Learning and Limitations of Heuristic Pricing: Evidence from a Reform

Avner Strulov-Shlain*

December 19, 2025

Abstract

Pricing is a core firm activity, yet it is unclear how firms learn to price. Do they learn about demand and respond accordingly, or do they use heuristics and adapt? This paper provides empirical evidence on the use, viability, and weaknesses of heuristic-based pricing. An Israeli reform banned prices ending in non-existent denominations (like .99), requiring supermarkets to adjust most prices while demand remained unchanged. If firms understood demand, they would have immediately avoided .00 endings. Instead, 20% of prices initially shifted to dominated .00 endings before gradually avoiding those again over the following year. This pattern suggests firms price through heuristics rather than by optimizing a known demand function: they learned that .99 endings work without understanding why, and thus could not immediately identify the new optimum. Forgone profits spiked at the reform but ultimately fell below pre-reform levels, suggesting forced exploration improved outcomes. These findings imply that heuristic-based pricing can lead firms towards the optimum but it requires enough exploration and is especially fragile under periods of change.

JEL Codes: D22, D83, L11, L81, M31
Keywords: pricing, heuristics, learning, left-digit bias, firm behavior, behavioral firms

*Strulov-Shlain: University of Chicago Booth School of Business, avner.strulov-shlain@chicagobooth.edu
I am thankful to Stefano DellaVigna for endless encouragement and support, and to Jean-Pierre Dube for honest discussions and detailed feedback. 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, 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. A version of this paper was previously circulated as “Firms have Partial Knowledge: Evidence From A Reform”

1 Introduction

Learning to set prices is important but complicated. The standard assumption in economics is that firms optimize a known objective function: they understand demand and choose prices accordingly. But how well can firms find good prices without actually knowing demand? And what happens when conditions change?

This paper addresses these questions by studying how supermarkets adjusted prices following a reform that restricted available price endings. The evidence suggests that firms can arrive at something close to optimal pricing through trial and error, using heuristics rather than optimizing against a known demand function. This has important implications: positive economics may provide a reasonable description of long-run behavior in stable environments, but it becomes unreliable when firms are shocked out of equilibrium. Since firms do not actually learn the objective *function*, there is an adjustment period where standard economic models get it wrong. In addition, it emphasizes the necessity of exploration in order to improve heuristics.

To illustrate the core identification idea of heuristics versus model-based optimization, consider a thought experiment. Imagine three options with an objective ranking of $a > b > c$. In steady state, the decision maker chooses a, consistent with the ranking. Now suppose option a is removed from the choice set. If the decision maker knew the full ranking, they would immediately switch to b. But if they had arrived at a through heuristics, learning that a works well without understanding *why*, they may not know that $b > c$. In that case, they might initially try c, and only through experimentation learn to switch to b.

The reform studied in this paper mirrors this thought experiment. 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, approximately 45% of supermarket prices ended in 99 (e.g., 2.99), while 00-ending prices were almost entirely absent. This distribution is consistent with pricing that accounts for left-digit biased demand[1]—a demand structure driven by consumers' distorted price perception, which leads to discontinuous drops in demand at round prices.[2] Under such demand, prices that would have ended with low digits in the absence of bias (e.g., 5.00 or 5.29) are better set with a 99-ending (e.g., 4.99). In the notation of the thought experiment, 99-ending prices represent option a, 90-ending prices option b, and 00-ending prices option c, with $a > b > c$.

When 99-endings were banned, model-based pricing predictions still implied that 00-ending prices should be avoided. The magnitude of left-digit bias that made 99-endings attractive also

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

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

makes 90-endings dominate 00-endings, even though the trade-off is stronger.[3] If firms understood the demand structure, they would have immediately shifted to 90-endings. Instead, at the time of the reform, 20% of prices ended in 00. In the years that followed, the share of 00-endings declined as firms explored and learned. Products whose prices shifted from 99 to 90 tended to remain at 90, while those that shifted to 00 typically reverted to 90 over time. About a year after the reform, firms had largely converged to 90-endings. Heterogeneity across products and stores provides additional texture: products with higher pre-reform shares of 99-endings, presumably reflecting stronger left-digit bias, were quicker to converge to 90-endings, while local competitive pressure had little discernible effect on learning speed.

This pattern is difficult to reconcile with standard optimization against a known demand function. A firm that knew demand would recognize immediately that 00-ending prices are dominated—the same logic that made 99-endings attractive before the reform applies directly. But the pattern is consistent with firms using heuristics to find good prices without actually knowing the underlying demand structure. Before the reform, the heuristic "use 99-endings" happened to produce something that appears like optimal pricing under left-digit bias. After the reform removed this option, firms had to learn anew, and the adjustment took time.

A structural model quantifies the costs of this adjustment process. Using estimated demand parameters and observed price distributions, I compare actual pricing to a counterfactual of full optimization. Before the reform, forgone profits were on the order of 2%. Firms were close but not at the optimum.[4] At the reform's onset, forgone profits spiked as firms chose dominated 00-endings. Over the following year, as firms learned, forgone profits declined and eventually fell below pre-reform levels. Forced exploration appears to have improved outcomes, with firms settling into better pricing patterns than before; on the other hand, the new heuristic still keeps firms closer to- but not at the- frictionless optimum.

This paper does *not* claim that heuristic-based pricing is necessarily suboptimal. Estimating demand is costly, and heuristics may be a reasonable response. The point is descriptive: firms find not-terrible prices without knowing the objective function, and this has implications for both researchers and practitioners. For researchers, positive economics may describe steady-state behavior reasonably well, but predictions about responses to shocks require caution. For practitioners, it highlights the value of experimentation and the risks of relying on heuristics when conditions change.

The main contribution is to document a case where firms learned through forced exploration, and to show that this learning was slow and costly. The setting is unusual: firms had to optimize against consumer behavior with left-digit bias, a problem with discontinuities that makes standard intuitions unreliable. That firms solved this problem heuristically, without knowing demand, is

[3]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-endings, but not 00-endings.

[4]Whether 2% is a large or small number is a matter of interpretation.

both reassuring (heuristics can work) and concerning (firms may not adapt quickly to change). Related work has documented firm learning in other contexts (Doraszelski et al. (2018); Huang et al. (2018)), but typically where firms must learn a new demand structure. Here, demand did not change, only the choice set did. Yet firms still required substantial time to adjust.

This paper also contributes to the literature on firm misoptimization. A growing body of work documents persistent deviations from profit maximization (Bloom and Van Reenen, 2007; Cho and Rust, 2010; DellaVigna and Gentzkow, 2019; Goldfarb and Xiao, 2011; Hortaçsu et al., 2019; Hortacsu et al., 2021; Huang, 2021; Shapiro et al., 2021; Strulov-Shlain, 2023). These findings raise a puzzle: how can firms persistently fail to optimize when the stakes are high? One answer is that firms face constraints that prevent optimization. Another, suggested here, is that firms may not actually know the objective function. They arrive at reasonable strategies through trial and error, but lack the knowledge to fully solve more complicated problems, and to immediately adapt when conditions change. This interpretation, which I call "partial optimization", provides a parsimonious explanation for both the steady-state behavior (close to but not at optimal) and the transition dynamics (slow adjustment with initial mistakes).

The paper proceeds as follows. Section 2 describes the setting, data, and left-digit biased demand. Section 3 presents the results on pricing patterns and demand response. Section 4 interprets the findings. Section 5 discusses implications.

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 at checkout; 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.

[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

For the purposes of this study, the reform is useful because it forced firms to adjust prices, without changing the underlying demand structure. The reform restricted the choice set by removing 99-endings as an option, but did not alter how consumers respond to prices. This creates an opportunity to observe how firms behave when their existing pricing strategies are no longer available.

2.2 Data

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.

Pricing data. 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. Appendix Figure A-1 shows that the main pricing patterns are qualitatively similar using raw data without these restrictions. 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%.

Table 1: Summary Statistics of Prices Data

Variable	ICC Original	ICC Main	CBS Original	CBS Main	StoreNext Original	StoreNext 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.

Scanner data. 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

[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.

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, focusing 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 Left-digit bias and pricing predictions

A key feature of supermarket pricing is left-digit bias: the tendency of consumers to overweight 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 round-number thresholds. The implication for pricing is that firms should avoid setting prices just above those discontinuities.

The literature supports this mechanism. Early work on psychological pricing reached mixed conclusions (Basu, 1997; Gedenk and Sattler, 1999; Ginzberg, 1936) and suggested multiple potential 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 derive pricing predictions, 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:

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

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. Assuming constant elasticity demand:

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

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 intuition is that a small price reduction from a round number (e.g., from 5.00 to 4.99) generates a large perceived price drop, making it profitable to price just below the threshold.

When should 00-endings be avoided? A key prediction concerns the *Next-Lowest Price P*: the lowest price above a just-below price q that could be optimal for some cost level. The just-below price ends with 99 before the reform and 90 after. If P exceeds the round number (i.e., $P > \lfloor P \rfloor$), then 00-ending prices are never optimal—they are strictly dominated by either the just-below price or a higher price.

Solving for the threshold level of left-digit bias θ_0 at which 00-endings become dominated yields remarkably low values. For an extreme elasticity of -7 and price level around 2, θ_0 is approximately 0.00075. For more typical elasticities of -1 to -4, the threshold is even lower.[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 implies that for any reasonable levels of left-digit bias and price elasticity, *if firms optimize against the true demand function*, we should observe no 00-ending pricing when 99-endings are available.

Predictions after the reform. After the reform, the just-below price changed from 99 to 90, and the next higher admissible price from 00 to 10. These changes make just-below pricing less profitable (the gap between 4.90 and 5.00 is larger than between 4.99 and 5.00) and upward adjustments necessarily larger. Re-solving for θ_0^{Post} under these conditions raises the threshold, but it remains low (e.g., 0.02 for $\epsilon = -4$ and price level around 2), still well below empirically observed levels of left-digit bias.

[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 $\lfloor P \rfloor$. Less elastic demand and higher prices therefore make 00-endings even less attractive.

Figure 1 shows the $\theta_0(\epsilon)$ values before and after the reform, alongside empirical estimates from Strulov-Shlain (2023). The key point is that the same logic that made 99-endings attractive before the reform also makes 90-endings dominate 00-endings after the reform. A firm that understood the demand structure would recognize this immediately.

