Review of Economic Studies (2023) **00**, 1–34 doi:10.1093/restud/rdac082

Advance access publication 19 December 2022

More Than a Penny's Worth: Left-Digit Bias and Firm Pricing

AVNER STRULOV-SHLAIN
Booth School of Business, University of Chicago

First version received December 2022*; Accepted December* 2022 (Eds.)

Firms arguably price at ninety-nine-ending prices because of left-digit bias—the tendency of consumers to perceive a \$4.99 as much lower than a \$5.00. Analysis of retail scanner data on 3500 products sold by twenty-five U.S. chains provides robust support for this explanation. I structurally estimate the magnitude of left-digit bias and find that consumers respond to a one-cent increase from a ninety-nine-ending price as if it were more than a twenty-cent increase. Next, I solve a portable model of optimal pricing given left-digit biased demand. I use this model and other pricing procedures to estimate the level of left-digit bias retailers perceive when making their pricing decisions. While all retailers respond to left-digit bias by using ninety-nine-ending prices, their behaviour is consistently at odds with the demand they face. Firms price as if the bias were much smaller than it is, and their pricing is more consistent with heuristics and rule-of-thumb than with optimization given the structure of demand. I calculate that retailers forgo 1–4% of potential gross profits due to this coarse response to left-digit bias.

Key words: Left-digit bias, Pricing, Behavioural firms

JEL Codes: D90, D12, D22, L11, L20, L81

1. INTRODUCTION

Companies often interact with consumers who are affected by various biases and heuristics. In such situations, companies need to take these behaviours into account and decide how to respond. However, these behaviours are often not appreciated or formalized by researchers nor firms, which can lead to sub-optimal responses. This paper studies one such leading example.

A common practice by firms, documented for at least eighty years (Ginzberg, 1936), is setting prices that end with ninety-nine.[1] A leading explanation is that these prices are set as a response to left-digit biased consumers. Left-digit biased consumers are insensitive to the cents component of the price, and therefore demand is relatively inelastic when the left-digits do not change, but very elastic when they do. In turn, firms are correct to price at ninety-nine-ending prices.

Yet, as much as left-digit bias had been studied in the past, we know very little about its implications on firm behaviour and welfare—in theory and in practice. For a start, we do not have an understanding of left-digit bias' magnitude, *i.e.*, what level of left-digit bias, if any, do

1. See also Aalto-Setala (2005), Anderson *et al.* (2015), Anderson and Simester (2003), Ashton (2014), Ater and Gerlitz (2017), Conlon and Rao (2016), El Sehity *et al.* (2005), Kalyanam and Shively (1998), Levy *et al.* (2011), Macé (2012), Schindler and Kibarian (1996), Strulov-Shlain (2019a), Snir *et al.* (2017), and Stiving and Winer (1997).

The editor in charge of this paper was Nicola Gennaioli.

consumers exhibit in their daily purchase decisions. This lack of quantification results in the common wisdom of "ninety-nine-ending prices might be a good idea", and not much beyond that. Furthermore, we know relatively little about how firms price in practice and how that compares to the optimum; we often assume they optimize in a particular way. Perhaps firms converged to optimal behaviour by decades of trial and error or internal research; alternatively, they might consistently deviate in specific ways.

The goal of this paper is to study firms' response to left-digit bias in three steps: estimating consumers' bias (demand), quantifying firms' pricing procedures (supply), and studying the implications of the supply response given demand. First, I parameterize left-digit bias into consumers' price perception, which allows me to incorporate left-digit bias into demand and estimate the magnitude of left-digit bias. Second, I solve for optimal pricing and heuristic pricing given left-digit bias. The prices set by firms reveal their beliefs about left-digit bias and the pricing procedures they use to set prices. Third, the first two steps allow me to compare how firms price in practice to how they should have priced, and study implications on welfare.

I use data from NielsenIQ (2020) on twenty-five retailers. While they all use ninety-nine-ending prices, their behaviour is consistently at odds with the structure of demand since they frequently use low-ending prices. Firms price as if the bias were much smaller than it is, and their pricing is more consistent with heuristics and rule-of-thumb than with optimization given the structure of demand. I estimate roughly 1–4% of forgone profits from pricing with underappreciation of the bias and ambiguous effects on consumer surplus. The persistent underreaction and mischaracterization of the bias, and the costliness of this deviation—similar to uniform pricing (DellaVigna and Gentzkow, 2019) and costlier than non-seasonal pricing (Butters *et al.*, 2019)—is the main finding of the paper.

