Firms have Partial Knowledge: Evidence From A Reform

Avner Strulov-Shlain*

October 31, 2025

Abstract

How can firms persistently misoptimize, and how can they learn? A pricing reform forced firms to adjust most prices. Because consumers exhibit left-digit bias, optimal pricing implies setting prices just below round numbers and avoiding round prices. The initial response revealed that firms operated with partial knowledge, even in the long run. Partial knowledge enables mistakes to persist. At the same time, firms were able to learn, which highlights exploration as a way to improve and lack of exploration as fostering partial knowledge. This paper encourages researchers to treat firms' misoptimization as a descriptive benchmark and urges practitioners to invest in learning processes.

*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, Alex Frankel, Wes Hartmann, Johannes Hermle, Ella Honka, Ali Hortacsu, Alex Imas, Avi Goldfarb, Tesary Lin, Sanjog Misra, Sarah Moshary, Muriel Niederle, Brad Shapiro, and Richard Thaler. 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

Like other decision makers in complex environments, firms often make mistakes. At the same time, the standard assumption in economics is that firms are sophisticated and, to a first approximation, do not err. How can they be wrong, the argument goes, if there are strong incentives to optimize?

This paper addresses two related questions: How can firms be persistently wrong? And how can they improve? Firms may deviate from a normative benchmark either because they know the optimal action but face real or perceived implementation constraints and frictions, or because they simply do not know what the optimal action is[1]. To improve, firms must either relax constraints (in the former case) or acquire knowledge (in the latter).

We distinguish between implementation frictions and partial knowledge by studying empirical patterns following a unique reform, and further quantifies errors using a simple structural model. Consider the following thought experiment, which illustrates this distinction. Imagine that there are three options with an objective ranking of $a \succ b \succ c$. In steady state, the decision maker chooses a, consistent with the ranking. Then, option a is removed from the choice set. If choosing b or c is equally costly, full knowledge implies choosing the next-best option, b. Choosing c instead suggests partial knowledge, as it reveals that the decision maker did not know that $b \succ c$. If, over time, the choice shifts to b, it indicates the ability to learn.

The reform studied in this paper mirrors this thought experiment by generating strong pricing predictions in the presence of left-digit biased demand (Strulov-Shlain, 2023). Firms shifted from a (99-ending prices) to c (00-ending prices) and then from c to b (90-ending prices), revealing that partial knowledge persisted before the reform, while post-reform behavior shows evidence of learning. A structural model quantifies deviations from a normative benchmark of profit maximization, showing an initial gap from optimal pricing before the reform, followed by a spike in forgone profits, which then declines over time—eventually settling at a lower level than before the reform.

In 2013, the Israeli government banned pricing at the level of coins that no longer existed—the equivalents of the U.S. penny and nickel—forcing companies to use prices ending in .X0 (Ater and Gerlitz, 2017). For example, 2.96 was banned, while 2.90 or 3.10 were allowed. Before the reform, the modal price ending in Israeli supermarkets was 99 (e.g., 2.99), accounting for approximately 45% of prices, while 00-ending prices were almost entirely absent. This pre-reform price-ending distribution is consistent with at least partial optimization, specifically with a pricing model that accounts for left-digit biased demand[2]—a demand structure driven by consumers' distorted price

[1] At a deeper level, lack of knowledge can be driven by real or imagined frictions. However, this is beyond the scope of the paper.

[2] That supermarkets recognize they face left-digit biased demand, at least to some degree, was also documented in U.S. supermarkets (Strulov-Shlain, 2023).

perception, which leads to discontinuous drops in demand at round prices.[3] Under such demand, prices that would have ended with higher digits in the absence of bias (e.g., 5.00 or 5.29) are better set with a 99-ending (e.g., 4.99). When 99-endings—previously a profit-maximizing choice under many cost and elasticity configurations—were banned, and the highest "just-below" price-ending became 90, pricing predictions still implied that prices should not end in 00.[4] Indeed, demand analysis using scanner data shows that changing a product's price from 4.99 to 5.00 versus 4.99 to 4.90 was associated with an excess demand loss of 5%–9%. In the notation of the earlier thought experiment, 99-ending prices represent option a, 90-ending prices option b, and 00-ending prices option c, with $a \succ b \succ c$. Thus, pre-reform behavior was consistent with at least an understanding that $a \succ b$ and $a \succ c$.

At the time of the reform, 20% of prices ended in 00—a high share inconsistent with the observed demand structure, where $b \succ c$. In the years that followed, the share of 00-endings declined. For some chains, a discrete drop in 00-ending prices occurred about eight months post-reform. The evolution of price-endings shows that products whose prices shifted from 99 to 90 remained at 90, while those that shifted to 00 quickly reverted to 90. In other words, we observe movement from a to both b and c, followed by a shift from c to b. Firms eventually learned that $b \succ c$, potentially more clearly than they had known before.

A structural model, building on Strulov-Shlain (2023), takes as input monthly price distributions, estimated price elasticities, and left-digit bias to quantify deviations from a counterfactual of full-knowledge, frictionless optimization. The model shows a pattern of consistent deviations before and after the reform (similar to List et al., 2023; Strulov-Shlain, 2023) on the order of 2% of profits, with a spike in losses of 0.75pp at the reform's onset. The magnitude of left-digit bias should have resulted in a higher prevalence of 99- and 90-endings than observed. However, about a year after the reform, firms are estimated to perform better than before, with losses falling below 2%. This improvement is primarily due to greater use of 90-endings and fewer lower-ending prices overall.

This paper does *not* claim that partial learning is necessarily normatively unjustified. Although forced change made learning costs unavoidable, the benefits may have been lower than (unobserved) costs. Nonetheless, the point stands: Firms operate under partial knowledge and deviate from a frictionless benchmark.

In sum, the paper presents evidence of persistent deviations from frictionless optimization, attributes these deviations to partial knowledge rather than implementation frictions, and shows that firms can learn through exploration. As such, it encourages researchers to take partial optimization seriously as a descriptive benchmark of firm behavior. It also shows that exploration can improve outcomes. While restricting actions does not constitute a viable policy lever, encouraging greater experimentation within firms may be a more promising path.

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

[4] There is indeed a trade-off introduced by the reform: pricing below the round price threshold (e.g., 4.90 vs. 4.99) incurs a larger revenue loss, and the discontinuity becomes "smoother." Yet, left-digit bias that justified many 99-endings also supports the use of 90-ending or higher prices, but not 00-endings.

The main contribution of this paper is to provide an example of partial knowledge as a root cause of persistent mistakes among large firms. That is, firms do not precisely know what to do, rather than knowing the optimal action but being unable to execute it. This partial knowledge may result from various mechanisms—ranging from rational or irrational inattention, principal-agent problems, managerial inertia, information gaps, or skill deficits. When firms are observed to mis-optimize, explanations in the literature include frictions and heuristics (DellaVigna and Gentzkow, 2019; Hortacsu et al., 2021; Huang, 2021; Strulov-Shlain, 2023), principal-agent misalignment in organizational economics (Hortacsu et al., 2021; Shapiro et al., 2021), management practices (Anderson et al., 2021; Bloom and Van Reenen, 2007, 2010; Bloom et al., 2013; Han et al., 2025), attention costs (Goldfarb and Xiao, 2011, 2019), and incorrect beliefs (Aguirregabiria and Jeon, 2020; Hortaçsu and Puller, 2008; Strulov-Shlain, 2023; Xie, 2018). All of these channels can result in or stem from partial knowledge, making it a parsimonious and useful descriptor of firm decision-making.

This paper also contributes to the literature on the nature, potential, and limits of firm learning (see review by Aguirregabiria and Jeon (2020)). Unlike in other settings (Doraszelski et al. (2018) and Huang et al. (2018)), firms did not need to learn a new demand structure; but like these settings, learning in a new situation happened, and took a long time. Further, unlike in those studies, , learning in this context did not lead to full optimization within the observed period. In terms of limitations to learning, prior work focused mostly on mis-learning by individuals. For example, Schwartzstein (2014); Hanna et al. (2014); Gagnon-Bartsch et al. (2021), show that in multi-attribute decision environments, individuals may persistently err by ignoring relevant dimensions. Field experiments in India and Uganda (Anderson et al. (2021); Bloom et al. (2013)) demonstrate that management and marketing training can increase profits and growth, attributing prior suboptimal behavior to information gaps. In the current study, the consistent use of 99-endings prior to the reform suggests that firms did attend to price endings but did not fully understand their implications. That is, learning appears to be partial—certain attributes are attended to, but only incompletely or temporarily.

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 interprets the findings. Section 5 concludes with a discussion of the implications of these findings.

2 Setting and Data

2.1 Reform details

On October 17, 2013, the Israeli Ministry of Economics announced that, starting January 1, 2014, it would begin enforcing an existing law banning pricing in denominations of coins that

no longer exist.[5] Israel had 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 denominations different than 10 Agorot multiples—most commonly with 9- and 5-endings. Therefore, most prices had to change due to the reform.

This policy was motivated by consumer concerns that they were being misled by non-existent prices (e.g., a posted price of 4.99 that is actually rounded to 5.00; prices in Israel are VAT-inclusive). Consumers either paid the total basket price electronically or the rounded price (to the nearest 10 Agorot) when paying in cash.

2.2 Data

The effects of the reform are evident in the dynamics of price setting and their impact on demand. I use data on prices and purchases in supermarkets, drawn from three sources: two are shelf price datasets from before and after the policy change, and the third is scanner data containing revenues and quantities sold. In all datasets, the unit of observation is a product in a store during a specific period. I first describe the pricing data, then the scanner data.

In the pricing data, the key variables are the shelf price as displayed in a store, and identifiers for the product, store, and retail chain. The main pricing dataset was collected by the Israel Consumer Council (ICC), a consumer advocacy group. ICC employees visited supermarket stores every 2–4 weeks between early 2013 and early 2015 and recorded the prices of a predetermined list of products. A "product" in this context corresponds to a UPC, uniquely defined by producer, weight, and size. The data were collected to help inform consumers which nearby store offered the lowest-price bundle in a given week. Consequently, stores and chains are identified by name and address. A second pricing dataset, from the Central Bureau of Statistics (CBS), is similar but defines "products" at the product-type level—for example, a carton of omega-3-enriched brown medium eggs (any brand). This dataset forms a balanced panel, but store identities are anonymized and chain identifiers are absent. Prices are sampled monthly. More details are provided in Appendix Section A.

