboptimal

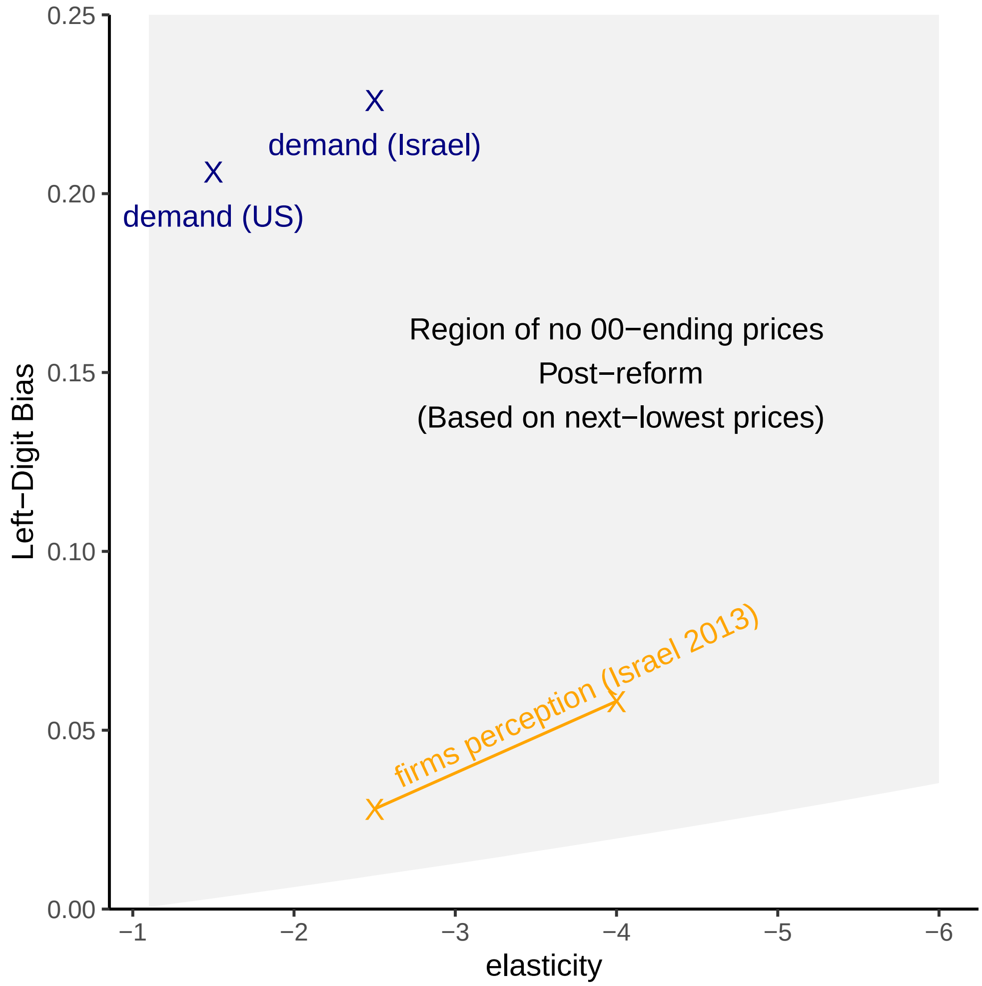

The figure shows the minimal levels of left-digit bias θ above which 00-ending prices are never optimal. The values are a function of the price elasticity and nominal price level. The dark line shows the threshold values in the pre-reform case (when 99-ending and 01-ending prices were allowed), and the light line shows post-reform values (where 90-ending is the highest price ending and 10-ending is the lowest price ending above 00). The gray area is the parameter set for which there should be no 00-ending prices after the reform. The "X demand (Israel)" is an approximation of the level of left-digit bias found in demand, given the drops in demand estimated from the reform and the price elasticity. "Firms perception (Israel 2013)" shows the range of estimates regarding firms' beliefs about the level of left-digit bias consumers have, that might rationalize their pricing behavior before the reform. The "X demand (US)" marks the demand-side estimates of left-digit bias, and "Firms perception (US)" the range of beliefs about left-digit bias that U.S. firms hold, as estimated in Strulov-Shlain (2022).