In order to study the effects on firms, I first define demand with left-digit biased consumers. As in DellaVigna (2009) and Lacetera *et al.* (2012), left-digit bias is defined as a distortion in the perception of numbers. In the setting of grocery prices which are almost entirely under \$10 (*i.e.*, a single dollar digit), the bias can be thought of as excess weight on just the dollar digit of the price versus the exact price. A price of \$4.99 is perceived as a mix between \$4.99 and "\$4 something". This distortion causes a dampened perception of an only-cents change, but over-perception of a left-digit change. The bias is modelled with a single parameter θ ranging between 0 (no bias) and 1 (full bias).[2] For example, with left-digit bias of $\theta = 0.2$, the difference between \$4.99 and \$5.00 is perceived as more than twenty cents, while the difference between \$5.00 and \$5.01 is perceived as 0.8 of a cent. In turn, this perception distortion translates to less elastic demand when only cents change, and more elastic demand when the left-digit changes.

I find strong left-digit bias in demand data. I analyse aggregate demand from retail data using a panel of sales data of almost 3500 popular products (UPCs, or a Universal Product Code) sold by twenty-five national U.S. supermarket chains. I estimate flexible models to find the residuals of quantity sold, *i.e.*, unexplained demand, by price. These residuals exhibit a sawtooth pattern: as prices increase, residual demand is increasing within each dollar-digit and then dropping when the dollar digit changes. The sawtooth pattern is consistent with the above-mentioned distortion of left-digit bias. I then quantify the level of left-digit bias that explains these patterns and find that it is about 0.2 on average.

2. For more digits, we can think of different weights for each digit, as in Lacetera *et al.* (2012). For simplicity sake, and because it is inconsequential given the empirical setting, I keep the one parameter formulation. I discuss the relative weight of dimes versus cents when appropriate.

Left-digit bias affects optimal pricing because it effectively creates discontinuities in demand, which lead to a bunching response in pricing and missing prices. Intuitively, if demand discontinuously drops at round prices, prices that were otherwise set just above the discontinuity should be lowered to just below it. Therefore, there should be masses of prices just below round numbers and regions of missing prices with low price-endings. I solve a model of monopolistic pricing facing left-digit biased demand to quantify the response. The model prescribes levels of excess ninety-nine-endings and ranges of missing prices from primitives of left-digit bias, price elasticity, and marginal costs.

Firms' pricing patterns imply under-response to the bias. Prima facie, firms respond to the bias with 30–40% of prices ending with ninety-nine. Yet, the level of response that these numbers represent is far from what is predicted by the model. With left-digit bias of 0.2, and any reasonable price elasticity, all prices should end with ninety-nine. Using identification stemming from the discontinuities in demand, I estimate retailers' perceived elasticities and left-digit bias needed to rationalize the pricing behaviour. I find that while the true bias on the consumer side is 0.20, firms perceive it to be as low as 0.01–0.03. Furthermore, *all* chains underestimate the bias they face. In an alternative interpretation of correct beliefs but with partial incorporation of the bias into pricing, I find that in these cases the bias is incorporated in 30–50% of pricing instances across chains. The best-fit rationalizations of pricing behaviour are those that allow for both partial incorporation and underestimation of the bias, or further, a rule-of-thumb pricing where in 40% of cases both lower and higher prices are rounded to the nearest ninety-nine-ending prices.

Finally, I find the above-mentioned pricing procedures lead to welfare losses. I compare gross profits under optimal price setting—when elasticities and left-digit bias levels are known—to cases in which there is mischaracterization of the bias as estimated above. Under a broad set of scenarios and assumptions, I find that underappreciation of left-digit bias leads to substantial losses of gross profits. In contrast, the effect on consumer surplus is undetermined: prices are higher than if firms were responding optimally, but the sign of the average effect depends on non-testable assumptions.

The contribution of the paper is providing an internally valid and robust quantitative study of optimal and actual behaviour of companies that face biased consumers. While the paper also contributes examination of the bias across thousands of products in two dozen chains, the main contribution of the paper is leveraging the demand estimates to quantify the gap between firms' predicted response and their actual pricing response to the bias. It serves as an example of "behavioural firms" that make consistent mistakes about demand.

More broadly, the paper questions the reality of firms' optimization. Big retail chains are making a costly mistake in pricing, one of their core activities. The assumption of optimization, extensively questioned for individuals in past decades, should be questioned for firms as well. Rather than optimizing, firms might be using heuristics or rely on partial learning.[3] Firms respond in the right direction, but stop short of optimization.

Left-digit bias and nine-ending prices have been previously explored. Some papers focus on documenting the prevalence of prices ending with nine (Levy *et al.*, 2011; Hackl *et al.*, 2014) or reversion to nine following currency changes (Aalto-Setala, 2005; El Sehity *et al.*, 2005; Strulov-Shlain, 2019a) as evidence of firms' revealed preference for nine-ending prices due to their benefits. Experimental papers explore underlying mechanisms (*e.g.*, Carver and

3. The paper by Strulov-Shlain (2019a) analyzes chains pricing response to a reform in Israel and reaches similar results. In particular, that firms operate under partial knowledge and persistently sub-optimize.