What we learn from 00-endings. The prediction that 00-endings should be avoided provides a useful diagnostic. If firms understand that they face left-digit biased demand, even approximately and even if they underestimate it, they should avoid 00-endings both before and after the reform. Observing substantial 00-ending prices after the reform, when 90-endings are available, suggests that firms did not fully understand the demand structure. They may have known that 99-endings worked well without knowing *why*, and therefore could not immediately infer that 90-endings would dominate 00-endings.

This is the sense in which the reform distinguishes heuristic pricing from model-based optimization. A heuristic like "use 99-endings" produces good outcomes before the reform but provides limited guidance afterward. Model-based optimization, by contrast, would immediately identify 90-endings as the new optimum.

3 Results

This section presents the empirical findings on pricing patterns and demand response. I organize the results around three questions: How did price endings change after the reform? How did individual products transition between price endings over time? And what was the demand consequence of choosing 00-endings over 90-endings? I defer interpretation to Section 4.

3.1 Price-ending shares over time

Figure 2 shows the shares of prices ending in 99, 90, and 00 by period. The shaded area represents the window between the announcement (October 2013) and enactment (January 2014) 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.[11] This pattern is consistent with the prediction from Section 2.3: under left-digit biased demand, 99-endings should be common and 00-endings rare.

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

Figure 1: Minimal levels of left-digit bias $\theta_0(\epsilon)$ above which 00-ending is never optimal

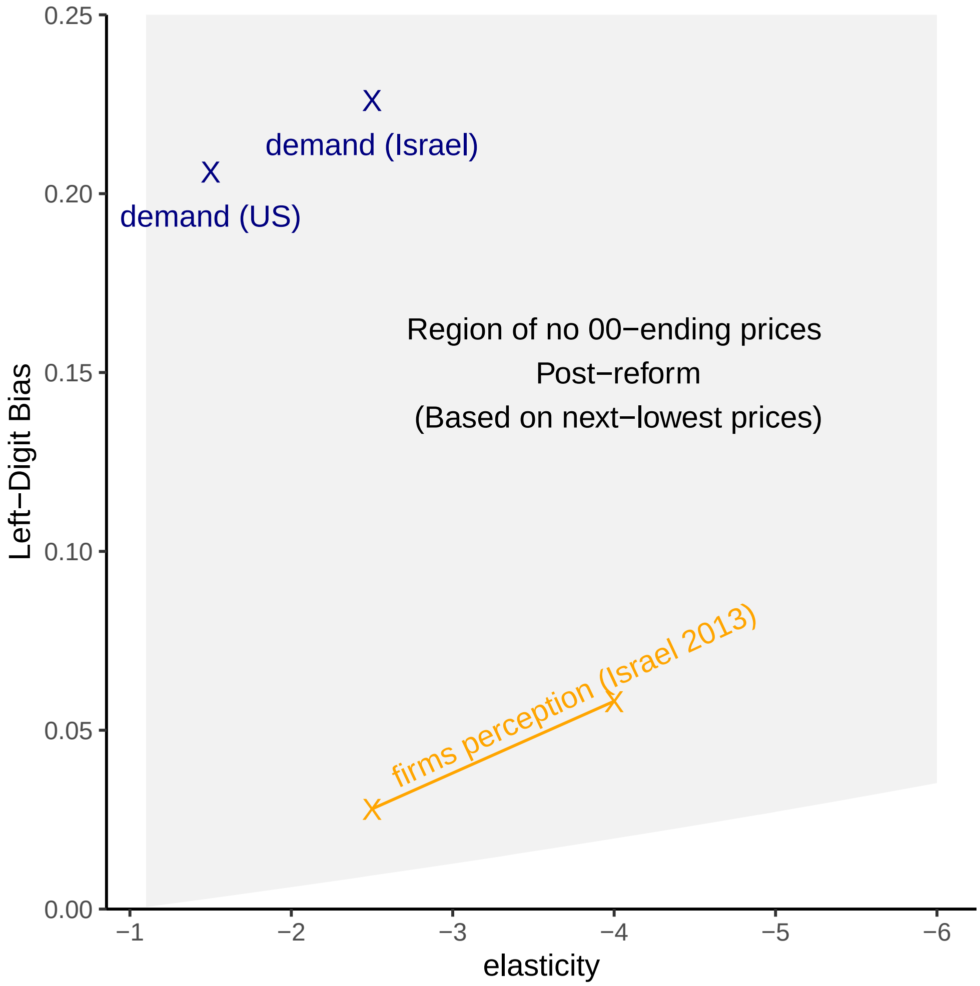

The figure shows the minimal levels of left-digit bias θ above which 00-ending prices are never optimal. The values are a function of the price elasticity and nominal price level. The dark line shows the threshold values in the pre-reform case (when 99-ending and 01-ending prices were allowed), and the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

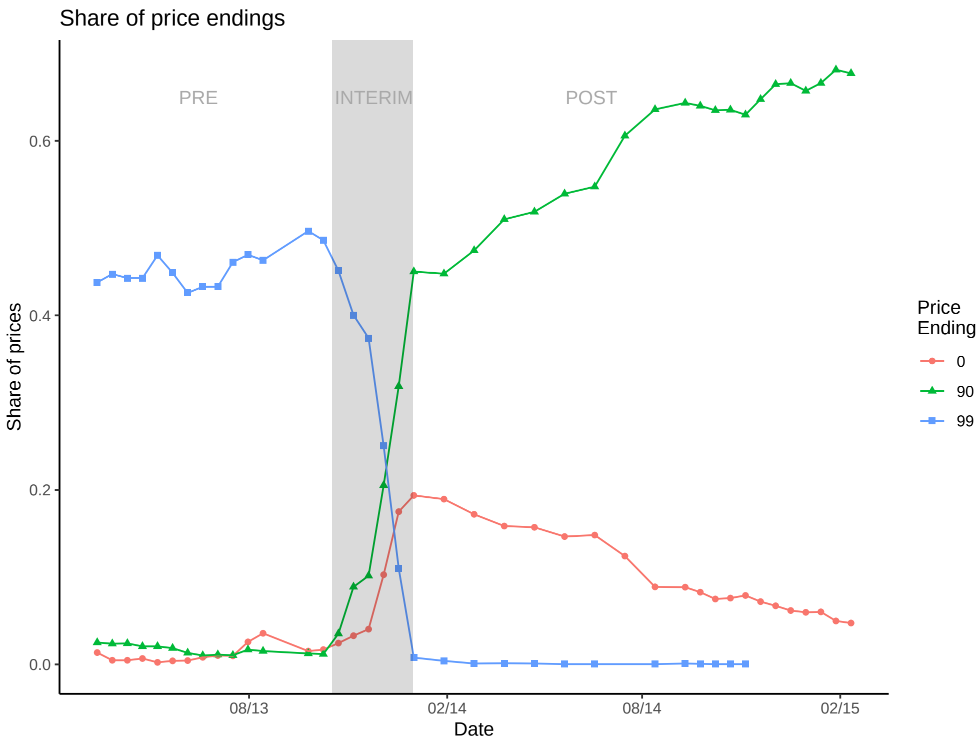

(b) Shares of products that end with 99, 90 and 00, CBS data

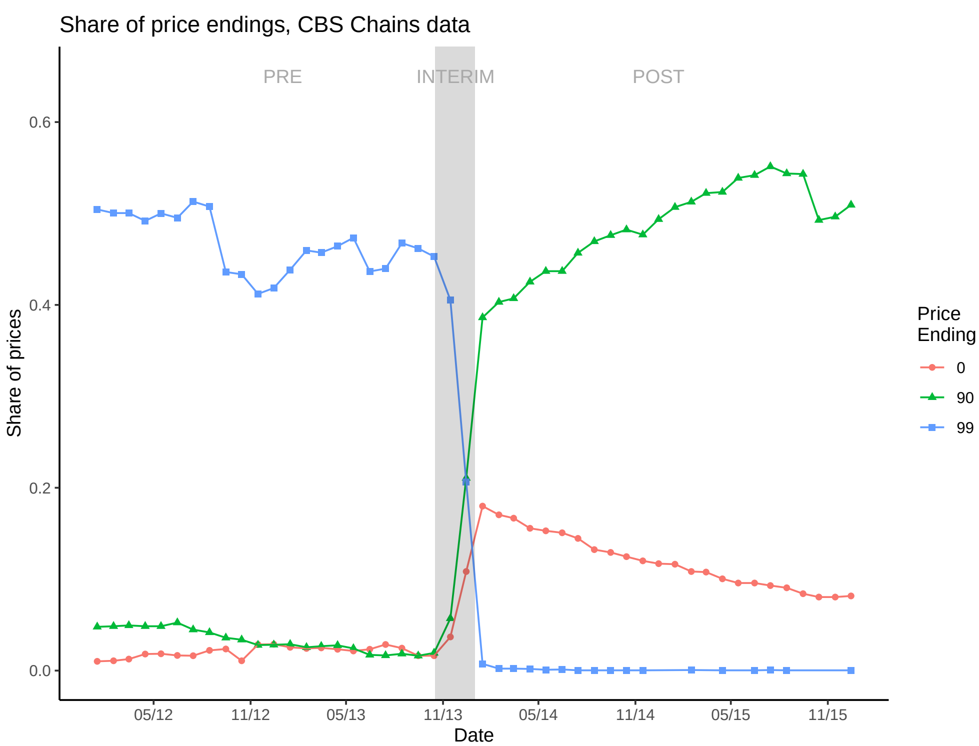

The figures above show that shares of 99-, 90-, and 00-ending prices across products and stores by sampling period. The shaded area represents the time between announcement of the reform on October 17, 2013, and its enactment on January 1, 2014. Squares represent 99-ending prices, triangles 90-ending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

Transition period. 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.

After the reform. When the policy took effect in January 2014, approximately 40% of prices ended in 90 and 20% in 00. The 20% share of 00-endings is notable: the model predicts that 00-endings should be avoided, yet one-fifth of prices were set at this dominated ending.

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 exceeded 65%, higher than the pre-reform share of 99-endings. 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-chain 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.

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.

The stepwise adjustments are informative. If the gradual aggregate decline were driven by slow price adjustment due to menu costs or other frictions, we would expect smooth declines at the chain level as well. Instead, some chains exhibit discrete drops, suggesting deliberate chain-wide policy changes rather than gradual price updating.

In summary, the share of 00-endings was low pre-reform, rose substantially at enactment, and then declined steadily over the following year. The 90-endings that replaced 99-endings proved stable, while 00-endings did not.