Due to missing observations and changes in the sampled goods, the pricing datasets are not balanced panels. I therefore focus on a consistent subset of products and impute missing prices as described in Appendix Section A. For the main analysis, I require that a product's price in a store be observed at least once during each reform period.

Table 1 shows the number of products, stores, and observations across samples. The main ICC sample includes 21 products across 143 stores, while the CBS sample includes 171 product types in 99 stores. The share of 99-ending prices in the pre-period ranges between 45% and 47%.

[5] With the exception of "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

The third dataset is a mix of daily and weekly scanner data, with the main variables being units sold and total revenue per UPC in a store. As in the ICC dataset, each store and chain has a unique (though anonymous) identifier. These data were provided by StoreNext Ltd., a market research firm that collects transaction-level data directly from store registers and aggregates it. The unit of observation is a product (UPC) in a store-week.[7] From an initial sample of 10 UPCs from 295 supermarket stores over three years (2013–2015), six products are usable for analysis.[8] I supplement this with another StoreNext dataset containing weekly data on 13 additional products from 287 national supermarket stores over the same period.[9]

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, 2023). StoreNext is the only research-accessible scanner data source for Israel, but it contains notable measurement error. Appendix A discusses cleaning procedures in detail.

The main scanner dataset, summarized in the "StoreNext Main" column of Table 1, includes 185,000 observations. Relative to the price data, it shows a lower share of 99-ending prices (23% vs. 43%) and a higher share of 00-ending prices post-reform (15% vs. 11%). In some analyses, I restrict to a "long price spells" sample, where I focus on price spells lasting at least two weeks to exclude likely promotional activity.

Together, these datasets allow for analysis of product-level price evolution in supermarket stores using both cross-sectional and panel variation, and they enable the study of how price endings affect demand.

2.3 Pricing under left-digit bias

Prior to the reform, many prices ended in 99; afterward, many shifted to 90—likely in response to left-digit biased demand. This section outlines the concept of left-digit bias and its implications for pricing.

Left-digit bias refers to the tendency of consumers to over-weight the leftmost digit of a price. Under such bias, consumers perceive prices with different leftmost digits as more different than they are in absolute terms, and prices with the same left digit as less different. This leads to downward-sloping demand curves with discontinuous drops at digit thresholds. The key implication is that firms should avoid pricing just above those discontinuities.

The literature supports this mechanism, identifying left-digit bias in both supermarket pricing and other domains. Early literature on psychological pricing reached mixed conclusions about profitability (Basu, 1997; Gedenk and Sattler, 1999; Ginzberg, 1936) and suggested multiple potential

[7] Due to the mix of daily and weekly data, I aggregate everything to the weekly level.

[8] One UPC has a single price throughout the sample (canned corn), another is a product introduced only after the reform (400g hummus), and two others are cottage cheeses with unusual pricing behavior (see Hendel et al., 2017).

[9] Although the store sets likely overlap substantially, store identifiers are unique to each dataset, preventing a match across sources.

explanations (e.g., Basu, 2006; Bizer and Schindler, 2005; Anderson and Simester, 2003). Experimental and empirical studies on 9 and 99 endings also yielded mixed evidence (Anderson and Simester, 2003; Bizer and Schindler, 2005; Sokolova et al., 2020; Thomas and Morwitz, 2005). However, Strulov-Shlain (2023) provides large-scale evidence from 25 U.S. supermarket chains and thousands of products supporting significant left-digit bias. Related work has established the presence of such bias in non-price contexts as well (Lacetera et al., 2012; Busse et al., 2013; Li and Qiu, 2023), and extended the idea to broader demand environments (Hilger, 2018; List et al., 2023; Repetto and Solis, 2018).

To study pricing implications, I use the model from Strulov-Shlain (2023), where biased consumers face monopolistic firms. A price p is perceived as a weighted average of the true price and a focal reference point:

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

where $\lfloor p \rfloor$ is the floor of the price, Δ is a constant focal price ending, and θ represents the degree of left-digit bias. For instance, if $\theta = 0.2$, then a 1-cent change from 4.99 to 5.00 is perceived as a 20.8-cent jump, while a change from 5.00 to 5.01 is perceived as just 0.8 cents.

Demand is a function of perceived prices, while consumers pay the true price when purchasing an item. Assuming constant elasticity demand, the demand function becomes:

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

Gross profits are then:

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

assuming marginal cost c. As shown in Strulov-Shlain (2023), this model predicts bunching of prices just below round numbers—i.e., prices ending in 99 or 90—depending on available denominations. The range of missing prices (e.g., 00 to 28-endings) is determined by model parameters.

Given the pricing function, we can focus on a key outcome, the *Next–Lowest Price*, that 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 satisfies:

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

Proposition 2 in Strulov-Shlain (2023) shows that for $\epsilon < -1$, higher θ leads to higher P. Therefore, since for $\theta = 0$ demand is smooth and $P = \lfloor P \rfloor$, there exists a threshold $\theta_0(\epsilon)$ such that for $\theta > \theta_0(\epsilon)$, $P > \lfloor P \rfloor$, implying 00-ending prices should not be observed. 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. Calibrations show 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[10]). 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.

After the reform, the "just-below" price changed from 99 to 90, and the next higher price from 00 to 10. These changes make just-below pricing less profitable and upward adjustments necessarily larger. Re-solving for θ_0^{Post} under these conditions (requiring $P > \lfloor P \rfloor + 0.1$) raises the threshold, but it remains low—e.g., 0.02 for $\epsilon = -4$ and $P = 2$—still below observed levels.

Figure 1 shows the $\theta_0(\epsilon)$ values before and after the reform, alongside empirical estimates from Strulov-Shlain (2023). These findings suggest that while firms may underestimate the bias, perceived bias levels and elasticities remain above the threshold at which 00-ending prices are dominated.

Therefore, theory and evidence suggest that with realistic levels of left-digit bias and elasticity, we should observe bunching at 99- or 90-ending prices and little to no 00-ending pricing—unless firms are do not realize that they face left-digit biased demand.

3 Results: Price Patterns and Demand Response

This section presents the effects of the reform on pricing and demand, without interpretation (deferred to Section 4).[11] First, I describe aggregate and chain-level shares of price endings before, during, and after the reform. Second, I analyze price-ending paths at the product-store level.

[10]Proposition 2 in Strulov-Shlain (2023) 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 $P - \lfloor P \rfloor$.

[11]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 its failure. The actual patterns in the data can further reject several descriptions of firms' knowledge and decision processes, to be discussed in Section 4.

Finally, I estimate the impact on demand of increasing a 99-ending price to a 00-ending versus lowering it to a 90-ending.

3.1 Price-endings shares

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

Before the reform, the share of products priced at 99-endings was high and stable, while 00- and 90-endings were rare. In both pricing datasets, the share of 99-endings before the policy announcement was about 45% and remained stable. In contrast, the shares of 00 and 90 were low, at 1.2% and 1.7%, respectively.[12]

Immediately after the policy was announced, but before it was enacted, firms began changing prices. Chains had 2.5 months to update prices, and 79% of products experienced a price change in December 2013. Chains did not wait until the policy bound; they promptly started moving away from 99-endings.

When the policy took effect in January 2014, approximately 40% of prices ended in 90 and 20% in 00. Soon thereafter, the share of 00-endings began a gradual and persistent decline.

A year after the reform, the 00-ending share had fallen to about 5% in ICC data and 8% in CBS data; the share of 90-endings was higher than the pre-reform share of 99-endings at above 65%. A subsample of StoreNext data extends to 2019 and shows the 00-ending share declining to roughly 3% by 2018–2019 (from about 20% in 2014–2015 in the same sample).

Between-chains heterogeneity While the aggregate decline in 00-endings is gradual, some chains exhibit stepwise adjustments. Figure 3 shows price-ending shares after the policy change for the six largest chains. For most chains, the 00-ending share drops in a stepwise fashion; for "Mega in the City" and "Rami Levy" (the second- and third-largest chains), the shares fall to nearly zero. "Shufersal" (the largest chain), with its subsidiaries "Deal" (discount stores) and "Shelli" (urban stores), reduces 00-endings to about 10% by mid-2014. "Wine Pavilion" and "Shufersal Deal" also appear to reduce 00-endings by 2015, though the data end too early to confirm this conclusively.

In summary, the share of 00-endings was low pre-reform, rose around enactment, and then declined steadily. In contrast, 99-endings—the pre-reform just-below prices—were quickly and persistently replaced by 90-endings after the reform. A natural question is if some products were priced at 90-endings and others at 00, or were many products priced at both price endings?

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

3.2 Price-endings paths

I now move from aggregate patterns to price-ending paths at the product-store level, which can be informative about firms' learning processes. Do the patterns resemble planned experimentation or adaptive learning?

Consider a product in a store initially priced at a 99-ending. For tractability, I classify price endings into three states: 90, 00, or Other. If a price p_{ist} of product i in store s at period t ends with 90, then the current state is $pe_{ist} = 90$. The state history $h_{ist} = (pe_{is1}, pe_{is2}, \ldots, pe_{ist})$ is then categorized into path histories as described by the automaton in Figure A-2.

Figure 4 reports the shares of different price-ending paths, with green representing 99, blue shades representing 90 (or Other after 90), and red shades representing 00 (or Other after 00).[13]

Three patterns emerge. First, products that switch to 90 are unlikely to change again, and specifically unlikely to move to 00. This is evident from the large, stable dark-blue "90, never 00" segment versus the small orange "00 after 90." Second, products priced at 00 are unlikely to remain at 00 and are very likely to move to 90: the shrinking dark-red "00, never 90" contrasts with the growing blue "90 after 00." Indeed, products that initially changed to 00 are 15 times more likely to end at 90 than at 00 in the last period of the data.[14] Third, a small set of products exhibit both transitions (00 after 90 and 90 after 00) and tend to eventually settle at 90 (e.g., $99 \to 90 \to 00 \to 90$). This is reflected in the growing light-blue "90 after switching" relative to the small light-orange "00 after switching." In the last period, switchers are 8.4 times more likely to end at 90 than at 00.