Padgett, 2012).[4] Fewer papers examine the effects of nine-ending prices on demand, either in field experiments (*e.g.*, Anderson and Simester, 2003; Ashton, 2014),[5] or in observational data (*e.g.*, Stiving and Winer, 1997).[6] These papers find that nine-ending prices increase demand, with some mixed results (*e.g.*, Backus et al., 2019). This paper adds to the literature by providing rigorous and robust evidence for the bias and its magnitude across multiple chains, products, and years.[7]

The paper also adds to the literature in behavioural IO in two ways. First, by studying a market where firms respond to biased consumers (see Heidhues and Köszegi, 2018, for review), and second, by studying a case where the firms are also "behavioural", in the sense that they are not profit-maximizing in a systematic way.[8] The closest paper, by List *et al.* (2021), follows a similar approach to mine. They look at left-digit bias for a highly sophisticated company offering two main products, and find substantial left-digit bias which is not taken into account in pricing decisions. The magnitude of left-digit bias and effects on revenues is similar to the magnitude I calculate. They then run a field experiment corroborating the counterfactual predictions of the effect of left-digit bias on revenues.

The structure of the paper is as follows. In Section 2, I define and document left-digit bias in demand data. Section 3 investigates how firms should price and quantifies how they actually price in practice. Section 4 studies the implications of the actual pricing patterns versus the expected pricing given demand. Finally, Section 5 concludes the paper.

2. DEMAND

To understand how firms should price in response to left-digit biased consumers we first need to understand demand. I specify a model of left-digit biased consumers, showing left-digit bias' effect on demand and its manifestation in the demand data. Next, in order to incorporate the bias in firms' pricing problem, I structurally estimate the bias parameter as prescribed by the model.

2.1. *Left-digit biased consumers and demand*

This section provides a simple parameterization of left-digit biased demand. The main idea is that consumers perceive prices with a distortion captured by a parameter θ. I use a similar approach to the parametric modelling of left-digit bias by DellaVigna (2009), Lacetera *et al.* (2012), and Busse *et al.* (2013). Like in Lacetera *et al.* (2012), I assume that this distortion is a primitive of consumers' behaviour and not a function of firms' actions, an assumption I

4. See Bizer and Schindler (2005), Carver and Padgett (2012), Schindler and Wiman (1989), Schindler and Kirby (1997), Schindler and Chandrashekaran (2004), Snir *et al.* (2017), and Thomas and Morwitz (2005). A summary table is available in Carver and Padgett (2012).

5. See Anderson and Simester (2003), Ashton (2014), Bray and Harris (2006), Dalrymple and Haines Jr (1970), Dube *et al.* (2017), Ginzberg (1936), and Schindler and Kibarian (1996).

6. See Blattberg and Wisniewski (1988), Hackl *et al.* (2014), Jiang (2020), Kalyanam and Shively (1998), Macé (2012), and Stiving and Winer (1997). Observational data papers usually find higher demand for ninety-nine- or nine-ending prices by including a dummy for these. Surprisingly, these papers rarely discuss the issue of price averaging that is common across these data.

7. A few papers document that prices tend to end with nine, or ninety-nine, and ask questions on the effect of this phenomenon on price stickiness (Levy *et al.*, 2011; Anderson *et al.*, 2015; Ater and Gerlitz, 2017) and the pass-through of taxes (Conlon and Rao, 2016). This paper adds to this literature by micro-founding this stickiness as an optimal pricing response of firms.

8. See Cho and Rust (2008), Cho and Rust (2010), DellaVigna and Gentzkow (2019), Goldfarb and Xiao (2011), and Hanna *et al.* (2014).

discuss in Section 4.3.[9] Although there can be multiple sources for left-digit bias, *e.g.*, inattention or categorization, the bias parameter has an intuitive interpretation as the extra weight put on left-most digits.[10]

Consider a product with a price p between 1 and 9.99.[11] Assume that a consumer perceives the price as $\hat{p}$

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

where $\lfloor \cdot \rfloor$ is the floor operator. The perceived price is a mix of the true price, with weight $1-\theta$, and the price with the correct left digits but a focal price ending $\Delta \in [0, 1)$ with a weight θ. The distortion increases with $\theta \in [0, 1]$, the *left-digit bias* parameter. For example, if $p = 4.99$, $\theta = 0.5$, and $\Delta = 0.19$, then $\hat{p} = 0.5 \cdot 4.99 + 0.5 \cdot (4.00 + 0.19) = 4.59$.[12]

Left-digit bias creates two effects on perceived price changes: discontinuities when the left-most digits change, and attenuated perception of changes in right-most digits. First, a 0.01 price increase is perceived as being roughly $0.01+\theta$ if the left-most digit changes (*e.g.*, 2.99–3.00); second, a 0.01 price increase is perceived to be lower than in the standard model, being $0.01(1-\theta)$, if only one digit changes (*e.g.*, 2.95–2.96). In turn, these price distortions strongly affect demand.