Figure 2: Shares of price endings over time

(a) Shares of products that end with 99, 90 and 00, ICC data

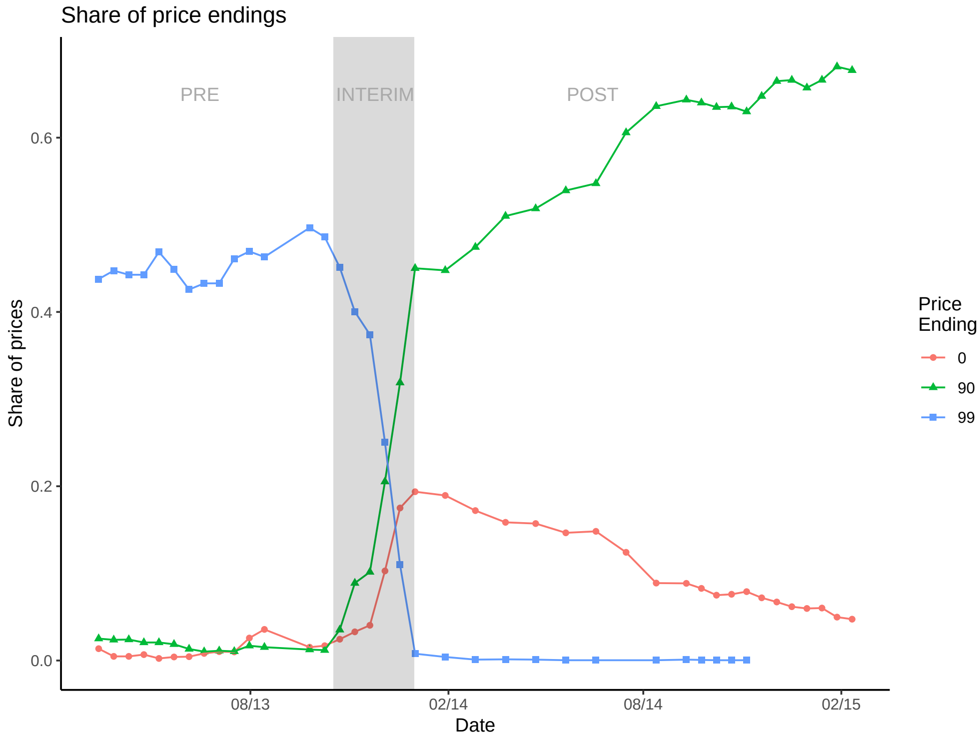

(b) Shares of products that end with 99, 90 and 00, CBS data

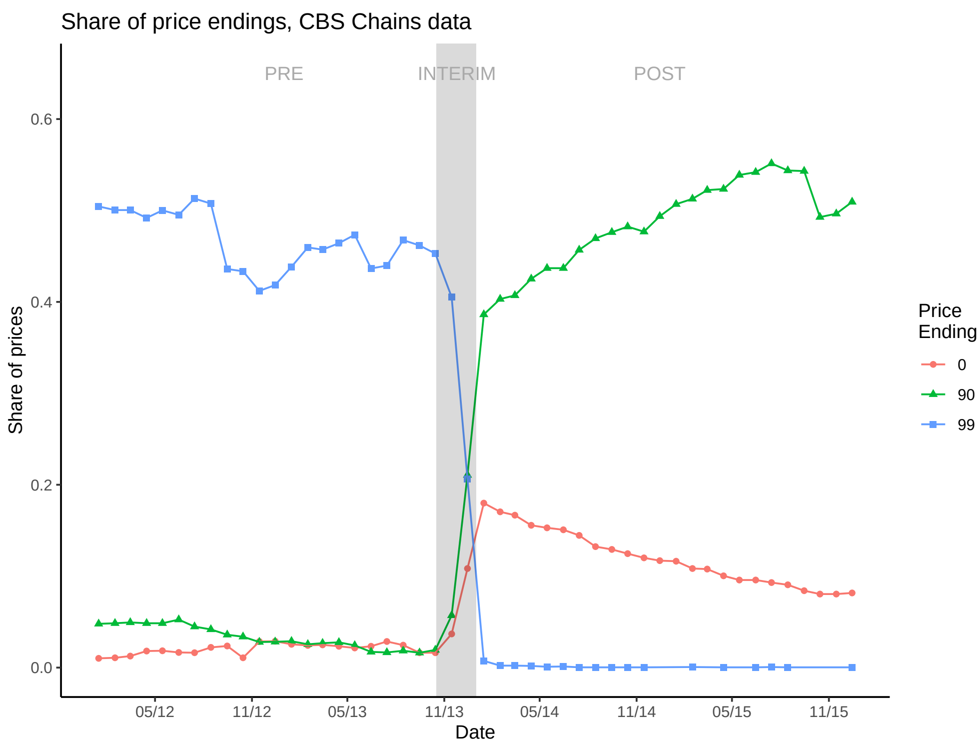

The figures above show that shares of 99-, 90-, and 00-ending prices across products and stores by sampling period. The shaded area represents the time between announcement of the reform on October 17, 2013, and its enactment on January 1, 2014. Squares represent 99-ending prices, triangles 90-ending prices and circles are 00-ending prices. The top panel shows these shares in the ICC data and the bottom panel in the CBS data.

Figure 3: Shares of price endings post-reform by chain

The figure shows that shares of 90- and 00-ending prices across products and stores by sampling period, broken down by chain. Data show shares at the post-reform period (starting January 2014). Results are shown for the 6 biggest chains in the data. Triangles represent 90-ending prices and circles 00-ending prices.