3.2 Heterogeneity across products and stores

The aggregate patterns mask meaningful heterogeneity across products and stores. Two dimensions are particularly informative: variation by product category and variation by local competitive environment.

Product-level heterogeneity. If firms understood the demand structure, products with stronger left-digit bias should be less likely to have 00-endings after the reform, following the same logic that made 99-endings attractive for those products should make 90-endings attractive post-reform. I proxy for left-digit bias intensity using the pre-reform share of 99-endings: products with more 99-endings presumably faced stronger bias (or at least stronger perceived bias).[12]

Figure 4 shows the relationship between pre-reform 99-ending shares and post-reform 00-ending shares by product. Before the reform (panel a), 00-endings were rare regardless of 99-ending prevalence, consistent with the prediction that 00-endings are dominated. Immediately after the reform (panel b), the relationship is more complex. Fresh produce—products with very high pre-reform 99-ending shares translated almost immediately to 90-endings, with virtually no 00-endings. For other products, however, there is a *positive* correlation: products with more 99-endings before the reform had *more* 00-endings immediately after. This suggests that for many products firms rounded 99-ending prices up to 00 rather than down to 90, without understanding the logic that made 99-endings optimal in the first place.

[12]What does it mean to have bias heterogeneity between products? Possible explanations include different clientele purchasing different products, other product characteristics receiving more attention than price (such as observable quality in fresh produce), or items purchased in bulk requiring price multiplication, which puts extra weight on leftmost digits.

Figure 4: 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)

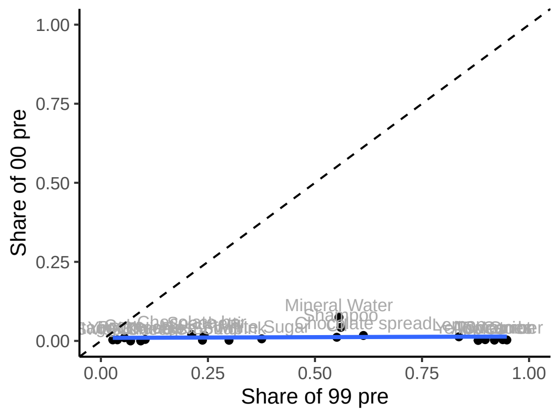

(b) Correlation between share of 99 pre and 00 immediately after policy change (1-2/2014)

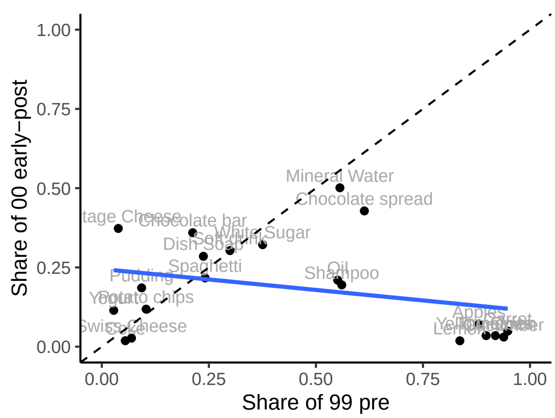

(c) Correlation between share of 99 pre and 00 a year after policy change (1-2/2015)

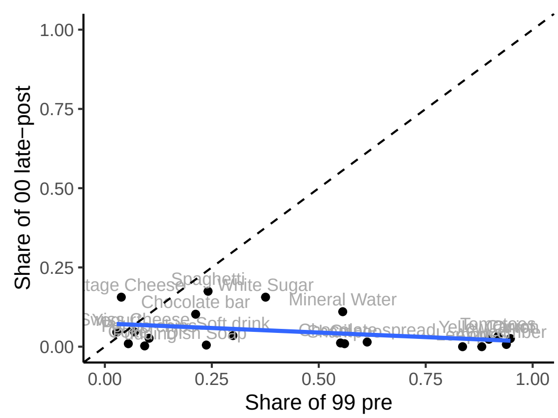

The figures show that shares of 00-ending prices at different periods against 99-ending prices at the pre-period, by product. The shares of 00-ending prices per product are calculated before the reform announcement, at the two months following the reform, and at early 2015 a year after the reform.

A year after the reform (panel c), the pattern shifts. Products with higher pre-reform 99-ending shares now have fewer 00-endings, consistent with learning: firms eventually figured out that 90-endings work better, and this learning happened faster for products where left-digit bias was stronger (or more salient). The CBS data show similar patterns, though less pronounced (Appendix Figure A-3).

Spatial competition. Does competition speed up learning? Stores facing more local competitors might learn faster if competitive pressure punishes mistakes more severely, or if firms observe and imitate successful competitors. Alternatively, competition might not matter if learning occurs at the chain level rather than the store level.

I examine this by comparing stores above and below the median number of competitors within a 5km radius.[13] Figure 4a shows that stores in more competitive areas had *lower* 00-ending shares initially. They were less likely to make the "mistake" of choosing 00-endings at the reform's onset. The differences are statistically significant in the second half of 2014.

[13]The choice of radius does not qualitatively affect the results. A 5km radius keeps the comparison meaningful within densely populated areas rather than simply separating urban from rural stores.

Figure 5: Spatial competition and shares of price ending

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

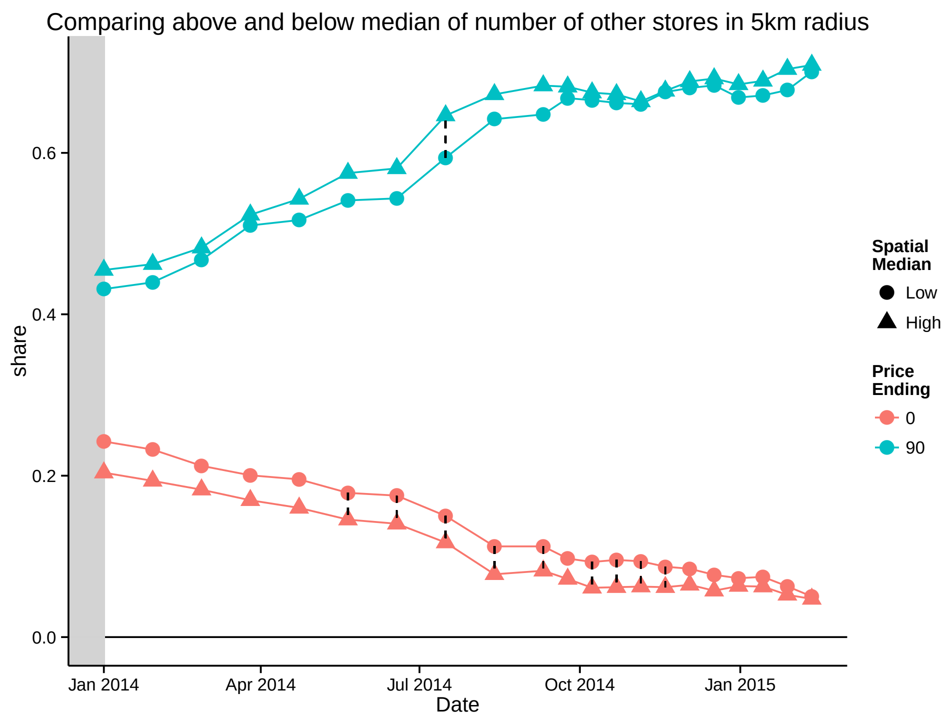

(b) Shares of 00 and 90 for "Shufersal Deal" stores above and below median of spatial competition

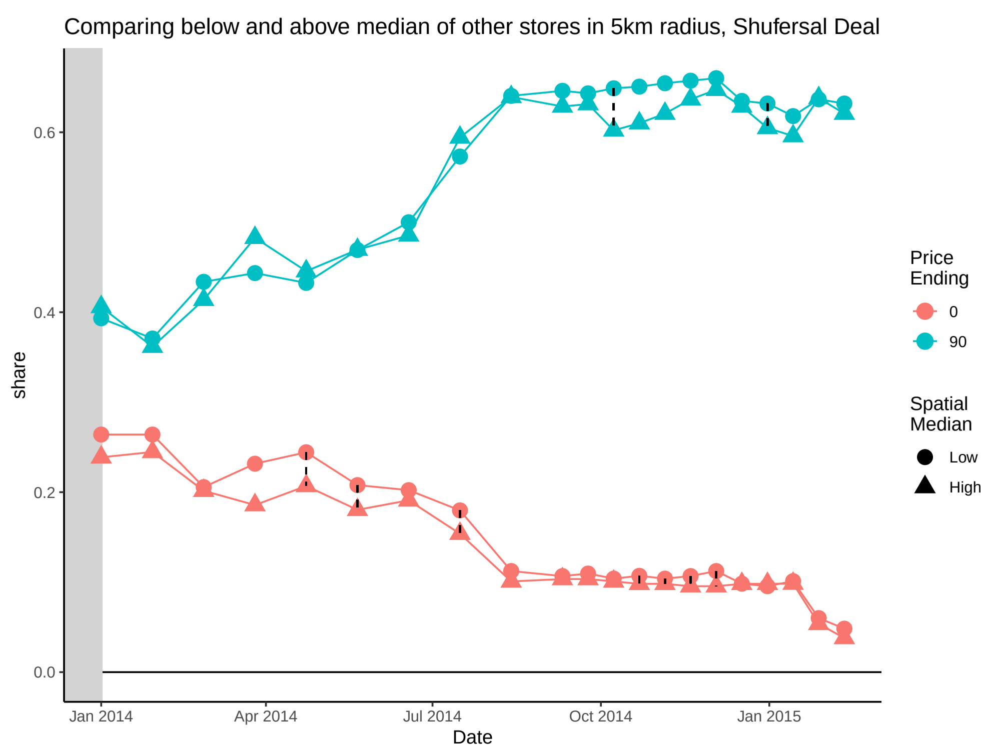

The figure shows the shares of 00- and 90-ending prices at the post-period for stores that are above and below the median of number of local competitors. The top panel show all stores, and the bottom panel only the stores of the Shfersal Deal chain.

However, this cross-sectional pattern could reflect chain composition rather than competition effects. Namely, different chains locate differently across areas. To address this, I restrict to stores within a single chain (Shufersal Deal). Figure 4b shows that within-chain differences are less pronounced. The gap is largest in January 2014 and persists somewhat, but there is no clear evidence that competition speeds convergence to 90-endings. If anything, stores in more competitive areas converged *more slowly*.