The focal price ending parameter, Δ, does not affect pricing but has first-order effects on welfare. Δ almost does not affect pricing because θ dominates these effects via the discontinuities. In contrast, Δ causes all prices to be perceived as lower ($\Delta = 0 \Rightarrow \hat{p} \leq p \forall p$) or larger (if $\Delta = 0.99$) than they are, hence shifting demand—causing consumers to over- or under-consume.

If a consumer has utility of the form $U(y, q) = y + (q/A)^{1+1/\epsilon}/(1+1/\epsilon)$ where y is the residual income, q is the quantity purchased of a good, and A is translating quantity to numeraire value, then overall demand will be of the form[13]

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

The effects of left-digit bias on demand lead to an intuitive visual and empirical test. Figure 1(a) illustrates constant-elasticity demand curves with and without the bias, *i.e.*, $\log Q = A + \epsilon \log(\hat{p})$, for $\theta = 0$ on the left panels and $\theta > 0$ on the right. On the left, without bias, demand is smooth. On the right, the slope is flatter within each dollar digit, and demand discontinuously drops when the dollar digit changes (highlighted by different shades). The dashed line is the fitted line from a log(quantity) on log(price) regression. Figure 1(b) shows the residuals by price from these regressions. If there is no bias, as on the left panel of Figure 1(b), residuals are a flat

9. This assumption is the main difference between this model and theoretical treatments of similar price perception biases (*e.g.*, Basu, 2006; Matějka, 2015)

10. In other papers, θ is termed "inattention" (see Gabaix, 2019). However, other forces might lead to observationally equivalent behaviour. Such mechanisms are imperfect price recall, price categorization, or a tendency to choose round numbers to represent internal values or reference-prices. Therefore, I use the term left-digit bias which captures all of these, and remain agnostic about the behavioural mechanism.

11. While the model can be extended to include more distortions, the empirical setting includes mostly prices that are below \$10. It also makes the exposition simpler. I discuss extensions in the final section of the paper.

12. An alternative interpretation is that a θ share of consumers replace the lower digits with the focal price ending. Though different in interpretation, these two models lead to similar qualitative predictions (Supplementary Material, Appendix B, solves for a pure price competition model, providing an example for when these two models differ). The former interpretation is a bit simpler to solve algebraically, and also captures the observation that "4.99 feels lower than 5.00" at the person-specific level. Furthermore, Morrison and Taubinsky (2019) find evidence against the pure multiple types interpretation.

13. Conversely, if there is type mixture then following the same formulation will lead to demand of the form $D(p; \theta, \Delta) = A[(1-\theta)p^{\epsilon} + \theta(\lfloor p \rfloor + \Delta)^{\epsilon}]$.

line at zero (if the fitted model is misspecified, it will exhibit some other general shape). On the right, left-digit biased demand residuals exhibit a sawtooth pattern with discrete drops when the dollar digits change. This shape is the key diagnostic of left-digit biased demand.

2.2. *Data*

This project relies on supermarket prices and purchase behaviour on a big set of products and chains, allowing a robust investigation of the bias and firms' response to it.

Data are weekly store-product quantities and prices of about 3500 consumer packaged goods sold by twenty-five supermarket chains in the U.S. from 2006 to 2019.

The exact price paid is a key variable, but it is not measured directly. The price in the data is a quantity-weighted average of all prices paid for a product in a store in a week. If the price changed mid-week, or different consumers paid different prices, the average price in the data is not a price any consumer actually paid. Therefore, while scanner data provide a broad coverage of products and a long panel, they require some data cleaning to determine which observations represent the exact price consumers paid. I discuss the main points here and elaborate on these issues in Supplementary Material, Appendix D, which goes into details about the data cleaning procedure; in addition to some standard restrictions, I focus on chains and stores that report "clean" prices. These are chains that are likely to be changing their prices on a weekly frequency and in alignment with the frequency with which NielsenIQ collects the data, and are also such chains that do not have different prices for members versus non-members, nor use personalized pricing. To operationalize this goal, I select chains and stores with (an admittedly arbitrary rule of) a maximum of 2% of prices with one, two, or three in the cents digit following the logic that these prices rarely appear on the shelf. In addition, because we need price variation to see how demand sensitivity changes when left-most digits change, I explicitly require some variation in dollar digits at the product level (this is implicitly done by other papers since products whose price does not change do not contribute to demand estimation). Since some of these issues are not discussed elsewhere, I encourage researchers interested in or concerned about exact prices in scanner data to read Supplementary Material, Appendix D, for a description of price construction procedures.