Figure 4: Price ending paths of products who ended with 99 before the policy change

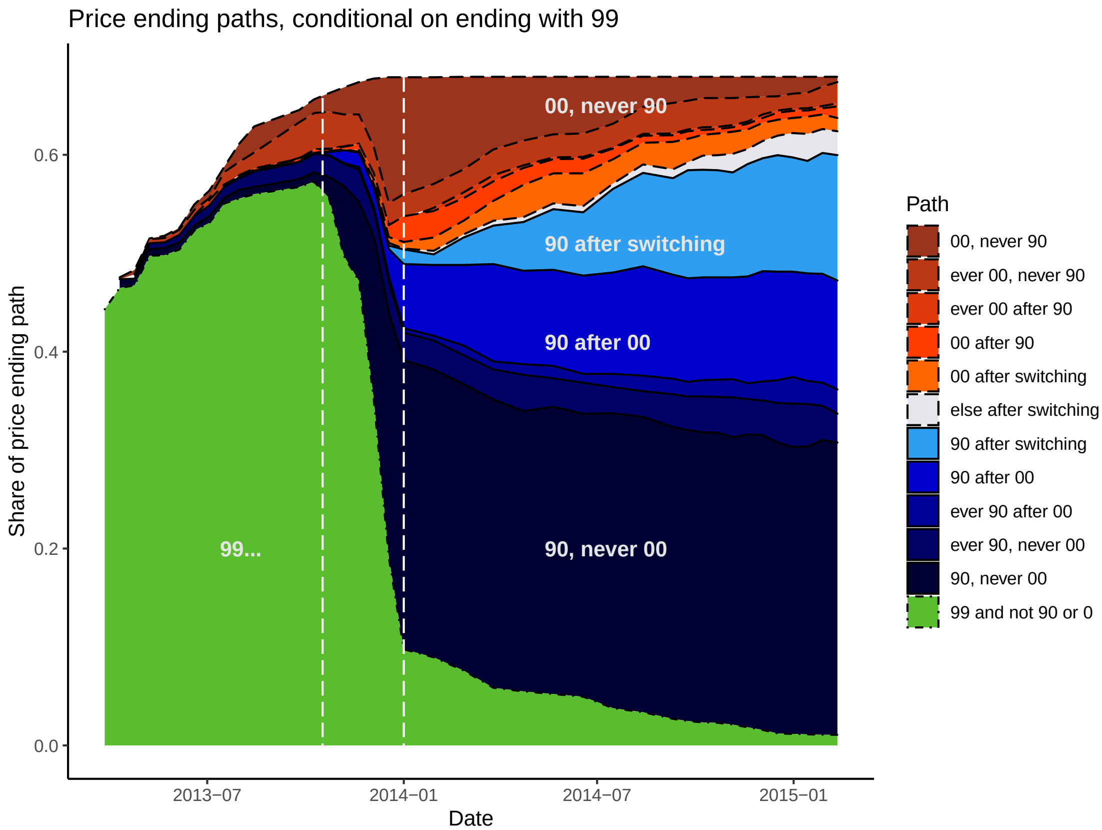

The figure shows the shares of each price-ending path (which are described in Figure A-2). The green area is the 99 or other price endings, the blue shades are the 90-ending prices, and the red shades are the 00-ending prices.

Online Appendix, Not For Publication

Firms have Partial Knowledge and they Partially Optimize: Evidence from a Reform

Avner Strulov-Shlain

A Data appendix

Prices, and their dynamics, are the core focus of this paper. This section elaborates on how prices are represented in the data, and the data-cleaning procedures taken to create the final samples assuring that inferences can be drawn from the data.

A.1 ICC

ICC data are collected by surveyors going to a set of predetermined stores with a product list. Each product is equivalent to a UPC. The surveyors record the shelf price of the item, and the total basket price was posted on the ICC website. The micro-data was generously provided to me by the ICC. The product list may change from week to week, but a core set of products remained constant. Surveyors went to the stores every 2-4 weeks (every 4 weeks in January-September 2014 due to budget cuts, and every 2 weeks otherwise).

On January 2015 Israel passed the "food law" mandating full price transparency starting May 2015 (Ater and Rigbi (2019)). This has made the ICC work redundant, and led to them canceling data collection.

Price imputation The ICC data are not a balanced panel, since the product list is changing over time, stores enter and exit the sample, and the sampling rate is inconsistent. Therefore I clean the data to be more balanced. First, I restrict the sample to products that are sampled at all periods. Second, I restrict it to stores that are sampled at least T times within each policy-relevant period. Specifically, I chose 4, 3, and 5 times for the pre-announcement (3/13-10/13), between announcement and enactment (10/13-12/13), and post enactment periods (1/14-2/15) respectively for the Main ICC data; and 15 times at the post period for the Post-only analysis, and at the last period in the data (on February 2015). Third, I balance the panel by completing missing price observations for a product in a store backwards (we also complete forward as a robustness test, but results did not change in a meaningful way). i.e., if a product's prices in 3 consecutive sampling periods are $\{3.9, Missing, 4.2\}$, I complete them to be $\{3.9, 4.2, 4.2\}$. By completing prices backwards I set an upper limit on learning rates.

The data cleaning makes the patterns cleaner and is crucial for exploiting the panel nature of the data for analysis, but does not qualitatively alter the main results. For example, Figure A-1 is equivalent to Figure 1a but is based on the raw data without any sample selection or price imputation.

Figure A-1: Price ending patterns using the full ICC sample

An important trait is that prices are recorded for the first unit purchased. That is, if there is a promotion such as "2nd item for 50% off", this is not taken into account. As such, the prices mostly represent non-sale prices.