Appendix B shows the changes in price-ending shares at the product level, and differences in price ending paths as they relate to the number of competitors.

Overall, product-stores that try a 90-ending tend to stick with it, while 00-endings are unstable. Some possible deliberate "experimentation" between 90 and 00 occurs but is relatively rare and typically resolves quickly in favor of 90.

3.3 Demand response

The prevalence of 00-ending prices early after the reform raises the question of whether these prices reduced revenues. The model predicts that 00-endings are detrimental to demand and should be avoided, and the observed price updating behavior, shying away from 00-endings over time, is consistent with this prediction.

The ideal design to identify the causal effect of price endings on demand would randomize price endings. In supermarkets, this would require randomizing a product's price within a store

[13] Green identifies products that ended with 99. Almost all of these eventually changed to either 90 or 00. Rarely does a product move from a 99-ending to something other than 90 or 00—in the figure, if it changes to Other it remains green in the post period.

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

across periods.[15] An alternative design randomizes comparable products' prices across stores (e.g., Chetty et al. (2009)). The reform creates a quasi-experimental version of both. Consider the same product priced at 5.99 several times in 2013 in two stores of the same chain. For the within-store design, imagine the product alternating between 5.90 and 6.00 after the reform. For the across-store design, imagine one store sets 6.00 and another sets 5.90. The analysis below leverages such variation.[16]

Using StoreNext scanner data, I estimate the following specification with instrumental variables for price:

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

where i denotes product, s store, t week, and c chain. Q is 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 α's are price-ending fixed effects: $\kappa^{99}_{ist} = 1$ if, in week t, the price equals p^{Modal99}_{is} (0 otherwise); in 2014, $\kappa^{90}_{ist} = 1$ if the price is 10 Agorot lower (ends with 90) than p^{Modal99}_{is} and $\kappa^{00}_{ist} = 1$ if it is 1 Agora higher (ends with 00) (0 otherwise). X includes controls: product-by-store fixed effects (baseline demand), an on-sale proxy (transitory sales),[17] same-price spell-length decile-by-product fixed effects (further transitory dynamics), month-by-product fixed effects (seasonality), and year-by-product fixed effects.

I instrument for the product-chain elasticity term using the leave-out average price of the product across stores in the same chain (Hausman-style IV; Hausman (1996)). Prices in other stores of the chain proxy for cost shocks but are plausibly orthogonal to local demand shocks.

In this setup, the α coefficients capture demand differences attributable to price endings, net of the overall price elasticity. The key comparison is α^{90} versus α^{00}: the difference reflects the demand gap between a 90-ending and a 10-Agorot-higher 00-ending in 2014, when price endings were more plausibly as-if random. The comparison between a 2013 99-ending and a 2014 00-ending (a 1-Agora increase) is also informative, though it conflates time with price endings.[18]

Table 2 reports the estimates of Equation (4). The contemporaneous excess demand difference between a 90-ending and a 10-Agorot-higher 00-ending, for prices that were nearest 99-ending pre-reform, is 5%–9%. That is, *over and above* the demand response to a 10-Agorot price increase, demand is about 5%–9% lower at a 00-ending. The difference between a 2013 99-ending and a 2014 00-ending (harder to interpret) is 1% to 4%. Column (2) most closely mirrors the within-store experiment (e.g., 5.99 in 2013, then 6.00 and 5.90 in the same store): the demand difference is 8.2%

[15] Randomization at the consumer level would also be ideal but is difficult in brick-and-mortar settings. In online marketplaces, List et al. (2023) exploit an RD design based on the smooth distribution of prices around 00-endings, and Dubé and Misra (2023) randomize prices in a field experiment.

[16] Strulov-Shlain (2023) test for these demand drops using larger and higher-quality data, but without randomized variation.

[17] The dummy equals 1 if the price is more than 3% lower than any price within ±3 weeks.

[18] If many prices end with 99 in 2013, year fixed effects may align the demand slope more closely with the 2013 level. Dropping year fixed effects would confound price-ending effects with other intertemporal factors.

($p = 0.025$). Columns (3)–(4) align with the across-store design: the difference is 5.8% ($p = 0.061$) in column (3), and 8.6% ($p = 0.0045$) in column (4).[19]

With an average elasticity of about −2.5, an average price of 11, and 5%–8% demand drops, the implied left-digit bias is $\hat{\theta} \approx 0.22$–$0.36$ (since $\theta \approx \frac{\%\text{ drop}}{\text{elasticity}} \cdot \text{price}$). This estimate is plotted as "X (Israeli demand)" in Figure 1. Sayag et al. (2024) replicate the current papers' methodology with an Israeli retailer data and estimate θ between 0.11 and 0.60, depending on sample restrictions, with 0.22 as the main estimate.

The model and prior literature predict no 00-ending prices, and the demand analysis suggests that 00-endings were a losing strategy. I quantify the magnitude in the next section.

4 Interpretation of Results

4.1 Framework

We interpret the findings in Section 3 assuming that left-digit bias is stable throughout the period and that firms *try* to optimize. The premise is to account for all patterns using the simplest consistent explanation. Three main patterns emerge: (1) high shares of 99-endings before the reform and of 90-endings after it, with low shares of 00-endings both before the reform and long after; (2) high initial shares of 00-endings with negative demand impacts; and (3) shares of 00-endings decrease consistently over time (and abruptly in some chains), and once a product-store is priced at a 90-ending, it is unlikely to be priced at a 00-ending later.

First, low shares of 00-endings and high shares of 99/90-endings before and long after the reform imply some acknowledgment of left-digit bias. Bunching at just-below prices and avoiding low-ending prices are qualitatively consistent with optimal pricing under left-digit biased demand. In the long run, firms recognize that just-below prices dominate 00-ending prices, even if they do not perceive the full magnitude.

Second, the initially high share of 00-endings and their negative demand consequences imply ignorance about the relative value of 90 versus 00. Firms behaved as if they understood the advantage of 99 over 00 before the reform, but that knowledge did *not* extend to predicting the relative value of 90 versus 00 after the reform. Thus, although pre-reform behavior seems consistent with an as-if optimal pricing model under left-digit bias, it fails to predict post-reform behavior.

Third, the downward trend in 00-endings, chains' abrupt changes, and the product-store price-ending paths are consistent with learning. Over time, chains appear to realize that 00-endings are worse than 90-endings. Abrupt shifts are consistent with "light-bulb" moments of realization and can rule out slow adjustment driven solely by price-changing frictions. Moreover, the observed

[19] Column (4) restricts to price spells lasting at least two consecutive weeks to reduce price measurement error.

paths—where 90-endings are effectively absorbing while 00-endings are not, and where 00 rarely follows 90 but the converse is common—align with model-free reinforcement learning (elaborated below).

Taken together, these patterns cannot be explained by profit maximization under a stable model of left-digit biased demand. Two consistent explanations stem from partial knowledge of demand. One is maximization with a wrong model that is insufficiently explored to be refuted. The other is model-free, under-exploratory learning that selects the currently believed best action without sufficiently updating beliefs about alternatives.

4.1.1 Model-based and model-free learning

Model-based decision making is the standard tool in economics. It has two components: a function mapping parameters and actions to outcomes, and a maximization operator that selects the best action subject to constraints. To identify the best action, the decision maker must learn the parameters of the target function. Model uncertainty is rarely considered, so learning typically means estimating parameters within a known model. For example, a monopolist must know price elasticities and marginal costs (and potentially the extent of left-digit bias) to set prices that maximize profits.[20] In this framework, learning updates the expected value of *every* possible action.

By contrast, model-free decision making—a private case of reinforcement learning (Sutton and Barto, 2018)—does not posit a structural model of the environment. The environment consists of states and actions, and the agent maintains: a reward signal (feedback from actions), a value function (the believed return to actions given states), and a policy (which action to take given beliefs and states). As in model-based learning, beliefs about action values are updated using feedback, and the policy selects the highest-valued action. However, resolving uncertainty means learning primarily about the value of the *actions taken*. Without a model, beliefs about untried actions are not updated by inference. For instance, if a firm expects \$10 in revenue at a \$1 price and \$12 at \$2, sets \$2, and earns \$8, it updates the value of \$2 but, absent a model, does not revise the value of \$1. Inference across actions requires a model that supports interpolation or extrapolation.

Rejecting a model rather than merely updating parameters is challenging both conceptually and practically; fully updating beliefs requires sufficient exploration.[21] Thus, whether firms are model-based with the wrong model or model-free with insufficient exploration, they can persistently make suboptimal decisions.

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

[21] In theory, both model-free and model-based procedures can converge to full learning given sufficient exploration. For example, Q-learning converges if each action is tried in each state infinitely often and updates use sufficiently small step sizes (Dayan, 1992; Watkins and Dayan, 1992). Distinguishing model-based from model-free learning often requires bespoke experimental manipulations (Daw et al., 2011; Kurdi et al., 2019).

4.2 Can Wrong Beliefs Explain Pricing?

This section seeks a descriptive model that reconciles all patterns. If firms strive to maximize some normative Π, what is the smallest deviation Π' that captures their behavior (like Kahneman et al. (1997))?

First, suppose firms know that demand exhibits left-digit bias but misestimate its magnitude. If firms believe the bias is small (and demand very elastic), 00-endings can be optimal. However, no beliefs about left-digit bias can simultaneously generate an excess mass of 99-endings before the reform and an excess mass of 00-endings after the reform (nor contemporaneous excess masses of 90- and 00-endings). In principle, left-digit bias might rationalize 99 versus 00 but not 90 versus 00 after the reform, because the post-reform trade-off is stronger. Yet, as Figure 1 shows, beliefs needed to rationalize 00-endings must be extreme. We can estimate perceived left-digit bias before the reform using the Method of Simulated Moments as in Strulov-Shlain (2023) and ask whether those beliefs are consistent with 00-endings after the reform.