NielsenIQ Retail Scanner (RMS) and Consumer Panel (HMS) data are provided by the Kilts Center at the University of Chicago Booth School of Business. RMS records weekly UPC-store level quantity and revenues. After selecting a subset of forty-one modules used by DellaVigna and Gentzkow (2019) and thirty-four chains with the potential of reporting accurate prices in such data, I start with 5.2 billion observations from 3602 stores, covering the period of 2006–2019. Summary statistics are described in Table 1. The final sample consists of 78 million observations of 3475 products, in 1502 stores, with a total of 1.18 billion units sold for $4.3 billion. The main level of analysis is at the product-chain level, and there are 9107 such pairs. The average price is $4.84 with an inferred transitory sales frequency of 16%. 41% of prices end with ninety-nine, and 87% end with nine as the last digit. I elaborate on the variables construction process in Supplementary Material, Appendix D.

2.3. *Reduced-form evidence*

The shape of demand residuals by price provides a test for left-digit biased demand (as in Figure 1). Specifically, focusing on the shape of demand residuals—net of factors such as seasonality, store and product characteristics, promotion mix, and price elasticities—allows to aggregate over more flexible demand structures, such as different elasticities and price distributions for different products.

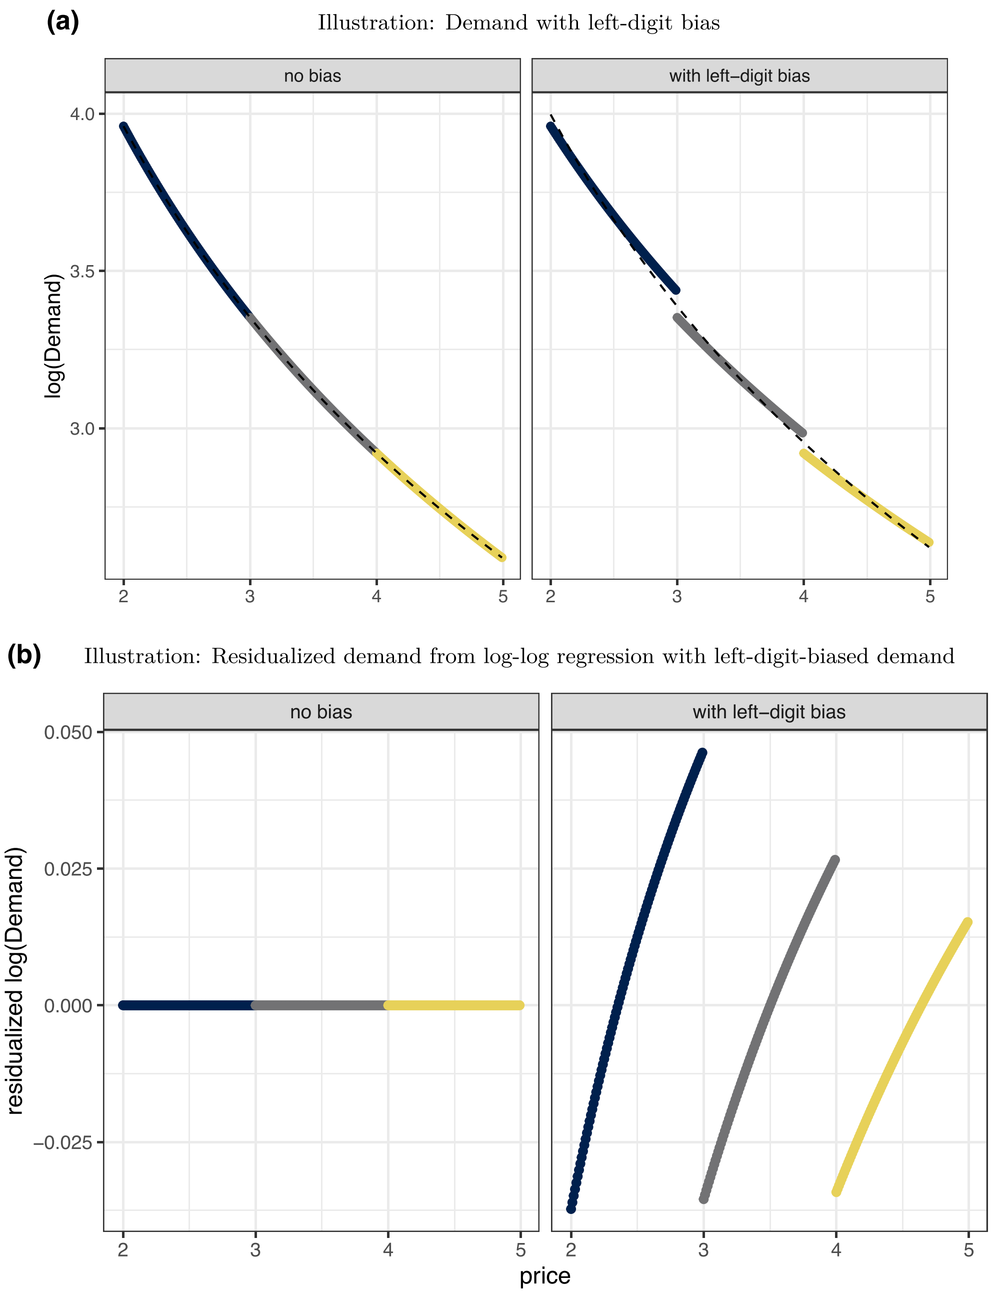

FIGURE 1

Demand and residuals with and without left-digit bias. (a) Demand with left-digit bias. (b) Residualized demand from log–log regression with left-digit-biased demand