These findings suggest that learning occurred primarily at the chain level rather than through local competitive dynamics. The between-chain heterogeneity documented in Figure 3, where some chains made discrete policy changes, is more consistent with centralized learning than with store-level adaptation to local competitive pressure.[14]

3.3 Price-ending paths at the product-store level

The aggregate patterns raise a question: were some products consistently priced at 90-endings and others at 00, or did individual products transition between endings over time? To answer this, I track price-ending paths at the product-store level.

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 Appendix Figure A-2. For example, the path "90, never 00" means that a product changed first to a 90-ending price and only 90-ending prices afterwards.

Figure 6 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).[15]

[14]Competition can still be a driving force of learning if different chains faced different intensity of competition.

[15]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.

Figure 6: Price ending paths of products that ended with 99 before the policy change

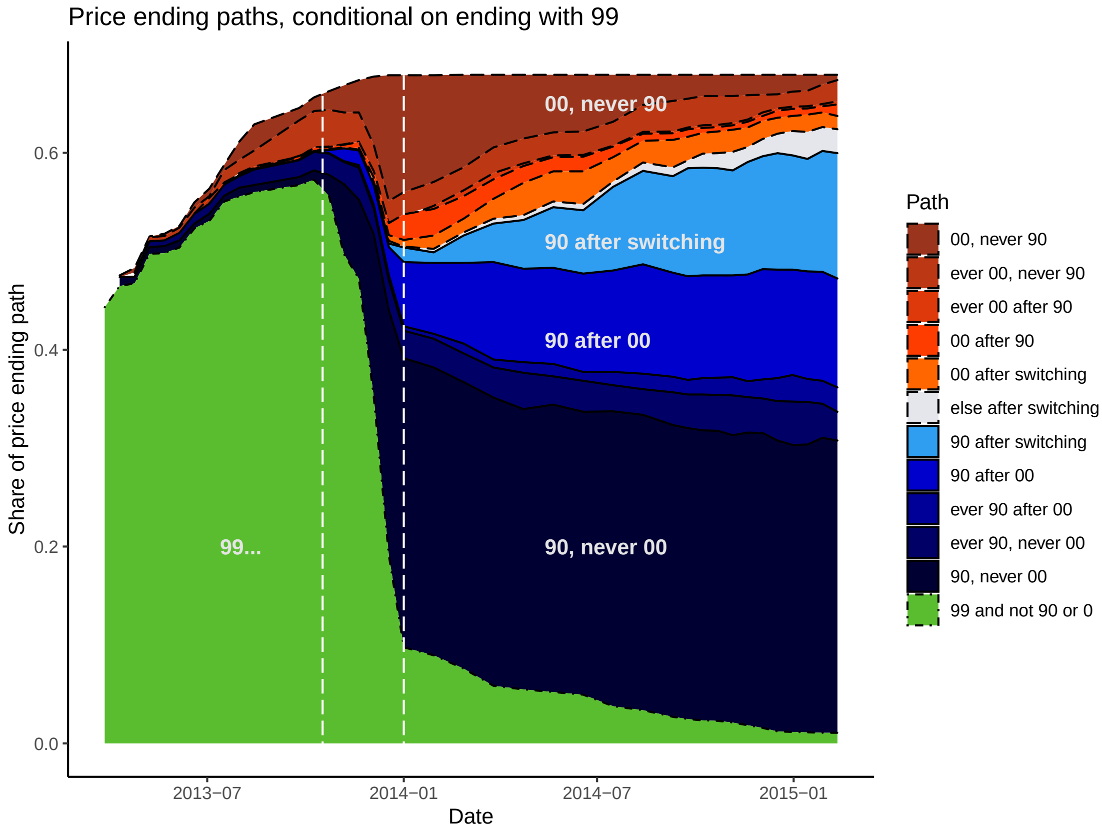

The figure shows the shares of each price-ending path (which are described in Figure A-2). The green area is the 99 or other price endings, the blue shades are the 90-ending prices, and the red shades are the 00-ending prices.

Three patterns emerge:

First, products that switch to 90 tend to stay at 90. This is evident from the large, stable dark-blue "90, never 00" segment versus the small orange "00 after 90." Once a product is priced at a 90-ending, it is unlikely to move to 00.

Second, products that switch to 00 tend to move to 90. The shrinking dark-red "00, never 90" segment contrasts with the growing blue "90 after 00." 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.[16]

Third, products that switch between endings eventually settle at 90. A small set of products exhibit both transitions (00 after 90 and 90 after 00), suggesting more deliberate exploration or experimentation (e.g., 99 → 90 → 00 → 90). These products tend to eventually settle at 90. In the last period, switchers are 8.4 times more likely to end at 90 than at 00.

[16]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.

Overall, product-stores that try a 90-ending tend to stick with it, while 00-endings are unstable and typically transition to 90. The asymmetry is striking: 90-endings are effectively absorbing, while 00-endings are not.

3.4 Demand response to price endings

The patterns above show that firms moved away from 00-endings over time. But did 00-endings actually hurt demand? The model predicts they should, and if firms were learning from experience, negative feedback from 00-endings could explain the observed transitions.

Identification strategy. 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 across periods.[17] An alternative design randomizes comparable products' prices across stores (e.g., Chetty et al. (2009)).

The reform creates a quasi-experimental version of both designs. Consider a product priced at 5.99 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 in one store. For the across-store design, imagine one store sets 6.00 and another sets 5.90. The analysis below leverages such variation.[18]

Specification. Using StoreNext scanner data, I estimate:

$$\log Q_{ist} = \alpha^{90}\kappa^{90}_{ist} + \alpha^{00}\kappa^{00}_{ist} + \alpha^{99}\kappa^{99}_{ist} + \beta_{ic(s)} \log P_{ist} + X'_{ist}\beta + \epsilon_{ist} \tag{3}$$

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}; 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). X includes controls: product-by-store fixed effects, an on-sale proxy,[19] same-price spell-length decile-by-product fixed effects, month-by-product fixed effects, and year-by-product fixed effects.

I instrument for the product-chain price elasticity term using the leave-out average price of the product across stores in the same chain (Hausman-style IV; Hausman (1996)). As long as pricing conduct takes cost into account, prices in other stores of the chain proxy for cost shocks but are plausibly orthogonal to local demand shocks.

[17] 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.

[18] Strulov-Shlain (2023) test for these demand drops using larger and higher-quality data, but without the quasi-random variation created by the reform.

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

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 3 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.

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.[20]

Results. Table 2 reports the estimates. The contemporaneous excess demand difference between a 90-ending and a 10-Agorot-higher 00-ending 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.

Column (2) most closely mirrors the within-store experiment (products that have both 90 and 00 prices in the same store): the demand difference is 8.2% ($p = 0.025$). Columns (3)–(4) align with

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

the across-store design: the difference is 5.8% ($p = 0.061$) in column (3), and 8.6% ($p = 0.0045$) in column (4).[21]

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

The demand analysis confirms that 00-endings were costly. Firms that chose 00-endings experienced lower demand than they would have with 90-endings, providing a potential mechanism for the learning documented in the price-ending paths.

3.5 Quantifying forgone profits

The previous results establish that (1) many firms initially chose 00-endings, (2) they gradually shifted to 90-endings, and (3) 00-endings hurt demand. I now quantify the profit consequences of these patterns by comparing observed pricing to a counterfactual of full optimization.

Approach. Quantifying forgone profits requires comparing actual pricing to a benchmark of what firms *should* have done. I use a two-product logit demand model with left-digit bias, calibrated to match the estimated price elasticity and left-digit bias from the data. For each of many simulated cost draws, I compute the profit-maximizing price and compare it to the price implied by a descriptive "rule-of-thumb" heuristic that matches observed pricing patterns.

The heuristic extends Strulov-Shlain (2023): firms compute a naive optimal price ignoring left-digit bias, then round to 99/90-endings or 00-endings with probabilities estimated month-by-month from the data. Forgone profits equal the percentage gap between profits under heuristic pricing and profits under optimal pricing. Appendix C provides complete details on the demand model, calibration, heuristic specification, and computation.

Results. Figure 7 shows the results. Before the reform, forgone profits hover around 2%—firms were not at the optimum, but not far from it.[22] At the reform's onset, forgone profits spike to about 2.75% as firms initially chose dominated 00-endings. Over the following year, forgone profits gradually decline as firms learn, eventually falling to roughly 1.7% by 2015—below the pre-reform level.

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

[22] Whether 2% is "large" or "small" depends on context. It is smaller than many estimates of pricing inefficiencies in other settings, but represents meaningful forgone revenue for a competitive industry with thin margins.

Figure 7: Estimated counterfactual forgone profits month by month

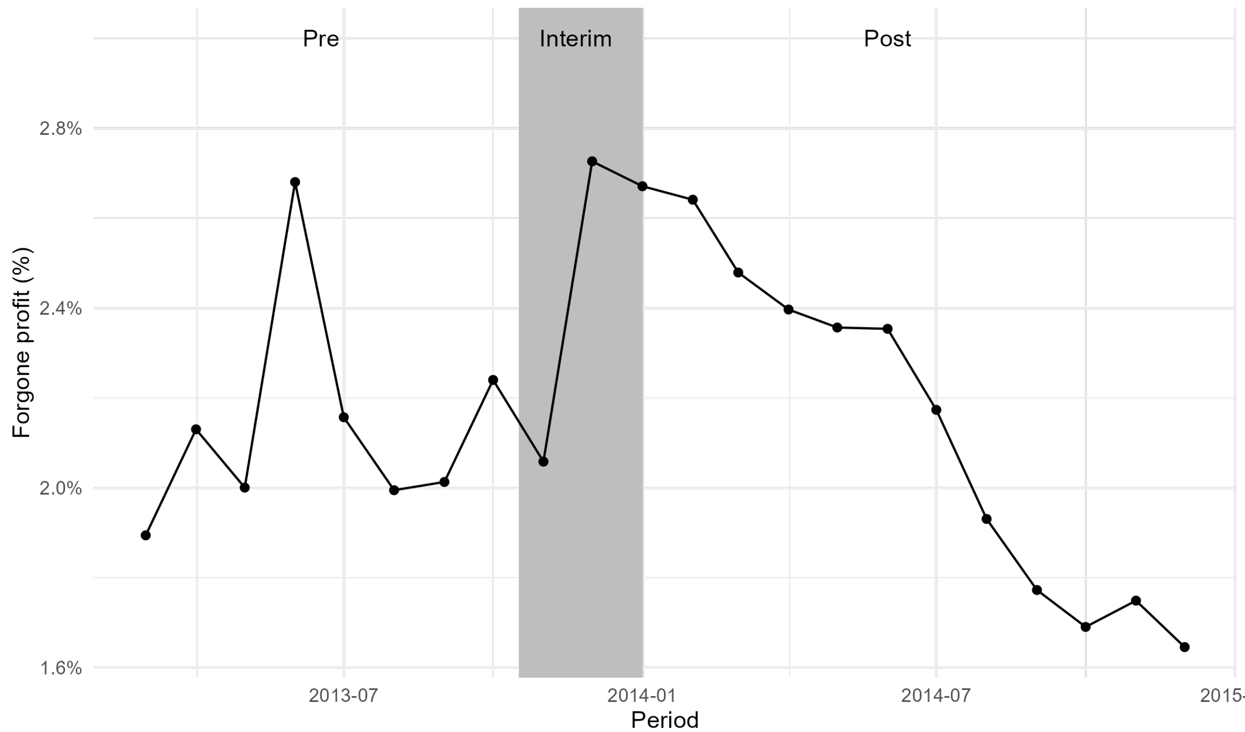

Figure shows counterfactual profits under a monthly estimated predictive rule-of-thumb pricing relative to full optimization.