The algorithm uses the empirical price distribution as moments and predicts the distribution implied by the model. Excess mass at 99-endings and the distribution of lower endings identify parameters. Appendix Section C details the algorithm.[22]

Although estimated perceived left-digit bias is much smaller than the demand-estimated level, it remains high enough to rule out 00-endings. I estimate perceived bias in the range 0.02–0.04, while threshold values that could support 00-endings are under 0.01. In other words, firms price as if they believe that the demand drop from a 1-cent increase from 4.99 to 5.00 is equivalent to a 3–5 cent increase, whereas a 2-cent-equivalent belief is sufficient to avoid 00-endings. Importantly, the demand-side estimates are 0.22–0.36 (i.e., a 1-cent increase from 4.99 to 5.00 is akin to a 23–37 cent increase). Underestimation of bias is common in U.S. markets as well (Strulov-Shlain, 2023), and is not surprising on its own. However, estimated beliefs before the reform are inconsistent with pricing behavior after the reform under the null of knowing the shape of demand.[23]

Finally, an *excess of both* 00- and 90-endings is inconsistent with any belief of left-digit bias. Excess 00-endings require beliefs implying an advantage to crossing the digit threshold (a "reverse" left-digit bias), whereas pre-reform pricing implies positive left-digit bias. Describing firms as optimizing following a disciplined belief that demand is left-digit biased is therefore a poor description of the observed behavior.

[22] Appendix Section C follows Strulov-Shlain (2023).

[23] The post-reform fact that 00 rarely follows 90 at the product-store level, while not conclusive on its own, is harder to square with model-based learning about a single parameter. If firms were estimating the left-digit bias parameter, they would need to randomize price endings systematically; yet most products remain at 90 even without direct comparison to 00. Cross-store experimentation could, however, still occur.

4.3 Forgone profits

I next quantify how far firms are from optimization by comparing observed pricing to counterfactual optimal pricing under the estimated true demand shape. Doing so requires two components: how firms price in practice and how should they price. Unlike flagging some prices as a mistake, which only require minimal and reasonable discontinuities in a downward sloping demand, this exercise requires more assumptions and a full model. For demand structure, I use a two-product logit with left-digit bias. I simulate costs and preferences (product utility intercepts, price sensitivity, and the left-digit bias parameter) and numerically solve for optimal prices for each draw.[24]

I then specify a descriptive pricing rule extending the "rule-of-thumb" model in Strulov-Shlain (2023). First, given marginal cost and elasticity, the firm computes a naive optimal price ignoring left-digit bias (via Lerner's rule). Second, if the naive price is "close enough" to a round number, it rounds to the nearest 99-ending with some probability. I extend beyond Strulov-Shlain (2023) to allow rounding to 00 with some probability. This heuristic can generate excess mass at both 99/90 and 00—something the optimizing left-digit-bias model cannot produce regardless of the bias level. The elasticity, cost, and parameters ω (probability of rounding to 99, to 00, and the rounding range) define a predictive function $p^r(\varepsilon, c, \omega)$.

I fit this rule to the price distribution month-by-month. As expected, it fits better than the optimizing model with left-digit bias. This yields monthly estimates of the pricing-rule parameters, $\hat{\omega}_t$. Using the same cost and preference draws as for the optimal-pricing calculations, I find the naive-optimal prices and then apply $\hat{\omega}_t$ to weight the implied profits across three actions: keep, round to 99/90, or round to 00.

Forgone profits are given by the ratio of profits under the rule-of-thumb prices p^r_{it} to profits under the optimal prices p^*_{it}:[25]

$$X = 1 - \frac{\sum_i \Pi(p^r(\varepsilon, c_i, \hat{\omega}_t); \varepsilon, \theta, c_i)}{\sum_i \Pi(p^*_i(\varepsilon, c_i, \theta); \varepsilon, \theta, c_i)}$$

Figure 5 shows the results. Forgone profits hover around 2% before the reform, rise to about 2.75% at implementation, and then gradually decline to roughly 1.7% by 2015—indicating that firms ultimately performed better than before the reform.

[24] Intercepts: 1 and 0.5; price sensitivities: -0.5; left-digit bias: 0.22. I draw 1,000 costs for product 1 from a Weibull distribution (shape 4, scale 6 - calibrated that the average price elasticity will match the estimated one), resample 50 costs for product 2, and generate 1,000 parameter–cost sets. I solve for the optimal price of each product once with to-the-cent pricing (pre-reform) and once with to-the-dime pricing (post-reform).

[25] Profits are given by

$$\Pi(p; \varepsilon, \theta, c) = D(p; \varepsilon, \theta) \cdot (p - c).$$

Two points merit emphasis. First, because forced exploration induced learning, firms both *could* learn and *had* something to learn—underscoring deviations from full knowledge pre-reform. Second, forced exploration—while not a typical policy lever—can facilitate learning. Nevertheless, by the end of the sample, meaningful deviations from the assumed optimum remain.

These magnitudes are sizable, though smaller than in Strulov-Shlain (2023) (likely due to using a two-product logit rather than a single-product monopoly). I view them as ballpark estimates rather than precise measurements. The key message is that the core phenomenon of distorted responses to discontinuities in perceived demand has a meaningful impact on profits.

In summary, pricing patterns following the reform indicate that firms did not optimize against the true demand model. They did not hold a stable or coherent belief about left-digit bias. However, a heuristic description—sometimes rounding to 99/90 and sometimes to 00—captures behavior well and allows quantification of forgone profits over time. Even in the best periods, profits fall short of the optimum; the reform initially increased losses but also spurred learning, leading to improved performance within about a year.

Overall, firms appear to price suboptimally due to incomplete or incoherent models and heuristic rules. Environmental shifts can amplify losses but can also jolt firms into exploration and learning.

5 Discussion

I find that, following a pricing reform, supermarket chains acted in ways inconsistent with full knowledge, which led to clear mistakes. The body of evidence on the dynamics of price setting before and after the reform, together with priors and direct evidence on demand response, suggests that firms faced left-digit biased demand yet behaved inconsistently and did not fully account for it. This paper joins other work documenting suboptimal behavior and argues that sustained partial knowledge can generate such gaps.

How do firms preserve partial knowledge? One mechanism is insufficient learning and exploration. For example, one possibility is that firms were maximizing against an incorrect model of demand that is non-monotonic at 99-ending prices—that is, the demand is generally downward sloping but exhibits spikes at 99-endings. Although this structure is refuted by the current data (as well as by List et al. (2023); Strulov-Shlain (2023)), it was suggested in some cases (e.g., by Stiving (2000) and Anderson and Simester (2003)). Alternatively, firms may be model-free learners who were able to discern the relative benefits of 99-endings versus 00-endings but not those of 90-endings versus 00-endings, only learning about the latter after the reform. Either explanation hinges on under-exploration. Had firms attempted to refute the demand-with-spikes model, they would have; had they experimented with more price endings, they would have learned the full structure of demand.

Exploration may be limited by constraints and opportunity costs (the exploration–exploitation trade-off), but that alone should not produce large, persistent mistakes. Other forms of under-exploration can lead to arbitrarily large, sustained mistakes, as in Hanna et al. (2014); Gagnon-Bartsch et al. (2021); Schwartzstein (2014), when decision makers deem some margins irrelevant and therefore never test them. In this setting, firms exhibit some awareness of the influence of price endings on demand, yet do not attempt to learn the precise form of that influence.

Critiques Several limitations apply. First, this is a single case. One objection is that price endings are too minor to matter, making the case unrepresentative of firms general knowledge levels and the counterfactual profit losses overstated. This is difficult to adjudicate without an external validation; quantification necessarily relies on structural assumptions. Indeed, the view that left-digit bias is unimportant is common, including among retailers. Interviews with retail executives in the U.S. and Israel suggest that many regard 9-endings as convention rather than deliberate optimization. This view is misguided: under-appreciation of left-digit bias can have large profit effects (List et al., 2023; Strulov-Shlain, 2023), and retailers themselves noticed that something went awry in early 2014. Consider the CEO of Mega, Israel's second-largest chain at the time: *"There was a 2%–3% decrease in revenues in Jan–Feb . . . No one can explain why that happens . . . I'm in business for 25 years, and had never seen such months."* Their subsequent change in price-ending behavior, and apparent learning, also suggests it mattered. Moreover, the perception that left-digit bias is unimportant is itself consistent with persistent under-exploration. The most extreme example of perceived unimportance is in List et al. (2023), in which a (very sophisticated) company, Lyft, 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 and its competitor, Uber, had changed its pricing to be mostly 9X-ending). Here, firms are aware to some degree, but only partially.

A related objection is that this is an outlier and that firms generally optimize. However, optimality is often assumed rather than tested. A growing literature documents persistent suboptimality in firms (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, 2023). Partial learning provides a coherent answer to the common question in that literature: how can firms persistently fail to optimize? Answer: because they do not know exactly what to do, so how can they optimize.

Another limitation is that alternative normative models could be constructed to rationalize the data as having no mistakes. For example, if beliefs about left-digit bias are state-dependent, allowing positive and negative bias simultaneously, then any pricing behavior become rationalizable. If, during periods of change, firms revert to a null of "anything goes" (akin to Enke and Graeber (2019)), the observed behavior could be made consistent with such a model. But expanding a model ad hoc to match any empirical pattern robs the notion of a "model" of its empirical bite.

. Some readers may find these conclusions hard to accept, while others might respond that misoptimization is obvious but that optimization remains a useful assumption.[26] I agree with both sentiments: firms may *try* to optimize as an ideal but face substantial challenges. To preserve the useful discipline of attempted optimization while acknowledging reality, in a separate paper, I suggest testing the sensitivity of results to deviations from the researcher's preferred normative (or best descriptive) model.