Notes: The figures illustrate simulated demand curves under a model of left-digit bias. The top figures show simulated log(quantity)-price demand curves. Shades represent different dollar digits. The dashed line is the curve of best fit from a log(quantity) on log(price) regression. The bottom figures show the residualized demand by price—actual demand minus predicted demand. The right figures are for left-digit biased demand with $\theta = 0.16$, and the left figures are for the standard case of no bias, $\theta = 0$.

TABLE 1
Summary statistics and data selection

	Full (1)	Clean stores (2)	Clean (3)
Observations (millions)	5177.6	181.56	77.87
Products	189,722	6187	3475
Stores	3602	1953	1502
Weeks	730	730	730
Chains	34	25	24
Product-chain	979,092	18,486	9107
Total dollar sale ($billions)	147.73	10.49	4.29
Total units sold (billions)	51.19	4.06	1.18
First date	2006-01-07	2006-01-07	2006-01-07
Last date	2019-12-28	2019-12-28	2019-12-28
Average price	4.27	4.08	4.84
Share on-sale	N/A	N/A	0.16
Share ninety-nine-ending	0.32	0.33	0.41
Share nine-ending	0.69	0.78	0.87
Share zero-ending	0.16	0.12	0.07

Notes: The table presents summary statistics for the NielsenIQ scanner data. Column (1) is the initial data set of the thirty-four chains with high data quality, and stores and products with long enough presence in the data. Column (2) shows the sample after keeping product-stores with few price-endings with one, two, or three cents. Column (3) is the main sample used in the paper for the pricing analysis.

Figure 2 shows a semi-parametric demand curve. The figure shows the result of regressing log-quantity on price dummies with a rich set of fixed effects (estimation details below). The figure shows estimates of the price fixed-effects for thirty-seven "ground and whole bean coffee" products, which were found to have a similar price elasticity and hence can be plotted together on the same curve. The shape of the demand curve closely resembles that of Figure 1(a), showing drops in demand between dollar digits and an attenuated slope within each dollar digit. To examine all data, I aggregate the residuals net of product-chain level price elasticity which then allows me to combine evidence from products with different price elasticities (slopes).

Corroboration of the model—a sawtooth pattern of residuals—is shown in Figure 3 for 9107 product-chain pairs. The horizontal axis is price and the vertical axis is residualized demand, where each dot is a ten-cent bin of prices. Like the illustration in Figure 1(b), residualized demand is increasing within a dollar-digit and drops across digits. I now turn to a detailed explanation of the generating process of this figure and show that this is a robust pattern.

The test requires estimating the exact shape of demand curves, a process which raises three challenges. Like in other demand estimation exercises, we first care about the overall slope of how quantity changes in response to price changes. Second, unlike other exercises, we need exact demand by precise price since otherwise discontinuities will be blurred. Third, we also require enough variation to separate within-dollar slopes from between-dollars change in demand. These three requirements inform the estimation procedure and data construction, and are elaborated on in Supplementary Material, Appendix D. First, to estimate demand I use a now common approach of a flexible log–log estimation on a sample of clean prices (*e.g.*, Hausman, 1996; Rossi, 2014; Hitsch *et al.*, 2017; DellaVigna and Gentzkow, 2019) justified by institutional knowledge of how retail chains set prices—well in advance, with quasi-random sales, and uniformly across stores. Second, the precise price measurement is addressed in the data cleaning process. Third, obtaining enough price variation across dollar digits is satisfied by data selection, and enough variation across price-endings is satisfied because chains use multiple price endings for all products in the sample.