The decline in forgone profits is notable: forced exploration appears to have improved outcomes. By removing 99-endings from the choice set, the reform forced firms to explore alternatives, and this experimentation led to better pricing than before. By the end of the sample, firms had not only recovered from the initial spike but achieved lower forgone profits than in the pre-reform steady state.

These magnitudes should be viewed as rough estimates rather than precise measurements. The quantification relies on functional form assumptions and simulated costs. The key message is qualitative: firms were not fully optimizing before the reform, the reform initially made things meaningfully worse, and learning brought improvements over time.

4 Interpretation

The key puzzle is the initial prevalence of 00-endings. If firms understood demand, they would have recognized immediately that 90-endings dominate 00-endings—the same logic that made 99-endings attractive before the reform. Yet many firms chose 00-endings initially and only learned to avoid them through experience. This section argues that heuristic-based pricing provides the most parsimonious explanation.

4.1 Heuristics versus model-based optimization

Consider two stylized descriptions of firm behavior.

Under *model-based optimization*, firms maintain a model of demand and choose prices to maximize profits given that model. Learning means updating beliefs about model parameters while the functional form remains fixed. Crucially, updating beliefs about one parameter updates the expected profitability of *all* prices, because the model links them together.

Under *heuristic-based pricing*, firms learn which actions work well through trial and error without maintaining an explicit demand model. A heuristic like "use 99-endings" can emerge from accumulated experience without the firm understanding *why* it works. Learning is local: observing the outcome of one price does not automatically update beliefs about untried prices.[23]

Pre-reform behavior is consistent with either description since both would produce 99-endings and no 00-endings. Post-reform behavior distinguishes them. Model-based optimization predicts immediate adoption of 90-endings; heuristic-based pricing predicts initial uncertainty followed by learning. The data support the latter.

4.2 Could wrong beliefs explain the patterns?

One might preserve the model-based framework by allowing firms to hold incorrect beliefs. Perhaps firms underestimated left-digit bias, leading some to believe 00-endings were acceptable.

This explanation ultimately fails. Using the Method of Simulated Moments as in Strulov-Shlain (2023) (details in Appendix D), I estimate that firms priced as if they believed $\theta = 0.02$–0.04, much smaller than the demand-side estimates of $\theta = 0.22$–0.36, but still positive. Even these underestimated beliefs are inconsistent with 00-endings post-reform: the threshold needed to rationalize any 00-endings is under 0.01. If firms held stable beliefs and optimized against them, they would have avoided 00-endings after the reform just as they did before.

Moreover, an excess of *both* 00- and 90-endings immediately after the reform is inconsistent with any coherent belief about left-digit bias. No stable belief can generate both patterns. The heuristic interpretation that firms knew 99-endings worked without knowing why is more parsimonious.[24]

[23]This description aligns with model-free reinforcement learning (Sutton and Barto, 2018). The agent maintains a value function over actions, updated based on observed rewards, but does not use a model to generalize across actions.

[24]An alternative that cannot be ruled out is that firms optimized against a misspecified model (e.g., believing 99-endings are special for reasons unrelated to left-digit bias). But if so, sufficient exploration over the decades preceding the reform should have led them to reject this model.

4.3 The nature of the optimization problem

What makes this case particularly informative is the nature of the underlying problem. Pricing under left-digit bias involves discontinuities at round-number thresholds and multiple local optima. Standard intuitions that optimal prices vary smoothly with costs and small price changes have small effects do not apply.

Yet firms moved towards the optimal strategy before the reform. Through accumulated experience, they learned to price just below round numbers. This cuts both ways: heuristic learning can be effective in stable environments, but the same firms struggled after the reform because they did not understand *why* 99-endings were optimal and could not immediately identify the new optimum.

4.4 Why exploration matters

The results highlight the importance of exploration for heuristic-based learners. Before the reform, firms had learned that 99-endings outperformed alternatives but rarely used 90- and 00-endings, so they had little information about how 90 compared to 00.

When the reform forced firms to abandon 99-endings, they faced a choice with limited information. In some cases they chose 90 and stuck with it; in others they chose 00 and eventually learned it was inferior. The asymmetry in price-ending paths, that 90 is absorbing while 00 is not, reflects firms learning from experience.

The forced exploration appears to have been beneficial: by the end of the sample, forgone profits were lower than before the reform. The pre-reform steady state, while reasonably good, was not fully optimal. The reform, by forcing firms to reconsider their pricing, led to better outcomes.

4.5 Where did learning occur?

The heterogeneity results shed light on the locus of learning. Product-level patterns suggest firms had some product-specific knowledge: products with more 99-endings before the reform converged to 90-endings faster afterward. But this knowledge was incomplete as many firms initially rounded 99 to 00 rather than to 90.

The spatial competition results suggest learning occurred primarily at the chain level. Stores in more competitive areas were not faster to learn; if anything, they were slower. The discrete policy changes visible in Figure 3, where entire chains shifted away from 00-endings at particular moments, are more consistent with centralized decision-making than with store-level adaptation (consistent with DellaVigna and Gentzkow, 2019; Hitsch et al., 2017; Tabanakov et al., 2024).

4.6 Summary

The patterns are best explained by heuristic-based pricing. Firms learned that 99-endings work well without understanding the demand structure that makes them optimal. When the reform removed 99-endings, firms could not immediately identify the new optimum and had to learn through experience. This process took about a year and involved costly initial mistakes.

The key implication: observing near-optimal behavior does not imply firms are optimizing a known objective function. Positive economics may describe steady-state behavior reasonably well, but deviations may persist and when firms are shocked out of equilibrium, behavior will deviate from optimization-based predictions.

5 Discussion

This paper studies how supermarket chains adjusted prices following a reform that restricted available price endings. The evidence suggests firms arrived at their pricing strategies through adjusting heuristics rather than by learning demand and optimizing. When the reform removed 99-endings, many firms initially chose dominated 00-endings before gradually learning to shy away from them. The adjustment took about a year and involved meaningful profit losses.

5.1 Implications for research

The results offer partial support for positive economics. In steady state, firms behaved approximately as if optimizing against left-digit biased demand. A researcher observing only pre-reform behavior might reasonably conclude that firms understood demand and priced accordingly. On the other hand, in order to better explain firms' behavior as optimizers we need to allow for incorrect beliefs about fundamentals. Future work can explore how to use normative models while allowing for deviations from the full optimum.

The "as if" assumption becomes more problematic when firms face changes. Because firms did not actually know demand, they could not immediately adapt. Researchers studying firm responses to policy changes, entry and exit, mergers and acquisitions, or demand shocks should allow for adjustment periods where behavior deviates from optimization-based predictions.

Further, even small steady-state deviations from optimum, sustained by heuristics, can cause large deviations in constructs like welfare or markups. A transition that shifts heuristics can affect not only optimal actions but also actual conduct before and after the change, potentially flipping the sign of counterfactual conclusions. Future work can incorporate learning or implementation costs to pricing strategies in order to derive better descriptive counterfactuals after transitions.

5.2 Implications for practice

The paper highlights the value of understanding demand. Heuristic-based pricing can produce good outcomes in stable environments, but creates fragility. Firms that rely on "what has worked" may not know *why* it works, leaving them vulnerable when conditions change. Indeed, the reform caused a spike in forgone profits as firms adjusted, and similar dynamics might occur more broadly following other changes: new competitors, shifting preferences, regulatory changes, or supply disruptions. Therefore, firms that understand their demand functions can adapt more quickly.

In addition, deliberate experimentation can also lower losses even under heuristic pricing. Testing alternatives even when current strategies seem adequate, can reveal whether current practices are truly optimal or merely good enough. Companies that never explore a specific margin can be arbitrarily wrong (Hanna et al., 2014; Schwartzstein, 2014; Gagnon-Bartsch et al., 2023; Dubé and Misra, 2023; List et al., 2023). The reform forced such experimentation; firms might benefit from inducing it voluntarily.

Finally, one might expect competitive pressure to discipline pricing mistakes by accelerating learning. The spatial competition analysis finds limited support for this at the store level. In settings where pricing is centralized, local competition may provide noisy signals that are slow to aggregate. Firms may wish to accumulate signals to reduce costly mistakes before triggering a policy change.

5.3 Limitations

Several limitations apply. First, this is a single case study in one industry and country. The specific patterns depend on the particular reform and demand structure. Whether similar dynamics occur in other settings is an empirical question.

Second, the magnitudes depend on assumptions. The forgone profit estimates rely on structural assumptions about demand and costs. The qualitative pattern of baseline deviation, spike at reform, gradual improvement is more robust than specific magnitudes.

Third, one might construct models that rationalize the data without invoking heuristics. While the demand analysis showing discontinuities immediately after the reform rejects that consumer left-digit bias vanished around the reform, maybe there were unobserved costs to 90-endings such that 00-endings might have been rational initially. These alternatives cannot be definitively ruled out, though they require additional assumptions to explain the eventual return to 90-endings and the firm-level patterns.

Finally, some may object that price endings are too minor to matter, making this case unrepresentative of firm decision-making more broadly. Two responses: First, the evidence suggests they do matter—5–9% excess demand losses from 00-endings, they matter across contexts and settings (List et al., 2023; Strulov-Shlain, 2023), and the CEO of Israel's second-largest chain, Mega, was quoted "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." Second, the perception that price endings are unimportant is itself consistent with the heuristic interpretation: firms may not know what they are leaving on the table.