A.2 CBS

The second database is prices collected by Israel's Central Bureau of Statistics (CBS) in order to create the Consumer Price Index. The data are similar in nature to ICC, in the sense that surveyors collect displayed prices of a list of products from stores every month. Since the goal of the CBS is to create a reflection of the representative shopping list, the CBS collects an extensive share of data from non-supermarkets (e.g., markets, specialty shops and convenience stores); each product is sampled from a somewhat different set of stores; and products are a "product-type" rather than an UPC. For example, the product "cottage cheese" may be of different size, manufacturer, and fat content between stores and within stores over time, and collected mainly from supermarkets in big cities. However, the panel is balanced, sampled monthly from 2012 to the end of 2015, and the set of product-types is an order of magnitude larger than in the ICC data (with 171 product-types in the final sample). Summary statistics can also be seen in Table 1. To make the data more comparable to the ICC data I restrict it in the same way and consider supermarkets only.

Most of the pricing pattern analysis is done using the ICC data, since it allows for analyzing the heterogeneity between chains which is of main interest. The basic analysis is conducted for both the ICC and the CBS data, showing very similar results.

A.3 StoreNext

The StoreNext data is a scanner data recording store-product-period level data on revenue and number of units sold. It also include store and chain anonymous identifiers. Prices in the data are therefore the quantity weighted average of prices paid during the period, and equal revenue over units.

Data have some clear measurement errors in it. To understand the issues, consider the following (true) series of prices in the data, of a chocolate in the same store from four consecutive days in 2015. The price is the division of revenue by units sold in that day. On Sunday the price was 10.9025 (4 units sold), on Monday exactly 10.90 (1 unit sold), and on Tuesday it was 10.902308 (13 units sold). That is, the price on Sunday seems like an average of 3 units sold for 10.90 and 1 unit for 10.91 (and on Tuesday of 10 units for 10.90 and 3 for 10.91). A same-day price difference of 1 Agora is extremely unlikely and 10.91 was not even an admissible price in 2015. Of course, when the price is so close to 10.90 it is easy to guess that 10.90 is the true price, but prices also fluctuate by a few Agorot from day to day, which still seems unlikely to be a true price update and make it harder to determine what the true price(s) might be. To get a sense of the magnitude of such noise in the data, consider that while weekly level data from a national US retailer exhibit 90% of prices that are "to-the-cent" (Strulov-Shlain (2022)), the *daily* Israeli data shows only 55% of those.

I however aggregate the data to the weekly level to merge it with another excerpt of 13 products[27]. To make the data comparable I am keeping records from 2013 to 2015.

Given the price measurement issues, I round prices to the nearest admissible price (at the 1-Agora level in 2013 and 10-Agorot level in 2014-2015)[28]. To be confident that the weekly price represents the true price that consumers paid, I keep observations where the price is indeed

[27]I am thankful for Itai Ater for making the data available.

[28]To create a balanced sample of product-stores, I keep product-stores pairs where there are any sales for at least 80% of the total weeks, and for 90% of the weeks of the pre-reform period of 2013. I then calculate for each price sequence of a product in a store the length of same-price spells, deeming the price to be the same if the differences between two consecutive weeks are at most 2-Agorot (unless the change is between 99- and 00-endings then I deem it different prices).

admissible absent rounding (34.6% of prices), or if it is also of a same-price spell of at least three weeks (35.8%), keeping a total of 47.26% of the popular product-store observations[29]. The descriptives are shown in column "StoreNext main" of Table 1. In the main sample, 72% of observations are of admissible prices and 84% are of same-price spells of at least 2 weeks