If firms are partial optimizers, they may behave persistently in a steady state while lacking *ex ante* knowledge of how to respond to regulation or market shocks. Market changes—like those studied here—can jolt firms out of a steady state and into a learning mode (Gagnon-Bartsch et al., 2021), or move them closer to or further from optimality depending on whether learning occurs. Such changes are common: establishments open and close, rivals enter and exit, tastes shift, trade agreements change, natural disasters and pandemics occur, and costs fluctuate. Treating firms as all-knowing risks missing important dynamics and drawing incorrect conclusions. Researchers should allow for potentially lengthy learning periods and consider the possibility that learning ends prematurely, leaving sustained suboptimization.

In conclusion, there is ample room for future research. Empirically, more cases and tests of partial knowledge should be explored, including how market forces shape firms' suboptimal behavior and learning. Methodologically, the usefulness and limits of assuming partial knowledge merit systematic evaluation. Theoretically, incorporating partial knowledge and optimization into mechanism design, regulatory analysis, and welfare raises important questions about policy efficacy and market outcomes.

[26] In a *Freakonomics* podcast episode discussing Shapiro et al. (2021), Steve Levitt remarked: "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

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.

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.

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.

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.

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.

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.

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.

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.

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.

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.

Han, Yi, David Huffman, and Yiming Liu, "Minds, models and markets: How managerial cognition affects pricing strategies," 2025.

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.

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.

Kahneman, Wakker, and Sarin, "Back to Bentham? Explorations of experienced utility," *The quarterly journal of*, 1997.

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.

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

Repetto, Luca and Alex Solis, "The Price of Inattention: Evidence from the Swedish Housing Market," 2018.

Sayag, Doron, Avichai Snir, and Daniel Levy, "Price setting rules, rounding tax, and inattention penalty," *arXiv [econ.GN]*, November 2024.

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.

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

Strulov-Shlain, Avner, "More than a penny's worth: Left-digit bias and firm pricing," *Review of Economic Studies*, September 2023, *90* (5), 2612–2645.

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.

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.

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

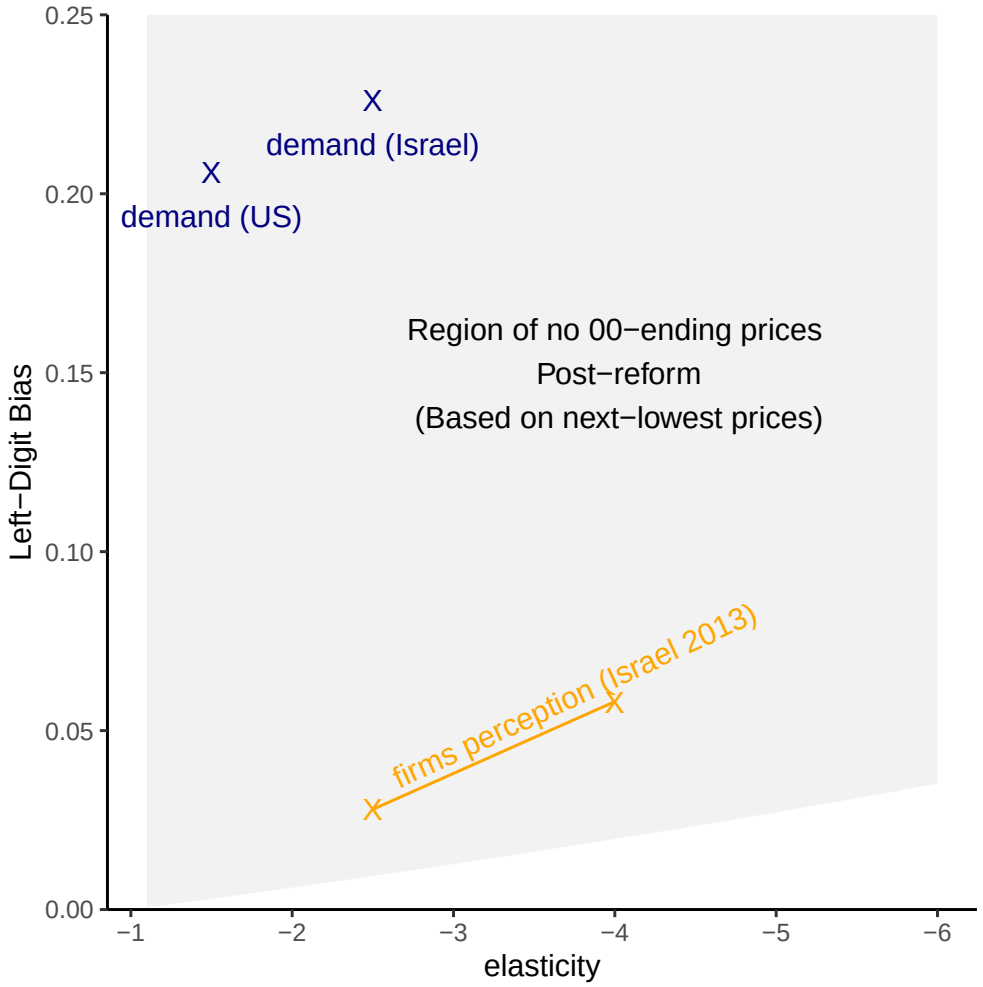
The table shows summary statistics of the three datasets used in the paper, in their raw and cleaned form. ICC, Israeli Consumer Council, collected a panel of prices at the item-store-biweek level. CBS, Israel's Central Bureau of Statistics, collected prices for the consumer price index at the "product 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.

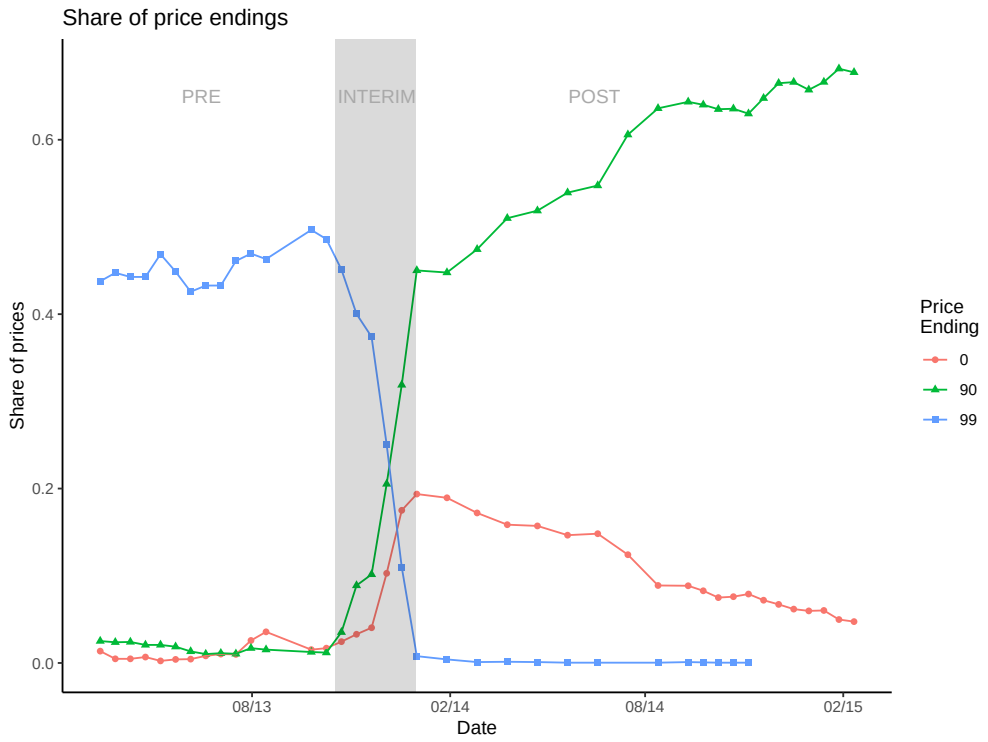
Figure 1: Minimal levels of left-digit bias θ_0 (ϵ) above which 00-ending is never optimal



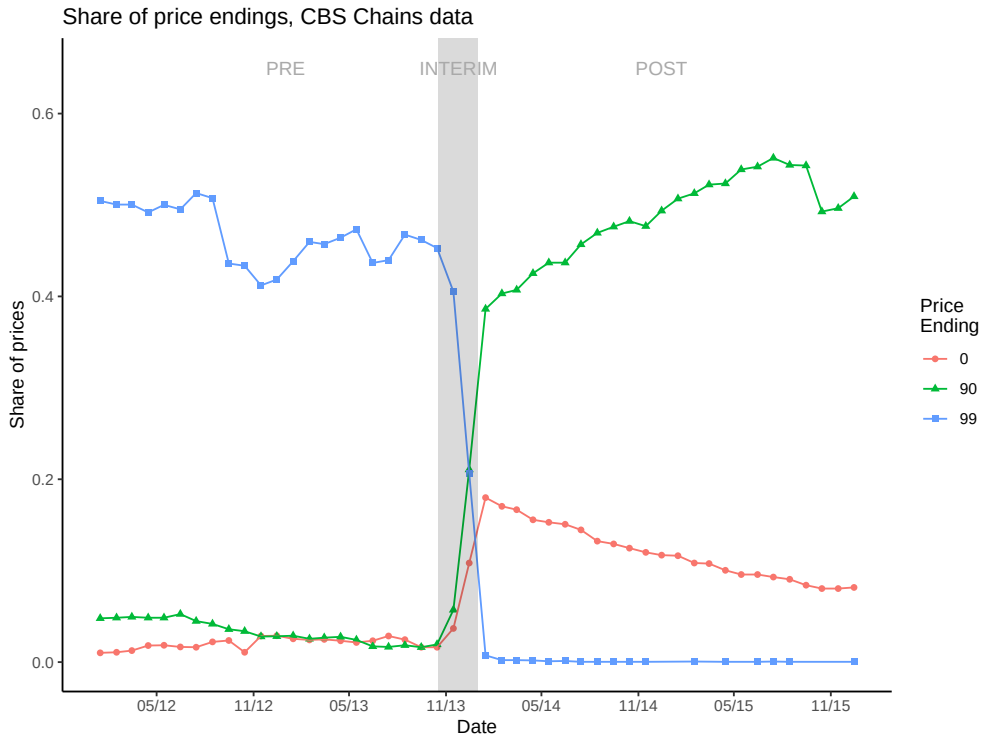
The figure shows the minimal levels of left-digit bias θ above which 00-ending prices are never optimal. The values are a function of the price elasticity and nominal price level. The dark line shows the threshold values in the pre-reform case (when 99-ending and 01-ending prices were allowed), and the 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 (2023).