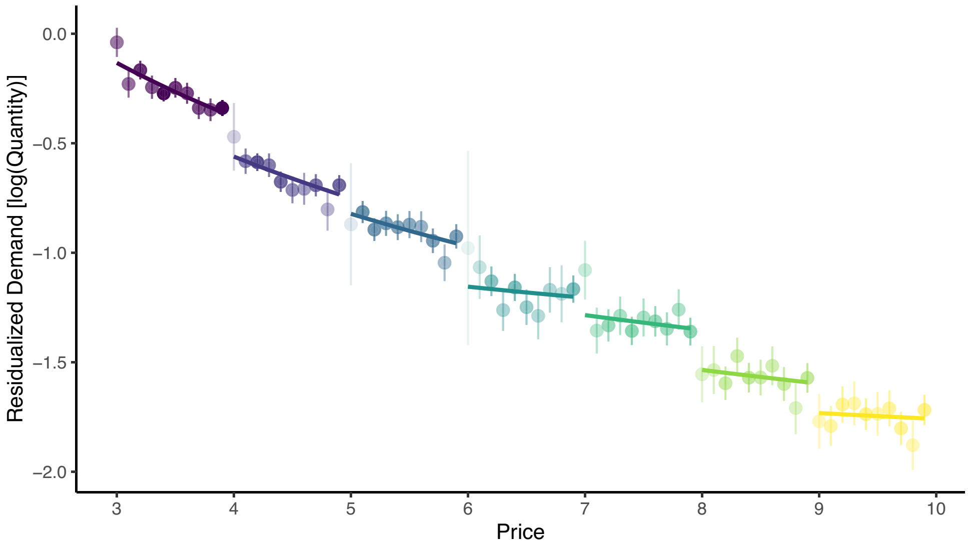

FIGURE 2

Demand curve of thirty-seven products, exhibiting drops at dollar digits

Notes: The figure shows a non-parametric demand curve. Data are for thirty-seven unique products with similar elasticities (between −1.35 and −1.65) from the module "ground and whole bean coffee". Each dot is the estimand from regressing log-quantity sold on a ten-cent dummy (*e.g.*, \$4.90–\$4.99), controlling for prices of competing products, and fixed-effects which are the interactions between UPC-chain and store, month-of-year, week-of-month, last digit, and promotional flags. Standard errors are clustered at the store level, and opaqueness in the figure is inverse to standard errors.

The key remaining identifying assumption is that there is no sorting of *price-endings* according to demand shocks. If products are priced at high-ending prices when demand is expected to be high, and proportionally so at low-ending prices when demand is expected to be low, it will generate the aforementioned sawtooth pattern.[14] Therefore, the analysis requires the explicit assumption that price-endings are not set in response to shifts in demand. Indeed, neither previous literature nor popular press recommends correlating the price-ending with demand fluctuations. Furthermore, interviews I conducted with retail executives support this assumption: according to them, price levels are determined by costs, competition, and price sensitivity, and price endings are the result of occasional rounding. Therefore, the observed sawtooth pattern of the residuals are characteristics of left-digit biased demand—rather than a result of firms' deliberate price-endings sorting according to demand shocks.

The above requirements and assumptions lead to the following regressions. I am using the following specification at the *product-chain* level, *i.e.*, I run 9072 separate regressions:

$$\log(Q+1)_{ist} = \epsilon_i \log P_{ist} + \alpha_{i0} \log P_{c(i)st} + \alpha_{i1} \log P_{1(i)st} + \alpha_{i2} \log P_{2(i)st} + \beta_{is} + \gamma_{i,year(t)} + \delta_{i,month(t)} + \psi_{i,week(t)} + \mu_i \text{Sale}_{ist} + \phi_i \text{Spell-length}_{ist} + e_{ist} \quad (3)$$

where i is the product-chain, s is the store, and t is the period. Q is number of units sold, and P is the price. ϵ captures the product-chain elasticity as the coefficient on log-price. The

14. An example of a violation of that assumption is if positive demand shocks increase prices from \$3.59 to \$3.99, but not from \$3.99 to \$4.29.