B Additional pricing patterns

B.1 Automaton Figure

Figure A-2: Automaton tracking price endings path

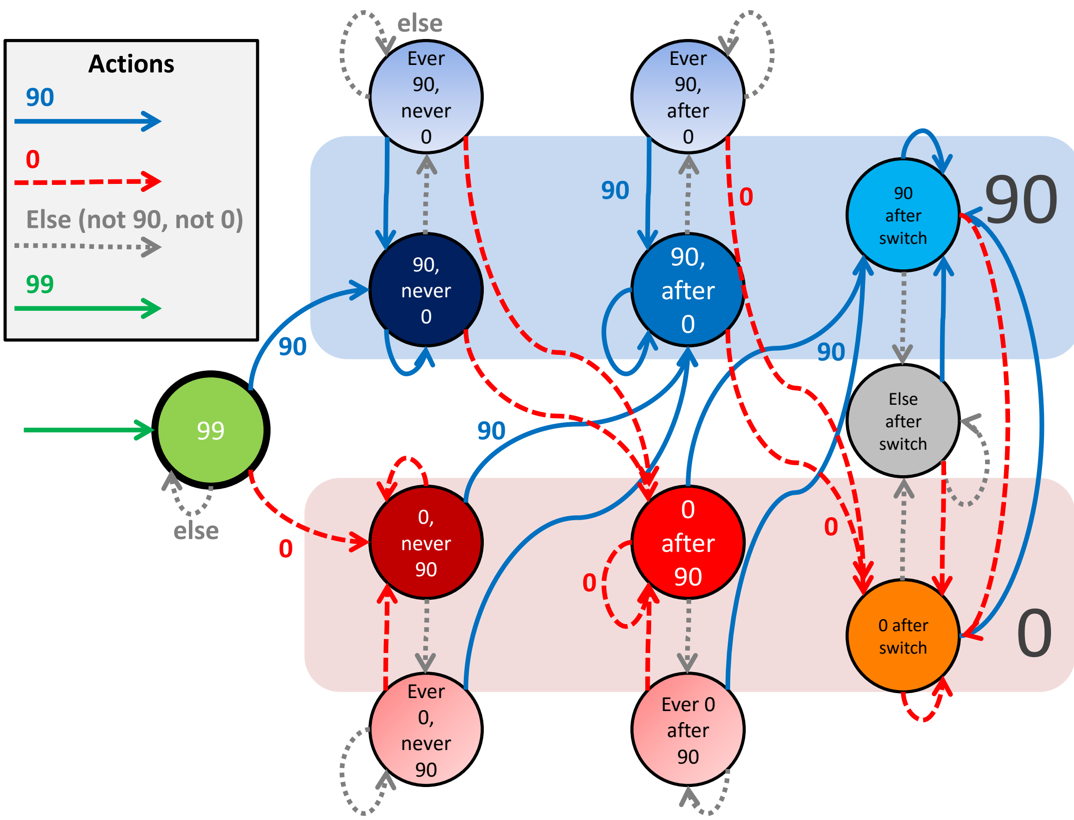

The figure illustrates a state-machine that assigns each observation into a "price-ending path" as a function of 4 possible price endings: 99, 90, 00, or other. For example, a price sequence of {4.99, 4.90, 5.00, 5.90} will have the corresponding states {99, 90 never 0, 0 after 90, 90 after switch}.

[29]For the same-price spells I code "the" price as the modal price of the spell.

B.2 Product heterogeneity

Figures 2 and 3 show price-endings shares aggregated across the different products, but we can learn more by thinking of heterogeneity. Consider the assumption that products with higher shares of 99-endings are associated with stronger left-digit bias, or at least with stronger perceived left-digit bias by the chain[30]. The correlation is stronger if different products have similar cost distribution and elasticities. To the extent that more 99-endings reflect stronger perceived left-digit bias, these should be associated with fewer 00-ending prices. Meaning, due to the true or perceived left-digit bias, if firms understand the model we would expect items that had more 99-endings to have fewer 00-endings post reform. Do these predictions hold in the data?

Figure A-3 gives an idea on how firms responded. The panels show for each product the share of 00-ending prices at different periods against the share of 99-ending for these products in the pre-period. First, Figure A-2a shows that even for products with few 99-endings (presumably due to lower left-digit bias), there was almost no 00-ending pricing taking place before the reform. Second, a year after the reform, as shown in Figure A-2c, there were more 00-ending prices for some products, and mostly for items with lower pre-reform shares of 99-endings. Finally, Figure A-2b shows the patterns immediately after the reform. The figure shows that there were many more 00-ending prices, and here too, mostly for items with lower shares of 99-endings though the negative correlation is driven by one group of products, namely fresh produce[31]. That is, for all products except for fresh produce, there is a strong positive correlation, suggesting that stores rounded 99 to 00 rather than the other way around. In Appendix Figure A-4, we see similar patterns in the CBS data, though less pronounced, where each point is a product-type rather than a specific product.

B.3 Spatial competition

Another dimension to look at heterogeneity of responses is through spatial differences. Can competitive forces increase the speed of convergence?

I examine whether stores that have more nearby competitors lower the share of 00 faster than those with fewer competitors. For each store in our data, I count the number of competing stores within a radius. For example, above and below median within a 5 km radius[32]. I then compare the differences in shares between stores above and below the median of spatial competition. Figure

[30]What does it mean to have bias heterogeneity between products? One explanation is that different clientele purchase different products, another is that other characteristics of the products receive more attention than the price (such as observable quality in fresh produce), or that items purchased in bulk or multiple units require price multiplication which puts extra weight on the left-most digits.

[31]One group of products, of fresh produce, was very high on 99-endings pre-reform and translated immediately to 90-endings post-reform. The reason for these items different pricing might rely on that they are purchased by weight (the pricing norm in Israel is to price per kg), but that is only a conjecture.

[32]Choice of other radii does not seem to matter. A choice of 5km keeps stores above and below median also within more densely populated areas, and is not a simple separation of urban versus sparse villages. I could not get too narrow due to measurement error in locations, which are based on store name and city, only sometimes with a precise address, and then extracted from Google Maps.

Figure A-3: Correlation between pre-shares of 99 and post-shares of 90 and 0. Each dot is a product in the ICC data.