Figure 2: Shares of price endings over time

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

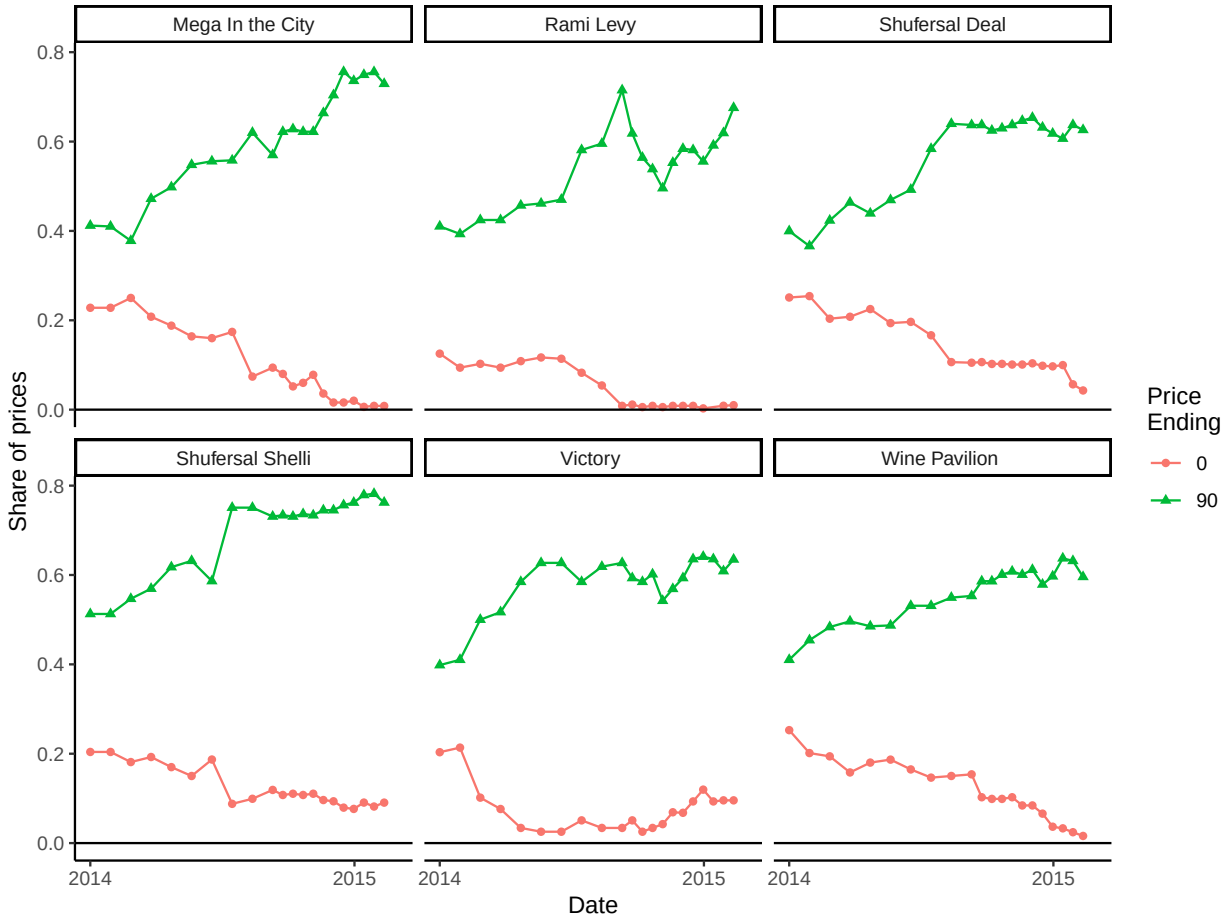


(b) Shares of products that end with 99, 90 and 00, CBS data



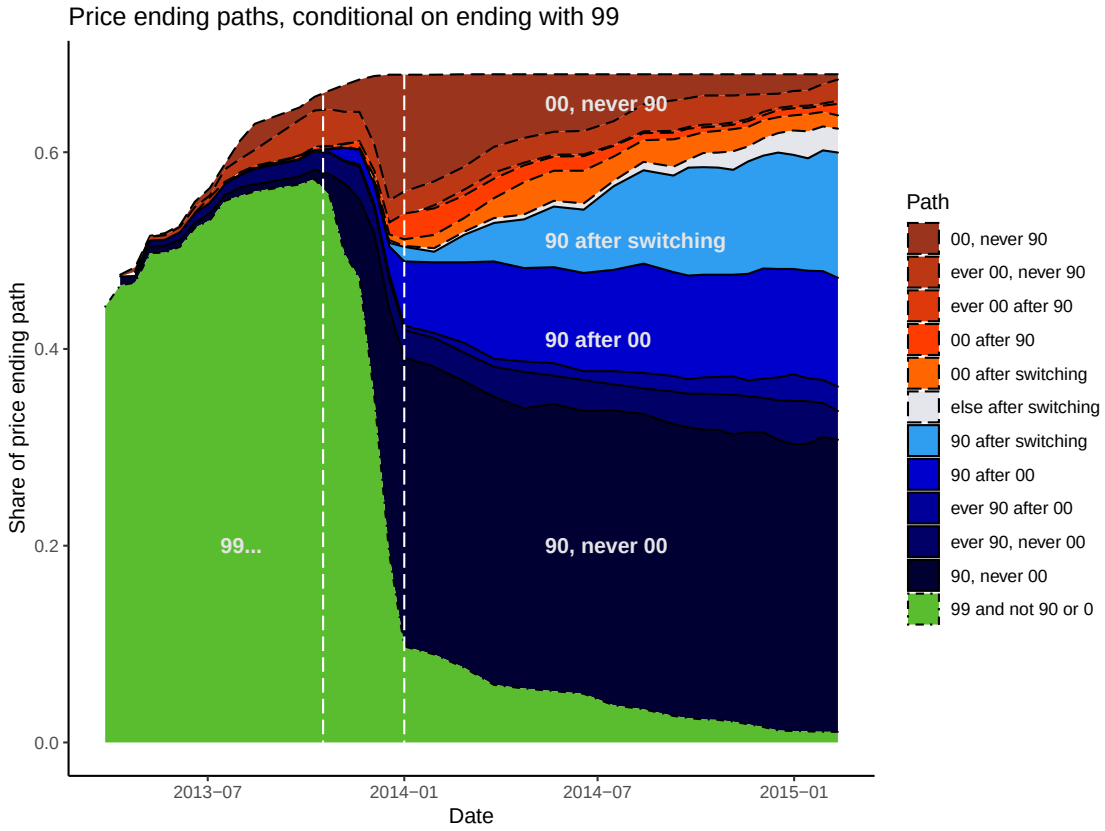
The figures above show that shares of 99-, 90-, and 00-ending prices across products and stores by sampling period. The shaded area represents the time between announcement of the reform on October 17, 2013, and its enactment on January 1, 2014. Squares represent 99-ending prices, triangles 90-ending 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 that 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.

Figure 5: Estimated counterfactual forgone profits month by month

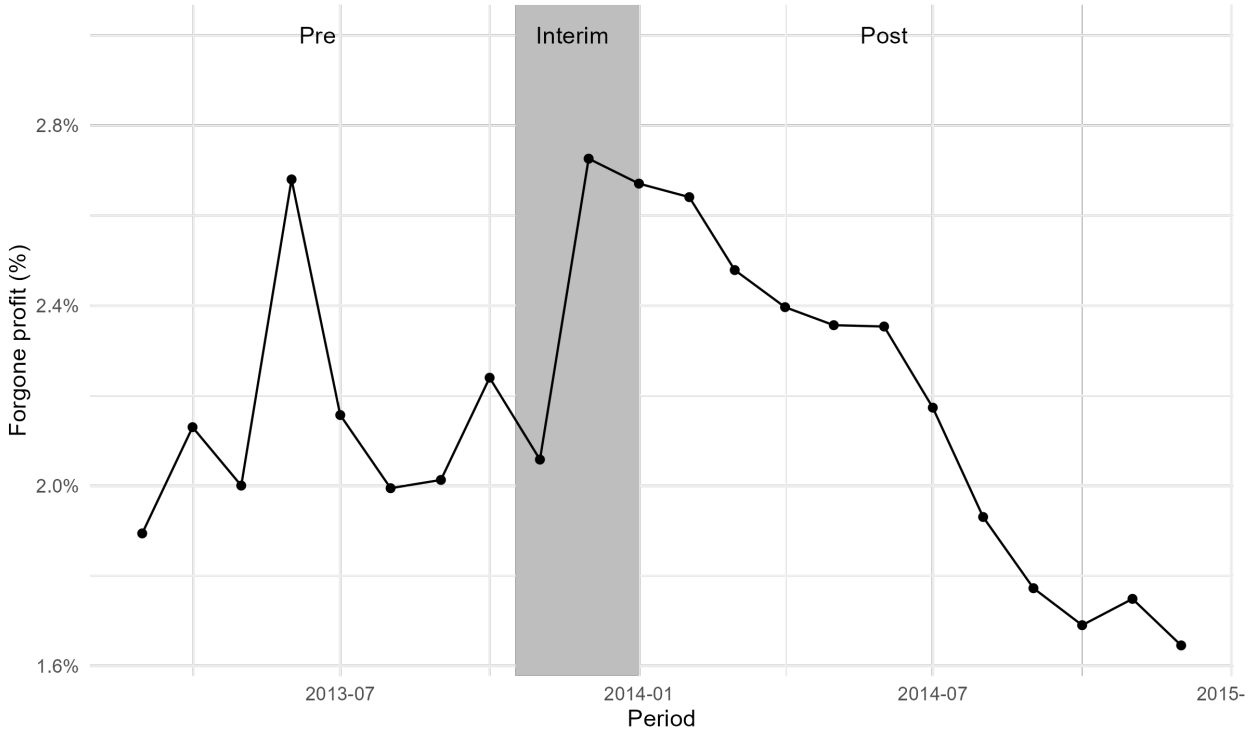


Figure shows counterfactual profits under a monthly estimated predictive rule-of-thumb pricing relative to full optimization.