5.4 Broader perspective

The finding that firms use heuristics rather than optimizing known objectives is not new. Surveys and interviews of businesspeople over the decades support heuristic-based pricing (Bewley, 2025; Haynes, 1964), and a growing literature documents persistent deviations from profit maximization (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). What this paper adds is evidence for a specific mechanism of heuristic-based learning that finds reasonable solutions without understanding why they work, and a demonstration of how this mechanism generates both steady-state partial-optimality and transition-period mistakes.

This provides a coherent answer to a recurring question: how can sophisticated firms persistently fail to fully optimize? The answer is that they may not know the structure of the profit function. Their objective can be to maximize profits, but they learn what works through local experience that does not generalize and can stop short of full optimization if their heuristic is not sophisticated enough. When conditions change, they must learn again.

Some readers may find this conclusion obvious; others may find it hard to accept.[25] Both reactions have merit. Firms do try to maximize profits, and the optimization framework provides useful discipline. But the gap between trying and succeeding can be substantial, especially for more inert decision makers and during periods of change.

5.5 Conclusion

Learning to set prices is important but complicated. This paper provides evidence that firms can find pricing strategies that resemble optimal behavior through heuristics, without knowing the demand function that makes those strategies optimal. This approach can work better in stable environments but creates vulnerability to change.

The broader lesson is that observing near-optimal behavior does not imply knowledge of the optimum. Firms may move toward the optimum through trial and error, leaving them unprepared

[25]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."

for new situations or far from optimum if they do not explore enough. For researchers, this suggests caution in assuming full optimization in steady-state and immediate optimization in response to changes. If conclusions are sensitive to deviations from the assumed optimum, communicating that sensitivity is warranted. For practitioners, it highlights the value of exploration in order to make heuristics "work", and of understanding *why* current strategies work, not just *that* they work.

References

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.

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.

Bewley, Truman F, *Price setting*, Oxford, England: Polity Press, May 2025.

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.

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.

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.

Gagnon-Bartsch, T, M Rabin, and J Schwartzstein, "Channeled attention and stable errors," *mimeo*, 2023.

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.

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.

Haynes, W Warren, "Pricing practices in small firms," *Southern economic journal*, April 1964, *30* (4), 315.

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, 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.

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*, November 2023, *90* (6), 3186–3237.

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.

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.

Tabanakov, Semyon, Ali Goli, and Pradeep K Chintagunta, "Retail pricing and organizational structure," *Kilts Center at Chicago Booth Marketing Data Center Paper*, 2024.

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.

Online Appendix, Not For Publication

Learning and Limitations of Heuristic Pricing: 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 (I 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

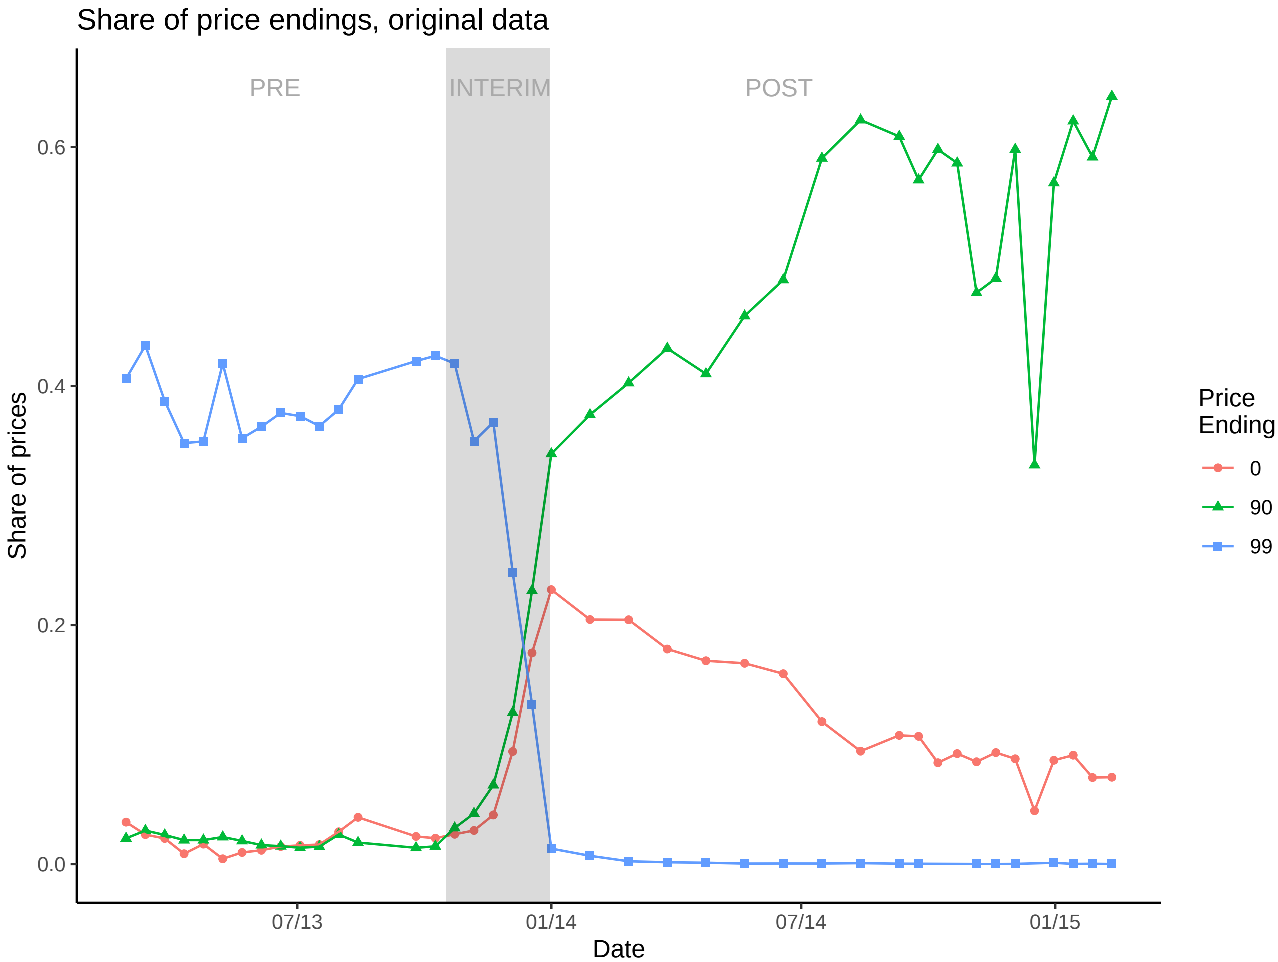

The figure shows shares of 99-, 90-, and 00-ending prices over time using the raw ICC data without sample selection or price imputation. The patterns are qualitatively similar to those in Figure 1a.

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 dataset is a scanner data recording store-product-period level data on revenue and number of units sold. It also includes 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.

The data includes some measurement errors. 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[26]. 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)[27]. To be confident that the weekly price represents the true price that consumers paid, I keep observations where the price is indeed admissible absent rounding (34.6% of prices), or if it is also of a same-price spell of at least

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

[27]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).

three weeks (35.8%), keeping a total of 47.26% of the popular product-store observations[28]. 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 Figures

B.1 Price-Ending Path Classification

Figure A-2 illustrates the state-machine used to classify price-ending paths at the product-store level, as described in Section 3.

Figure A-2: Automaton tracking price endings path

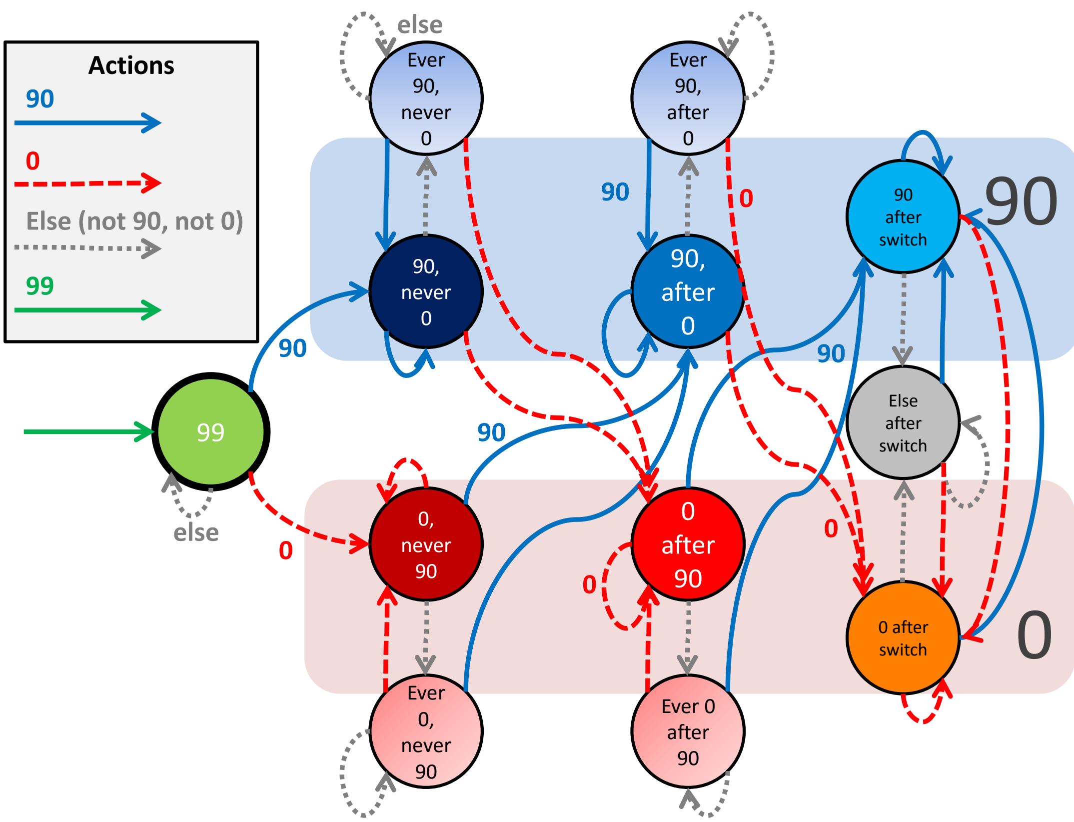

The figure illustrates a state-machine that assigns each observation into a "price-ending path" as a function of 4 possible price endings: 99, 90, 00, or other. For example, a price sequence of {4.99, 4.90, 5.00, 5.90} will have the corresponding states {99, 90 never 0, 0 after 90, 90 after switch}.