(a) Correlation between share of 99 pre and 00 pre (3-10/2013)

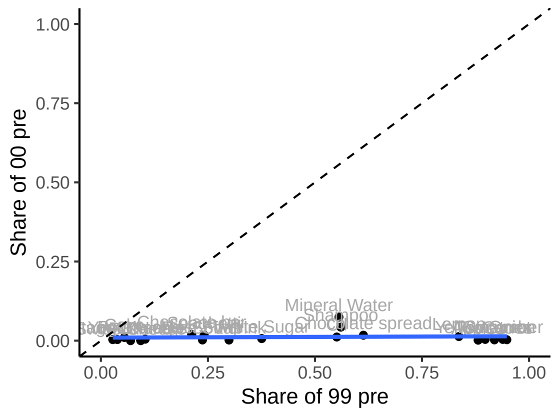

(b) Correlation between share of 99 pre and 00 immediately after policy change (1-2/2014)

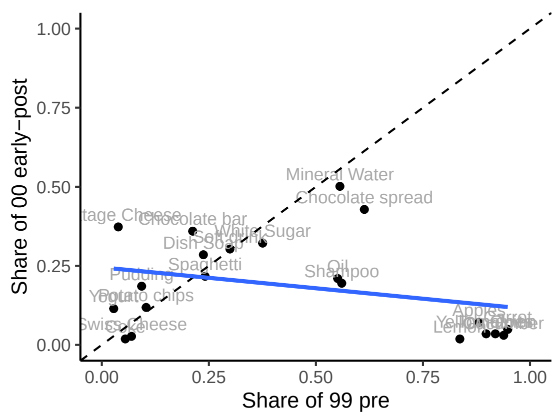

(c) Correlation between share of 99 pre and 00 a year after policy change (1-2/2015)

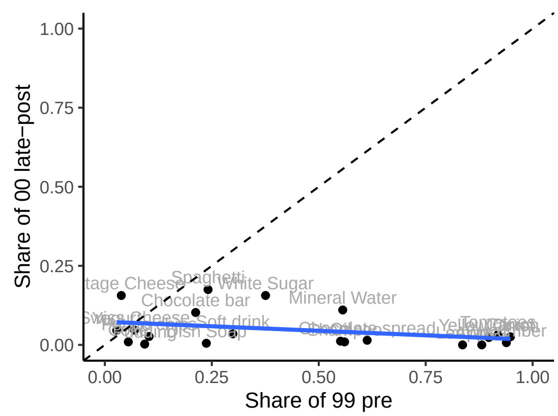

The figures show that shares of 00-ending prices at different periods against 99-ending prices at the pre-period, by product. The shares of 00-ending prices per product are calculated before the reform announcement, at the two months following the reform, and at early 2015 a year after the reform.

A-4a shows that stores in a denser area have consistently a lower share of 00, at least initially. The dashed lines represent significant differences when controlling also for product-level shares, and clustering standard errors at the store level. The differences are significant mainly at the second half of 2014, implying that maybe differential learning is in place. However, different chains are located differently between more or less dense regions, and the between-chain heterogeneity can be the source of differences. Therefore, I do the same exercise only for the 25 stores of "Shufersal Deal" (Figure A-4b), where differences in shares of 00 above and below the median[33] are less pronounced but still similar qualitatively. Most of the gap is on January 2014, and it is somewhat persistent, suggesting if anything there is a negative effect on the speed of convergence.

[33] The allocation to medians is still based on all chains

C Perceived left-digit bias algorithm

A step-by-step algorithm explains the mapping between parameters of optimal pricing and moments in the data. Let S_p be the price share of prices set at a price p resulting from one such price aggregation. I use minimum-distance estimation to match the empirical price densities, $\{S_p\}$, to the predicted price densities, $\{\hat{S}_p\}$ (where p belongs to a grand price vector $\mathcal{P}$). For example, take $\mathcal{P} = \{2.29, 2.39, ..., 3.69\}$. The algorithm works as follows: (1) Fit a logistic polynomial to the empirical price cumulative density function using prices that are not bunching at 99 or missing due to the bias, and call that *price* distribution $\hat{F}_p$. The parameters from that regression give the shape of price distribution. Then, for each pair of elasticity ϵ and left-digit bias θ: (2) take all relevant 99-ending prices, $\{q_i\}$, and calculate the matching next-lowest prices $\{P_{q_i}\}$ using Equation 3 (for example, say $\{q_i\} = \{1.99, 2.99\}$ and that $\{P_{q_i}\} = \{2.51, 3.53\}$). (3) Expand the vector of prices to include $\{q_i\}$, $\{P_{q_i}\}$, and all 9-ending prices between each pair of P_{q_i} and q_{i+1}, giving a vector of prices $\mathcal{P}'$ (e.g., $\mathcal{P}' = \{1.99, 2.51, 2.59, 2.69, ..., 2.99, 3.53, 3.59, 3.69\}$). (4) Given θ and ϵ, calculate for each $p_i \in \mathcal{P}'$ the cost $\underline{c}_{p_i}$ such that p_i is the profit-maximizing price. By construction, this is given by the first order condition as $\underline{c}_{p_i} = p_i \frac{1+\epsilon}{\epsilon} + \frac{\theta}{1-\theta}\frac{\lfloor p_i \rfloor}{\epsilon}$. (5) For each price, the share of observations at that price is $\hat{S}_{p_i} = \hat{F}_c\left(\underline{c}_{p_{i+1}}\right) - \hat{F}_c\left(\underline{c}_{p_i}\right)$, where $\hat{F}_c$ is the *cost* distribution. I recover $\hat{F}_c$ using the following identity coming from the model:

$$\begin{aligned}
\hat{F}_c(\underline{c}_p) &= Pr(c \leq \underline{c}_p) = Pr\left(p\frac{1+\epsilon}{\epsilon} + \frac{\theta}{1-\theta}\frac{\lfloor p \rfloor}{\epsilon} \leq \underline{c}_p\right) \\
&= Pr\left(p \leq \underline{c}_p\frac{\epsilon}{1+\epsilon} - \frac{\theta}{1-\theta}\frac{\lfloor p \rfloor}{1+\epsilon}\right) \\
&= \hat{F}_p\left(\underline{c}_p\frac{\epsilon}{1+\epsilon} - \frac{\theta}{1-\theta}\frac{\lfloor p \rfloor}{1+\epsilon}\right)
\end{aligned}$$

(6) Expand $\mathcal{P}'$ to $\mathcal{P}$ by adding the missing prices between q_i and P_{q_i} and assigning predicted zero shares to them.[34]

To estimate parameters, I minimize the sum of squared differences between predicted moments and actual moments.The above procedure generates the predicted price moments $\hat{S}_{p_i}$ as a function of elasticity ϵ and left-digit bias θ. Since I am using the densities of some prices to fit the shape price distribution, I exclude these moments from the minimization problem.

[34]relabel non-9-ending next lowest prices as the lower 9-ending price, without changing the share. For example, 2.51 will be relabeled as 2.49 with a share $\hat{S}_{2.49} = \hat{F}_c\left(\underline{c}_{2.59}\right) - \hat{F}_c\left(\underline{c}_{2.51}\right)$

Figure A-4: Correlation between pre-shares of 99 and post-shares of 90 and 0. Each dot is a product in the CBS data.

(a) Correlation between share of 99 pre and 00 pre (3-10/2013)

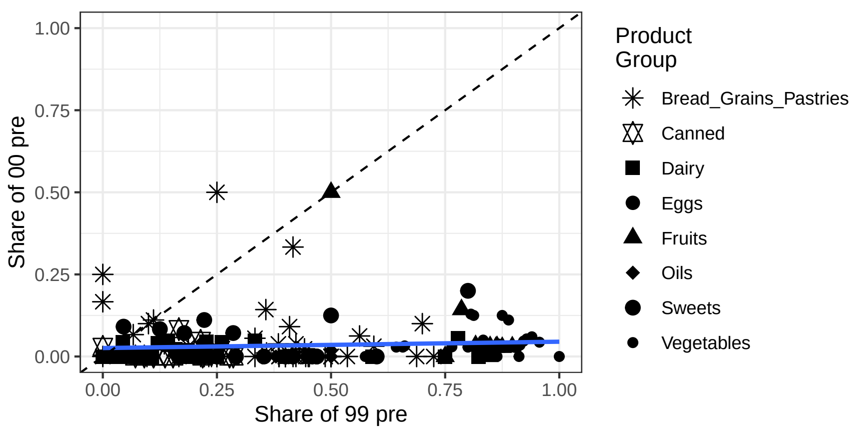

(b) Correlation between share of 99 pre and 00 immediately after policy change (1-2/2014)

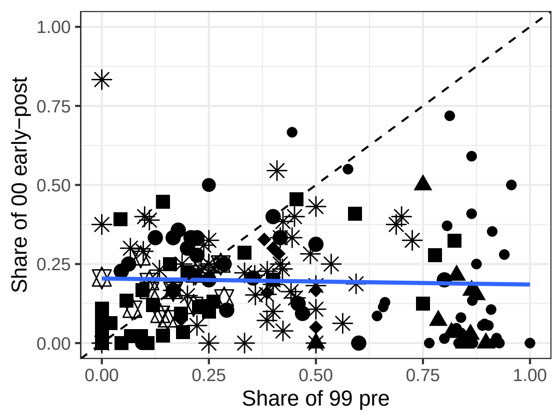

(c) Correlation between share of 99 pre and 00 a year after policy change (1-2/2015)

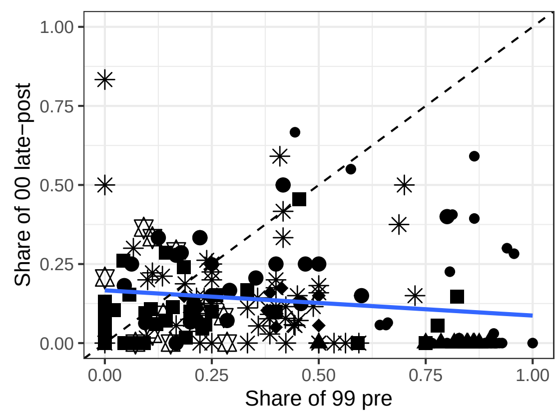

The figures show that shares of 00-ending prices at different periods against 99-ending prices at the pre-period, by product-type in the CBS data. The shares of 00-ending prices per product are calculated before the reform announcement, at the two months following the reform, and at early 2015 a year after the reform.