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

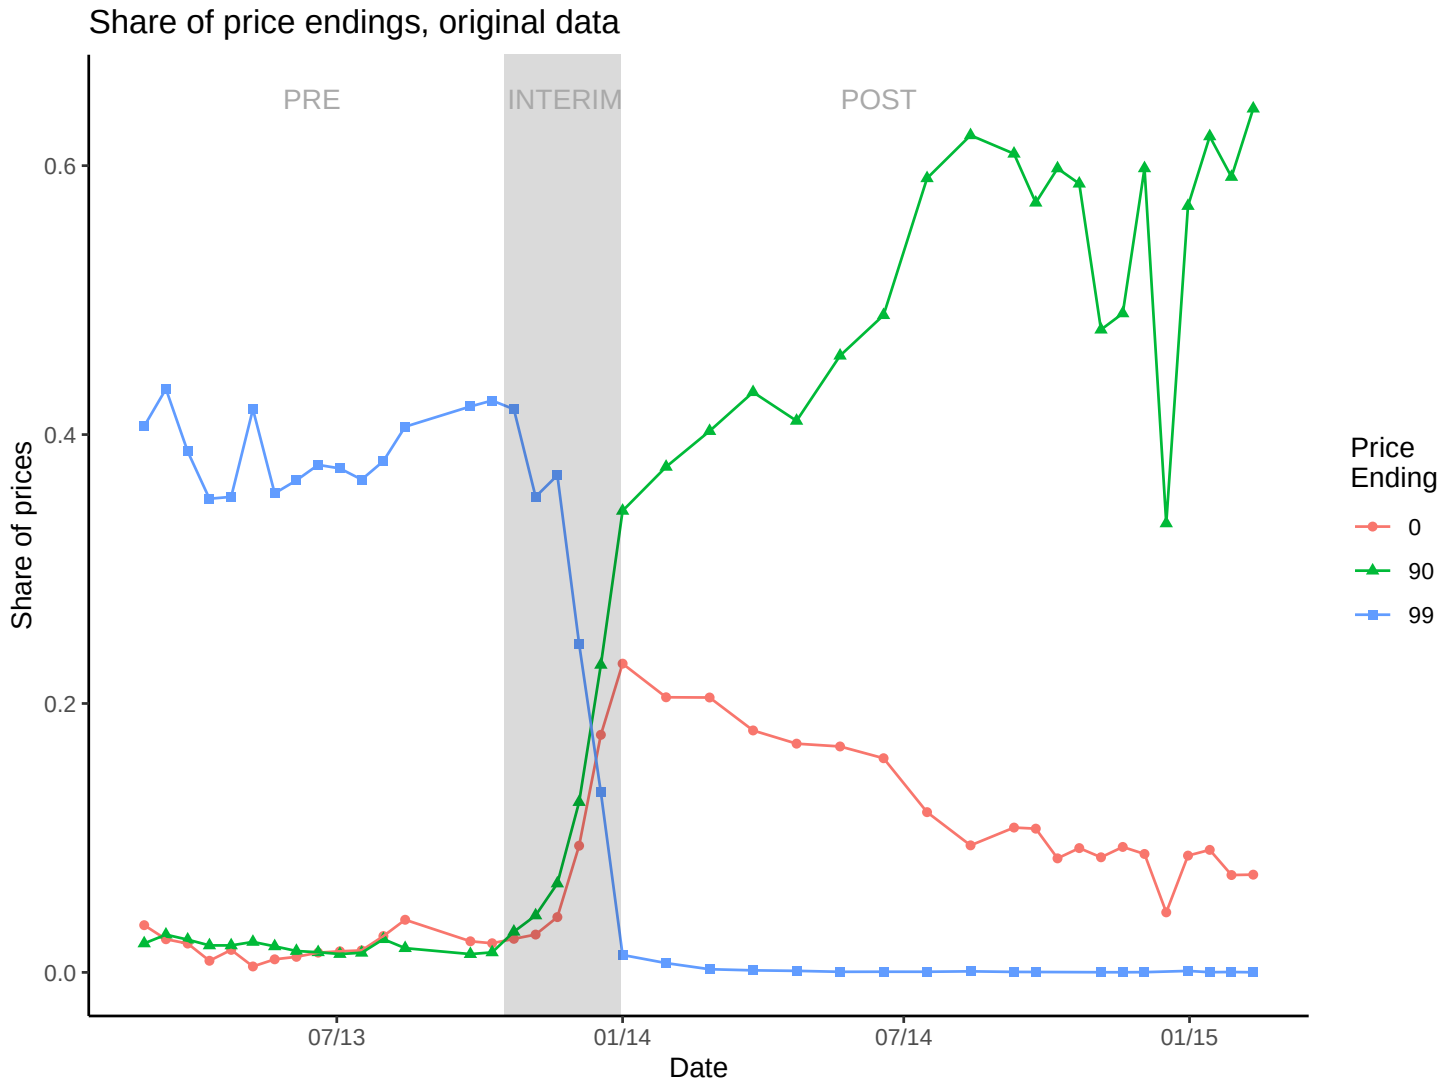
ICC data are collected by surveyors going to a set of predetermined stores with a product list. Each product is equivalent to a UPC. The surveyors record the shelf price of the item, and the total basket price was posted on the ICC website. The micro-data was generously provided to me by the ICC. The 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 (2023)), 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.

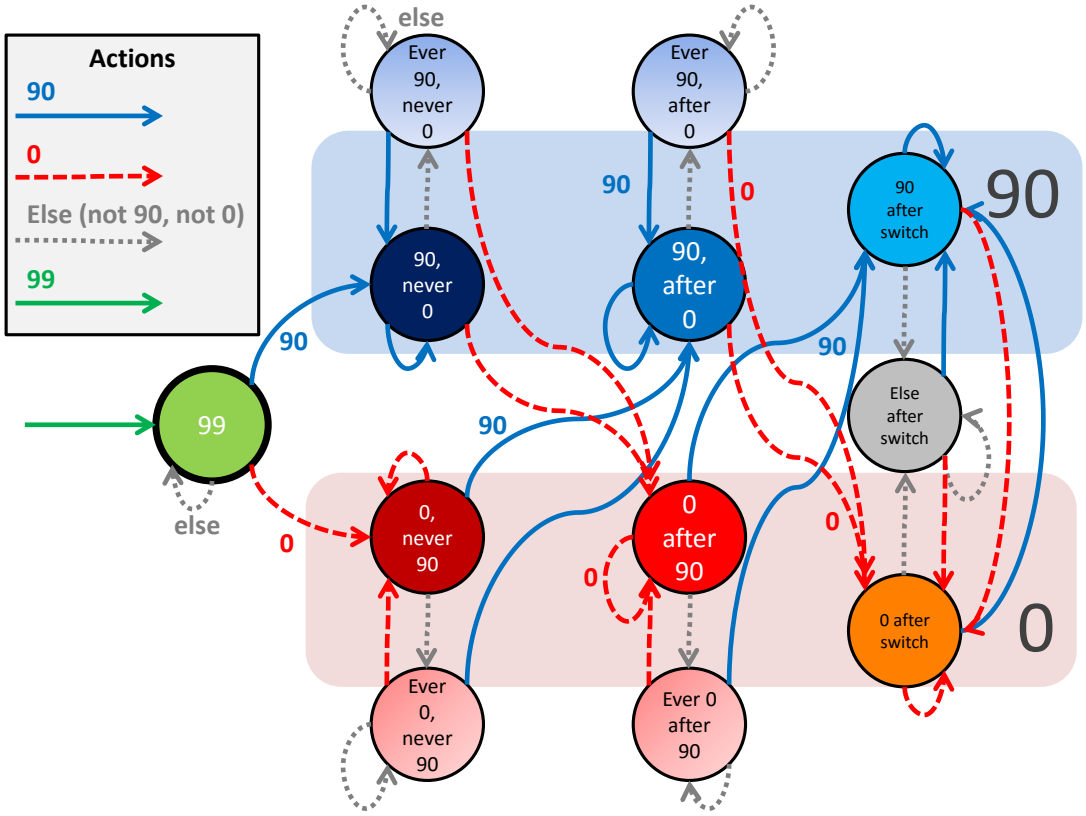
[28] To create a balanced sample of product-stores, I keep product-stores pairs where there are any sales for at least 80% of the total weeks, and for 90% of the weeks of the pre-reform period of 2013. I then calculate for each price sequence of a product in a store the length of same-price spells, deeming 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

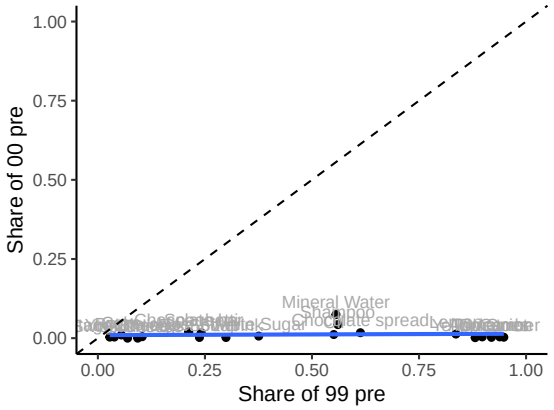
[30] What does it mean to have bias heterogeneity between products? One explanation is that different clientele purchase different products, another is that other characteristics of the products receive more attention than the price (such as observable quality in fresh produce), or that items purchased 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.