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

B.2 CBS Replication of Product Heterogeneity

Figure A-3 replicates the product-level heterogeneity analysis from Figure 4 using CBS data, where each observation is a product-type rather than a specific UPC.

Figure A-3: Correlation between pre-shares of 99 and post-shares of 00: CBS data replication.

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

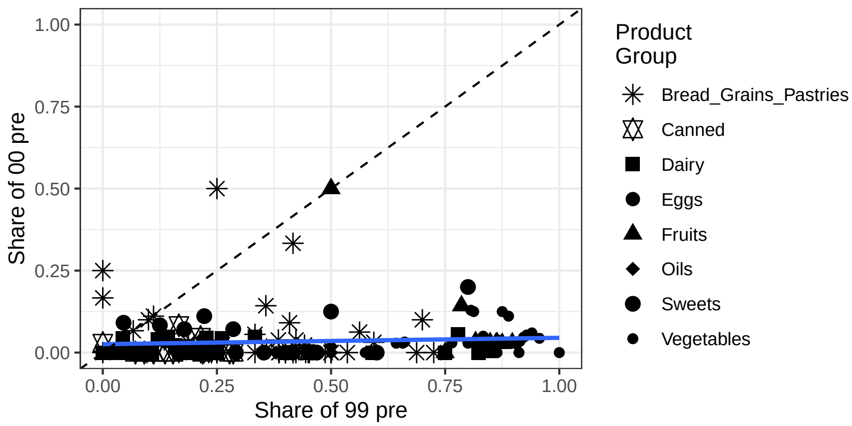

(b) Correlation between share of 99 pre and 00 immediately after policy change (1-2/2014)

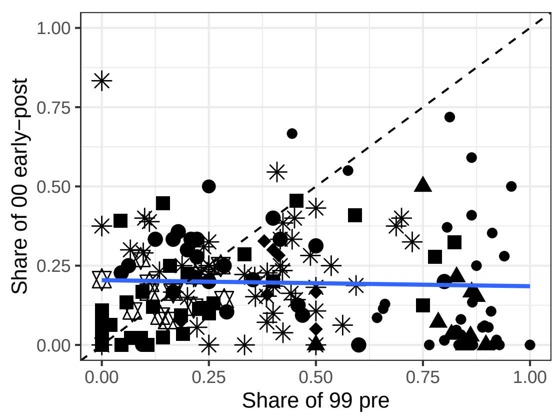

(c) Correlation between share of 99 pre and 00 a year after policy change (1-2/2015)

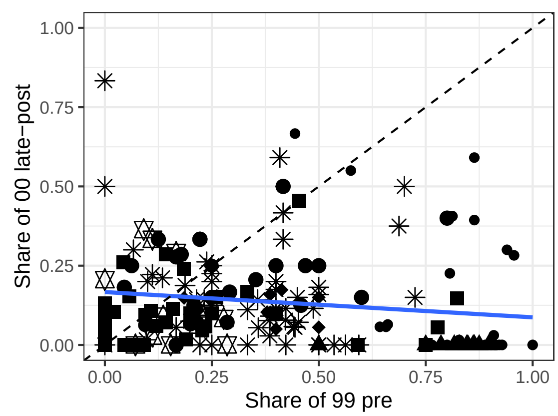

The figures show shares of 00-ending prices at different periods against 99-ending prices in the pre-period, by product-type in the CBS data. The patterns are similar to those in Figure 4, though less pronounced. Each dot represents a product-type rather than a specific UPC.

C Forgone Profits Computation

This appendix provides details on the computation of forgone profits reported in Section 3.

C.1 Demand Model

I model demand using a two-product logit with left-digit bias. Consumer j chooses product $i \in \{1, 2, \text{outside}\}$ to maximize:

$$u_{ij} = \delta_i - \alpha \hat{p}_i + \varepsilon_{ij} \tag{4}$$

where δ_i is the product intercept, $\alpha > 0$ is price sensitivity, $\hat{p}_i = (1 - \theta)p_i + \theta(\lfloor p_i \rfloor + \Delta)$ is the perceived price incorporating left-digit bias, and ε_{ij} is Type-I extreme value. The outside option has $\delta_0 = 0$ and $p_0 = 0$.

Market shares follow the standard logit form:

$$s_i(p_1, p_2; \theta) = \frac{\exp(\delta_i - \alpha \hat{p}_i)}{1 + \exp(\delta_1 - \alpha \hat{p}_1) + \exp(\delta_2 - \alpha \hat{p}_2)} \tag{5}$$

I use logit rather than constant-elasticity (log-log) demand because it produces more conservative estimates of forgone profits. Under log-log demand, losses from suboptimal pricing are larger; qualitative patterns are similar.

C.2 Calibration

Parameters are calibrated as follows:

- **Product intercepts**: $\delta_1 = 1$, $\delta_2 = 0.5$. These are normalized; the key is that products differ in baseline appeal.
- **Price sensitivity**: $\alpha = 0.5$. This is calibrated so the average price elasticity across simulated prices matches the estimated elasticity of approximately -2.5 from Section 3.
- **Left-digit bias**: $\theta = 0.22$. This is the demand-side estimate from Section 3.
- **Focal ending**: $\Delta = 0$.
- **Cost distribution**: Costs are drawn from a Weibull distribution with shape parameter $k = 4$ and scale parameter $\lambda = 6$. This distribution is chosen so that the resulting price distribution (under optimal pricing) roughly matches the observed price range and distribution shape.

C.3 Simulation Procedure

Cost draws. I draw 1,000 costs for product 1 from the Weibull distribution. For each product-1 cost, I independently draw 50 costs for product 2, yielding 50,000 (cost_1, cost_2) pairs.

Optimal pricing. For each cost pair, I solve for the prices (p_1^*, p_2^*) that maximize total profits given both prices:

$$p_i^* = \arg\max_{p_i, p_j} \; s_i(p_i, p_j; \theta) \cdot (p_i - c_i) + s_j(p_i, p_j; \theta) \cdot (p_j - c_j) \tag{6}$$

The solution is computed twice for each cost pair:

1. **Pre-reform**: prices can be set to the cent (e.g., 2.99, 3.00, 3.01, ...)
2. **Post-reform**: prices restricted to dimes (e.g., 2.90, 3.00, 3.10, ...)

Profits under optimal pricing. For each cost draw i:

$$\Pi_i^* = s_1(p_1^*, p_2^*; \theta) \cdot (p_1^* - c_{1i}) + s_2(p_1^*, p_2^*; \theta) \cdot (p_2^* - c_{2i}) \tag{7}$$

C.4 Rule-of-Thumb Heuristic

The descriptive model of actual pricing is a "rule-of-thumb" heuristic with three parameters estimated each month: $\omega_t = (\omega_t^{99/90}, \omega_t^{00}, r_t)$.

Step 1: Naive price. Given cost c and price elasticity ε (computed from the logit model at average prices), compute the naive optimal price ignoring left-digit bias:

$$p^{\text{naive}} = \frac{c \cdot \varepsilon}{1 + \varepsilon} \tag{8}$$

This is Lerner's rule for a firm facing constant-elasticity demand.

Step 2: Rounding decision. Let $p^{\text{round}} = \text{round}(p^{\text{naive}})$ be the nearest integer to the naive price (higher or lower). If $|p^{\text{naive}} - p^{\text{round}}| \leq r_t$ (the naive price is "close" to a round number):

- With probability $\omega_t^{99/90}$: set price to $p^{\text{round}} - 0.01$ (pre-reform) or $p^{\text{round}} - 0.10$ (post-reform)
- With probability ω_t^{00}: set price to p^{round}
- With probability $1 - \omega_t^{99/90} - \omega_t^{00}$: keep p^{naive}

If the naive price is not close to a round number, keep p^{naive}.

Estimating ω_t. For each month t, I estimate $(\omega_t^{99/90}, \omega_t^{00}, r_t)$ by matching the model-implied price distribution to the observed distribution. The procedure is:

1. Take the empirical shares of 99/90-endings and 00-endings in month t
2. For candidate parameter values, simulate the heuristic pricing for each cost draw
3. Compute the implied shares of each price ending
4. Choose parameters to minimize the distance between implied and empirical shares

This procedure assumes the cost distribution is stable over time (estimated once from pre-reform data using the method in Appendix D) while the heuristic parameters ω_t vary month-by-month.

C.5 Computing Forgone Profits

For each month t and cost draw i:

1. Compute optimal prices (p^*_{1i}, p^*_{2i}) and optimal profit Π^*_i
2. Apply the heuristic with estimated $\hat{\omega}_t$ to get rule-of-thumb prices (p^r_{1i}, p^r_{2i})
3. Compute rule-of-thumb profit: $\Pi^r_i = s_1(p^r_{1i}, p^r_{2i}; \theta) \cdot (p^r_{1i} - c_{1i}) + s_2(\cdot) \cdot (p^r_{2i} - c_{2i})$

Forgone profits in month t:

$$\text{Forgone}_t = 1 - \frac{\sum_i \Pi^r_i}{\sum_i \Pi^*_i} \tag{9}$$

This measures the percentage of potential profits lost due to heuristic (rather than optimal) pricing.

D Perceived Left-Digit Bias Estimation

This appendix provides a detailed description of the Method of Simulated Moments (MSM) procedure used to estimate the "perceived" left-digit bias parameter θ implied by observed pricing behavior. The goal is to find the value of θ that, if firms were optimizing against left-digit biased demand with that belief, would best rationalize the observed distribution of prices. This procedure follows Strulov-Shlain (2023).

D.1 Setup and Notation

Model. Firms face constant-elasticity demand with left-digit bias. The perceived price is:

$$\hat{p}(p; \theta) = (1 - \theta)p + \theta(\lfloor p \rfloor + \Delta) \tag{10}$$

where p is the actual price, $\lfloor p \rfloor$ is the floor (largest integer less than or equal to p), Δ is the focal price ending (set to 0 throughout), and $\theta \in [0,1]$ is the left-digit bias parameter. Demand is $D(p;\theta) = A\hat{p}(p;\theta)^{\epsilon}$ where $\epsilon < -1$ is the price elasticity.