Figure A-5: Spatial competition and shares of price ending

(a) Shares of 00 and 90 for stores above and below median of spatial competition

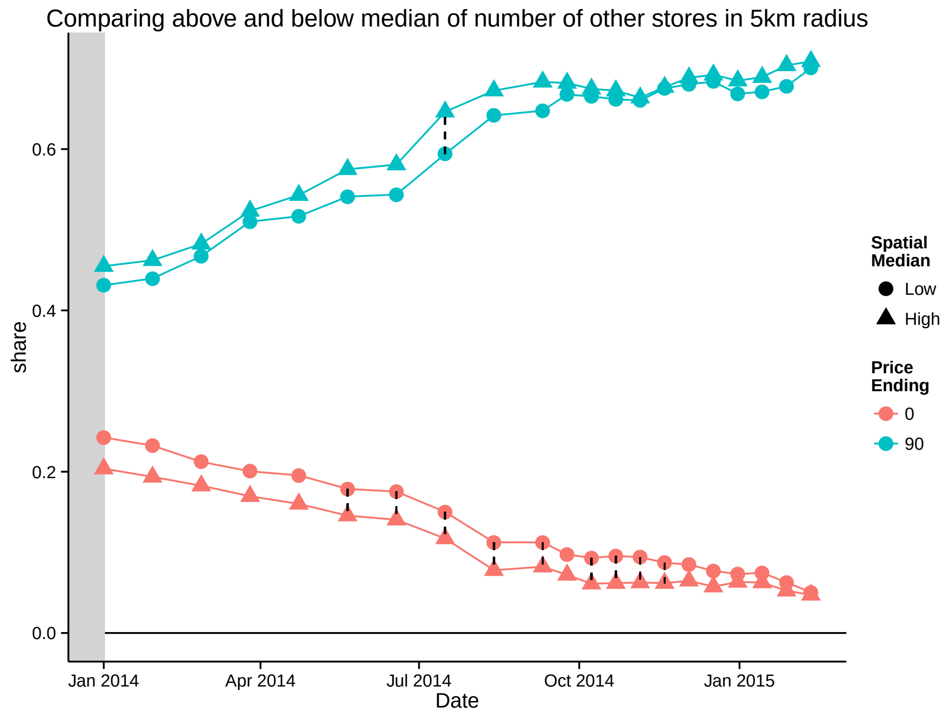

(b) Shares of 00 and 90 for "Shufersal Deal" stores above and below median of spatial competition

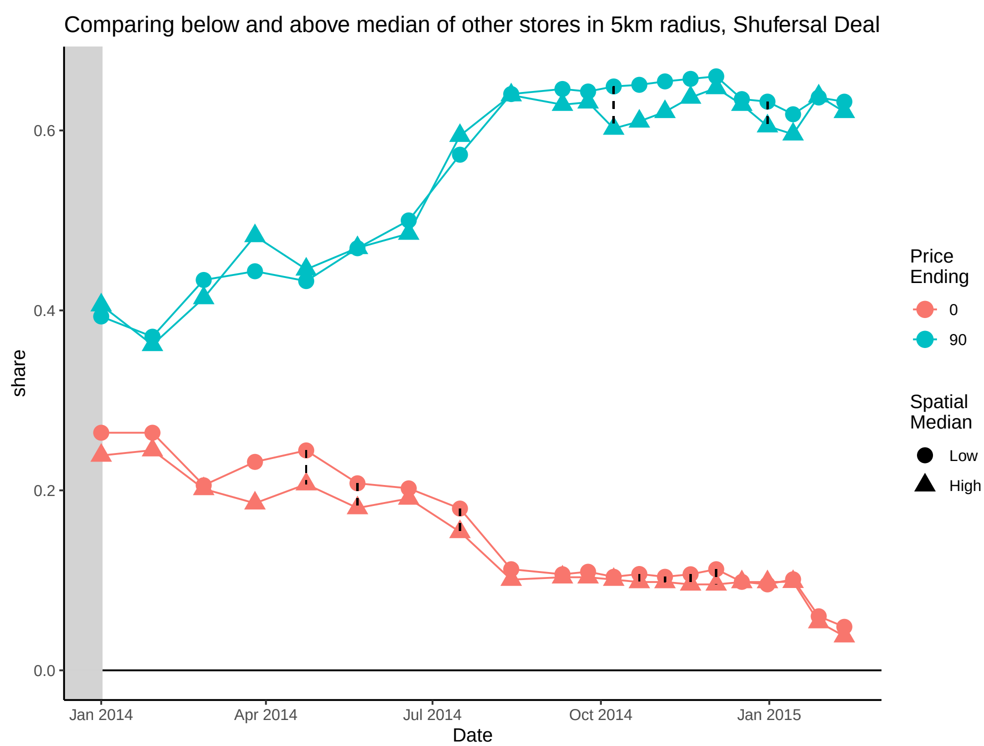

The figure shows the shares of 00- and 90-ending prices at the post-period for stores that are above and below the median of number of local competitors. The top panel show all stores, and the bottom panel only the stores of the Shfersal Deal chain.