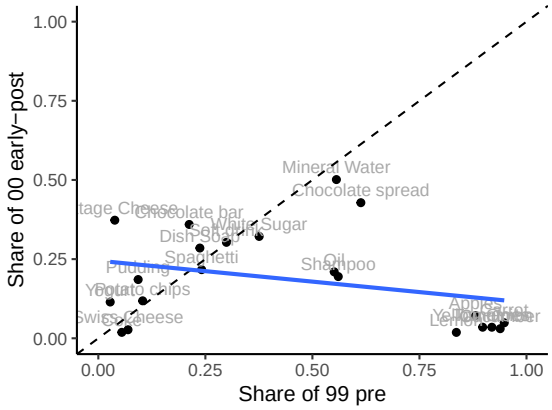
[32] Choice of other radii does not seem to matter. A choice of 5km keeps stores above and below median also within more densely populated areas, and is not a simple separation of urban versus sparse villages. I could not get too narrow due to measurement error in locations, which are based on store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

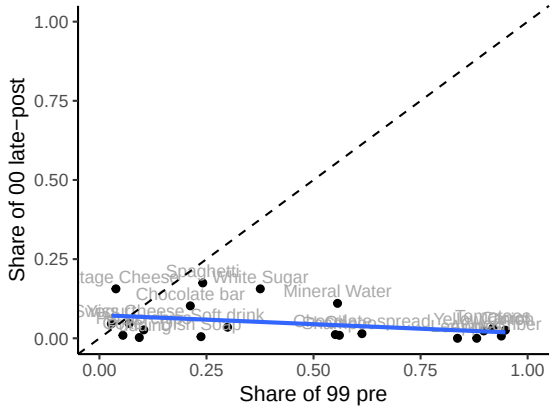
(a) Correlation between share of 99 pre and 00 pre (3-10/2013)



(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:

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

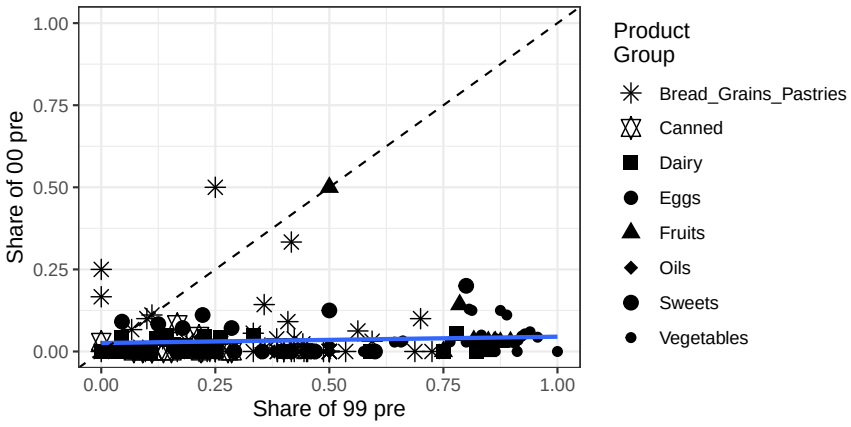
(6) Expand $\mathcal{P}'$ to $\mathcal{P}$ by adding the missing prices between q_i and P_{q_i} and assigning predicted zero shares to them.[34]

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.

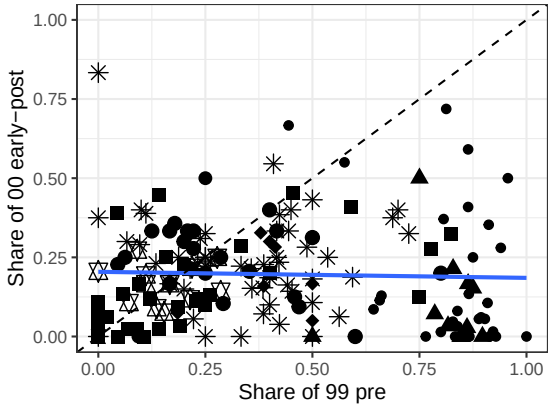
[34] relabel non-9-ending next lowest prices as the lower 9-ending price, without changing the share. For example, 2.51 will be relabeled as 2.49 with a share $\hat{S}_{2.49} = \hat{F}_c\left(\underline{c}_{2.59}\right) - \hat{F}_c\left(\underline{c}_{2.51}\right)$

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.

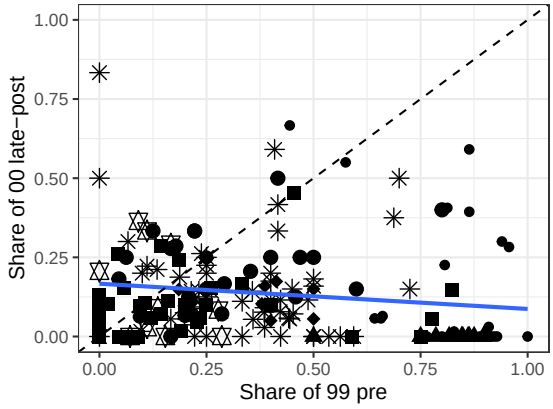
(a) Correlation between share of 99 pre and 00 pre (3-10/2013)



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



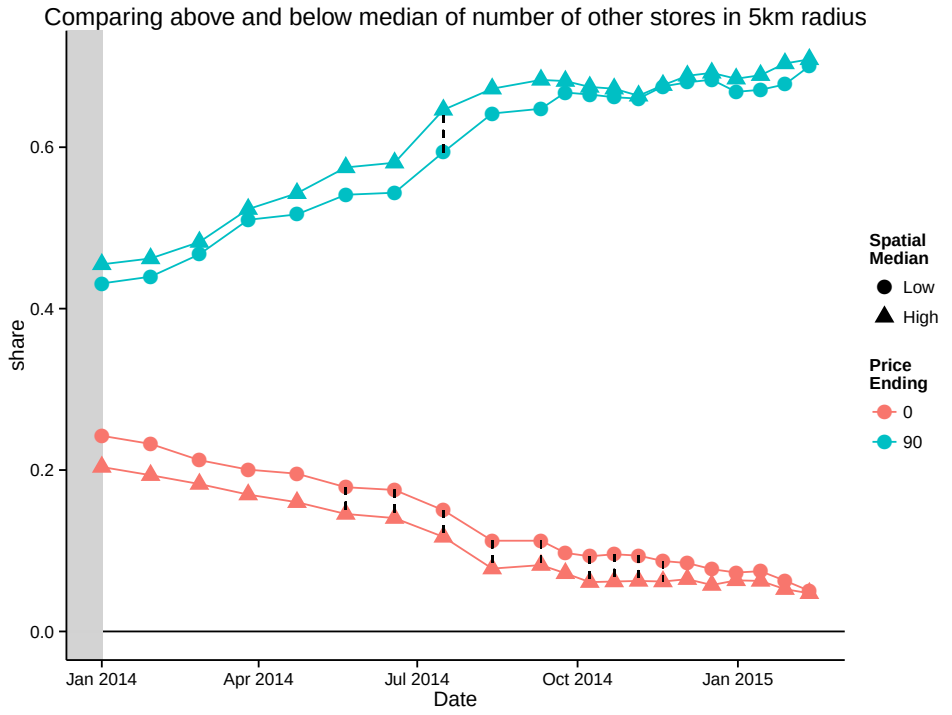
(c) Correlation between share of 99 pre and 00 a year after policy change (1-2/2015)



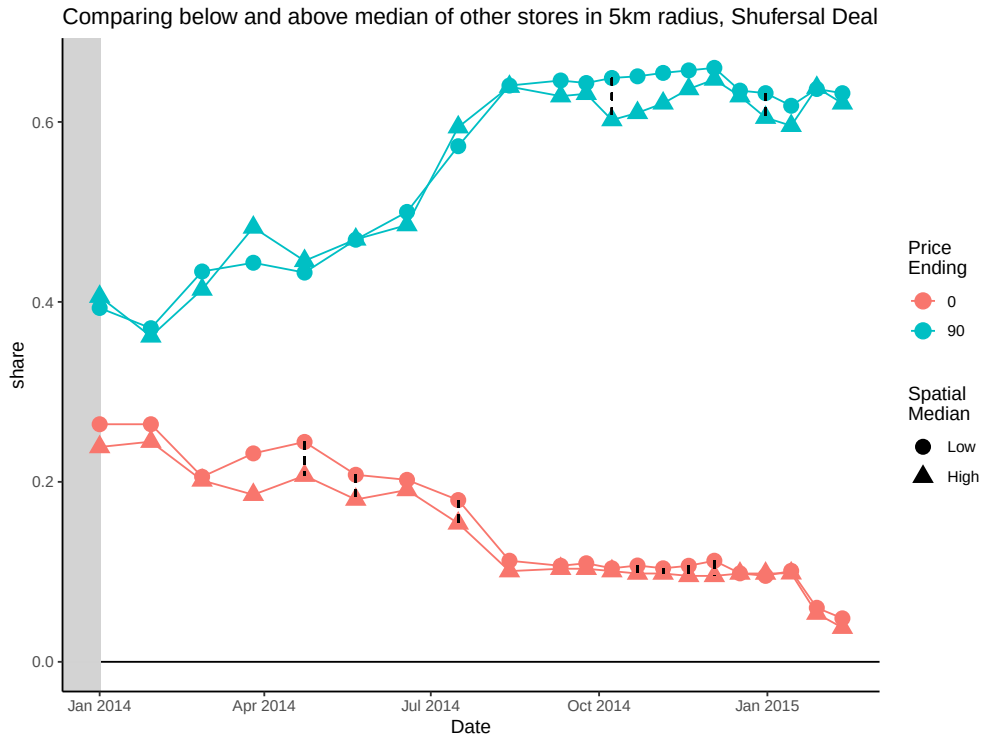
The figures show that shares of 00-ending prices at different periods against 99-ending prices at the pre-period, by product-type in the CBS data. The shares of 00-ending prices per product are calculated before the reform announcement, at the two months following the reform, and at early 2015 a year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