Firm optimization. A profit-maximizing firm with marginal cost c solves:

$$\max_{p} \Pi(p;\theta,c) = D(p;\theta) \cdot (p - c) \tag{11}$$

The first-order condition for an interior optimum (away from the discontinuity at round numbers) yields the relationship between optimal price and cost:

$$c(p;\theta,\epsilon) = p + \frac{\hat{p}(p;\theta)}{\epsilon(1-\theta)} = p\frac{1+\epsilon}{\epsilon} + \frac{\theta}{1-\theta}\frac{\lfloor p \rfloor + \Delta}{\epsilon} \tag{12}$$

This equation maps each non-just-below price p to the unique cost c for which p is optimal (given θ and ϵ).

Key implication: bunching and missing prices. Under left-digit bias ($\theta > 0$), firms should bunch at "just-below" prices (e.g., X.99) and avoid prices just above round numbers. The *Next-Lowest Price* P_q is the lowest price above a just-below price q that is optimal for some cost. For any just-below price q (e.g., 2.99), there exists a threshold $P_q > \lfloor q \rfloor + 1$ such that:

- Prices in the interval $(\lfloor q \rfloor + 1, P_q)$ are never optimal for any cost (these are "missing" prices)
- The just-below price q is optimal for costs in some interval $[\underline{c}_q, \bar{c}_q]$
- Prices at or above P_q become optimal for sufficiently high costs

The Next-Lowest Price P_q is found by solving for the price $P > \lfloor q \rfloor + 1$ such that there exists a cost c where the firm is indifferent between q and P, and P satisfies the first-order condition at that cost. This is given implicitly by Equation (3) in the main text.

D.2 Data and Moments

Price data. The estimation uses the empirical distribution of prices from the ICC dataset. Let $\{p_1, p_2, \ldots, p_N\}$ denote the observed prices (one observation per product-store-period).

Price grid. Define a price grid $\mathcal{P} = \{\rho_1, \ldots, \rho_K\}$ covering the relevant range. For example, $\mathcal{P} = \{2.00, 2.10, 2.20, \ldots, 3.90, 4.00\}$ at the cent level. In practice, I use prices from the data rounded to the nearest cent.

Empirical moments. For each price $\rho_k \in \mathcal{P}$, compute the empirical share:

$$S_p^{\text{data}} = \frac{\sum_i \{i : p_i \in [\rho_k, \rho_{k+1})\}}{N} \tag{13}$$

These empirical shares are the moments to be matched.

D.3 Estimation Algorithm

The algorithm proceeds in two stages: (1) estimate the shape of the underlying price/cost distribution using prices unaffected by bunching, and (2) use the model to predict price shares as a function of (θ, ϵ) and find the parameters that best match the data.

D.3.1 Stage 1: Estimate the Price Distribution Shape

Identify unaffected prices. Under the model, prices ending in X.99 exhibit excess mass (bunching), and prices just above round numbers (e.g., X.00 to X.28) are missing. Prices in the middle of each dollar range (e.g., X.29 to X.89) should be unaffected by left-digit bias and reflect the underlying cost distribution.

Fit a smooth distribution. Using only the unaffected prices, fit a parametric distribution to the empirical CDF. I use a logistic polynomial:

$$\hat{F}_p(p) = \frac{1}{1 + \exp(-(\beta_0 + \beta_1 p + \beta_2 p^2 + \cdots))} \tag{14}$$

Specifically:

1. Select prices in the "unaffected" range: for each dollar interval $[n, n+1)$, use prices from $n.29$ to $n.89$ (or a similar range that excludes bunching at X.99 and missing prices near X.00).

2. Compute the empirical CDF at these prices.

3. Regress the logit of the empirical CDF on a polynomial in price: $\text{logit}(\hat{F}_p^{\text{emp}}) = \beta_0 + \beta_1 p + \beta_2 p^2 + \ldots$

4. Use the fitted coefficients to define $\hat{F}_p(p)$ for all prices in $\mathcal{P}$.

This gives a smooth estimate of what the price distribution *would* look like in the absence of left-digit bias.

D.3.2 Stage 2: Predict Price Shares Given (θ, ϵ)

For each candidate parameter pair (θ, ϵ):

Step 2.1: Identify just-below prices and compute Next-Lowest Prices. For each just-below price q in the data (e.g., $q \in \{1.99, 2.99, 3.99, \ldots\}$):

1. Solve for the Next-Lowest Price P_q using the indifference condition. Specifically, find $P_q > \lfloor q \rfloor + 1$ such that:
$$P_q + \frac{\hat{p}(P_q;\theta)}{\epsilon(1-\theta)} = \frac{\hat{p}(P_q;\theta)^{\epsilon} \cdot P_q - \hat{p}(q;\theta)^{\epsilon} \cdot q}{\hat{p}(P_q;\theta)^{\epsilon} - \hat{p}(q;\theta)^{\epsilon}} \tag{15}$$
This is Equation (3) in the main text. Solve numerically (e.g., using bisection or Newton's method).

2. Record P_q. For example, if $\theta = 0.03$ and $\epsilon = -3$, solving for $q = 2.99$ might yield $P_{2.99} = 3.27$.

Step 2.2: Construct the predicted price grid. The model predicts three types of prices:

1. **Just-below prices** $\{q_1, q_2, \ldots\}$: these have excess mass due to bunching
2. **Missing prices**: prices in $(\lfloor q_i \rfloor + 1, P_{q_i})$ for each q_i should have zero mass
3. **Regular prices**: prices at or above P_{q_i} (up to the next just-below price) follow the smooth distribution

Construct the augmented price grid $\mathcal{P}'$ that includes:

- All just-below prices $\{q_i\}$
- All Next-Lowest Prices $\{P_{q_i}\}$
- All regular prices between P_{q_i} and q_{i+1}

Step 2.3: Map prices to costs. For each price $p \in \mathcal{P}'$, compute the cost $\underline{c}_p$ such that p is the profit-maximizing price:
$$\underline{c}_p = p\frac{1+\epsilon}{\epsilon} + \frac{\theta}{1-\theta}\frac{\lfloor p \rfloor + \Delta}{\epsilon} \tag{16}$$

Step 2.4: Recover the cost distribution. The cost distribution is derived from the price distribution using the change-of-variables implied by the first-order condition. Since $\hat{F}_p$ gives the CDF of prices (from Stage 1), and the FOC maps prices to costs, we have:
$$\hat{F}_c(c) = \hat{F}_p\left(p(c)\right) \tag{17}$$

where $p(c)$ inverts the FOC. Substituting:

$$\hat{F}_c(\underline{c}_p) = \hat{F}_p\left(\underline{c}_p \cdot \frac{\epsilon}{1+\epsilon} - \frac{\theta}{1-\theta}\frac{\lfloor p \rfloor + \Delta}{1+\epsilon}\right) \tag{18}$$

Step 2.5: Compute predicted price shares. For each price p_i in the augmented grid $\mathcal{P}'$:

$$\hat{S}_{p_i}(\theta, \epsilon) = \hat{F}_c(\underline{c}_{p_{i+1}}) - \hat{F}_c(\underline{c}_{p_i}) \tag{19}$$

This gives the predicted share of prices at p_i: it equals the mass of costs for which p_i is optimal.

For **just-below prices** q: the predicted share includes all costs in the bunching region:

$$\hat{S}_q(\theta, \epsilon) = \hat{F}_c(\underline{c}_{P_q}) - \hat{F}_c(\underline{c}_q) \tag{20}$$

where $\underline{c}_q$ is the lowest cost for which q is optimal, and $\underline{c}_{P_q}$ corresponds to the Next-Lowest Price.

For **missing prices** (between $\lfloor q \rfloor + 1$ and P_q): set $\hat{S}_p = 0$.

Step 2.6: Aggregate to the empirical price grid. The augmented grid $\mathcal{P}'$ may have prices (like $P_q = 3.27$) that don't correspond to actual price points in the data. Aggregate predicted shares to match the empirical grid $\mathcal{P}$:

- Assign the share at P_q to the nearest price ending in 9 below it. For example, if $P_{2.99} = 3.27$, assign its share to 3.19 (or the next admissible 9-ending price).
- Sum shares for prices that map to the same grid point.

D.3.3 Stage 3: Minimize Distance

Objective function. Choose (θ, ϵ) to minimize the sum of squared differences between predicted and empirical shares:

$$(\hat{\theta}, \hat{\epsilon}) = \arg\min_{\theta, \epsilon} \sum_{p \in \mathcal{P}^*} \left(\hat{S}_p(\theta, \epsilon) - S_p^{\text{data}}\right)^2 \tag{21}$$

where $\mathcal{P}^*$ excludes the prices used to fit the distribution in Stage 1 (to avoid using the same data twice).

Optimization. Use grid search or numerical optimization over (θ, ϵ). Typical ranges: $\theta \in [0.001, 0.5]$, $\epsilon \in [-7, -1.5]$.

D.4 Implementation Details

Focal ending. Set $\Delta = 0$ throughout (the focal point is like ignoring the end digits. In practice the perceived parameters are largely insensitive to this choice).

Polynomial degree. A cubic or quartic polynomial typically suffices for $\hat{F}_p$.

Price range. Focus on prices in a range where the density is reasonably high (e.g., \$2–\$15 for supermarket products). Exclude outliers.

Pre- vs. post-reform. For pre-reform estimation, just-below prices end in .99 and missing prices are .00 to approximately .28 (depending on θ). For post-reform, just-below prices end in .90 and missing prices are .00 to approximately .09.

Elasticity. In some specifications, ϵ is fixed at an externally estimated value (e.g., from demand estimation) and only θ is estimated. In others, both are estimated jointly.

D.5 Interpretation

The estimated $\hat{\theta}$ represents the level of left-digit bias that *firms behave as if they believe consumers have*. This "perceived" bias can differ from the true bias estimated from demand data:

- If $\hat{\theta} < \theta^{\text{true}}$, firms underestimate left-digit bias and price as if the demand discontinuity is smaller than it actually is.
- If $\hat{\theta} > \theta^{\text{true}}$, firms overestimate the bias.

In this paper, I find $\hat{\theta} \approx 0.02$–$0.04$ from pre-reform pricing, while demand-side estimates suggest $\theta^{\text{true}} \approx 0.22$–$0.36$. Firms substantially underestimate left-digit bias, but their perceived bias is still positive. If we wish to keep the model with incorrect beliefs as a good positive description of how they price, we would interpret their actions as if they recognize that just-below pricing works, but not how well.