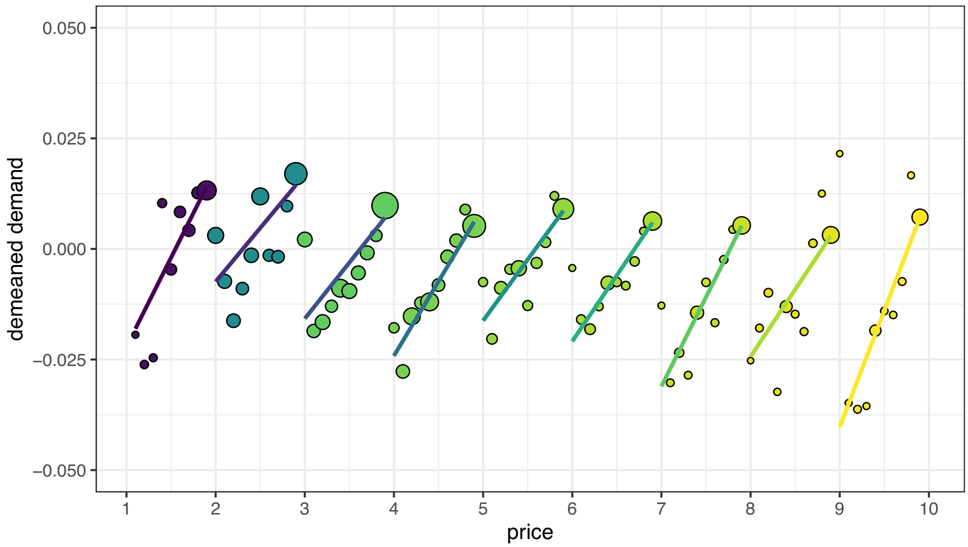

FIGURE 3
Sawtooth patterns of residualized demand by price

Notes: The figure plots demeaned log-demand by price, residualizing UPC-chain level price elasticity, seasonality, cross elasticities, store fixed effects, and promotion effects. Each of the dollar digits is represented by a different shade. The size of each circle represents the number of observations. Within-dollar linear fits, weighted by number of observations, are added as solid lines. Plotted fixed effects estimated at ten-cent bins (*e.g.*, \$3.90–\$3.99).

αs capture cross-elasticities, where $\log P_{c(i)st}$ is the average log-price of other products of the same product category $c(i)$ in the same store s at the same period t, and similarly $1(i)$ and $2(i)$ are the first and second most popular items in the category of i (excluding i itself). β is a store fixed effect capturing a time-invariant product average demand in a store; γ is a year fixed effect capturing product-level shifts in demand; δ is month-of-year fixed effect capturing product specific seasonality; and ψ is a week-of-month fixed effect capturing micro-seasonality within a month.[15] Finally, μ and ϕ capture transitory and dynamic effects of transitory sales: Sale_{ist} is an indicator for a transitory sale (whether a product's price is lower than the inferred base price in week t in store s); $\text{Spell-length}_{ist}$ is the number of consecutive weeks for which the same price was set for product i in store s at week t—to address issues of stockpiling and different pricing strategies.

The residuals from these regressions are, by definition, the unexplained component of demand. I take these residuals, $\hat{e}_{ist}$, and calculate a simple average across products, stores, and time, at ten-cent bins.[16] I annotate the different price bins as $\tilde{p}$ (*e.g.*, $\tilde{p} = \{\$2.90, \$2.91, \ldots, \$2.99\}$ is one bin):

$$\hat{Q}(\tilde{p}) = \frac{\sum \hat{e}_{ist} \mathbf{1}(p_{ist} \in \tilde{p})}{\sum \mathbf{1}(p_{ist} \in \tilde{p})}$$

15. See, *e.g.*, Hastings and Shapiro (2018).
16. Due to some extreme outliers, I winsorize the residuals at the 1% and 99% of the distribution.

Figure 3 shows the estimated $\hat{Q}(\tilde{p})$, the average residuals of demand per price bin (ten-cent bins). It shows that quantity purchased, netted of price, product, store, sales, and seasonality effects, is increasing within each dollar-digit and drops at dollar crossings.

Next, I conduct several tests with different methods and samples to verify that the findings are robust. First, I show the results from the same procedure as described above, only modifying specification (3) to use leave-one-out price as instrumental variables (commonly referred to as Hausman instruments). That is, the log-price of a product i in store s at time t is instrumented with the average price of product i at time t but in other stores of the same chain in other markets (other designated market areas (DMAs)). The idea is that changes in prices of the same product in the same chain outside of the DMA reflect changes in costs that are shared across markets but are not associated with shocks to local demand.[17] The results are shown in Supplementary Material, Appendix Figure A-1(a). Next, the patterns are robust for other demand structures at the product level—constant semi-elasticity which results from regressing log-quantity on prices rather than log-prices (Supplementary Material, Appendix Figure A-1(b)), or a flexible fifth-degree polynomial in prices (Supplementary Material, Appendix Figure A-1(c)). Finally, if the model is correct, regressing log-quantity on price and floor(price) incorporates the bias as modelled and should eliminate the sawtooth pattern. Indeed, Supplementary Material, Appendix Figure A-1(d), shows that parameterization of left-digit bias is a good description since adding the floor of the price as a regressor flattens the residuals (cf. Supplementary Material, Appendix Figure A-1(c)).

Furthermore, I test whether these discrete drops in demand when the dollar digits change are statistically significant. I regress demand at the UPC-chain level by adding two dummy variables capturing discrete jumps in demand at two round prices—the floor and ceiling of a product-chain's median price.[18] For example, if the median price is \$3.42, I estimate discrete changes at \$3 and \$4. Although this estimation is somewhat noisy due to the small number of prices at the product-chain level, it shows that demand drops on average by about 6–7% when dollar digits change (over and above the price elasticity). Further details on the distribution of drops are shown in Table 2, panel A, and in Supplementary Material, Appendix Figure A-2.

In summary, sawtooth patterns are exhibited across various specifications, samples, products, and retailers, and are eliminated when the structure of the bias is incorporated into the model. Taken together, these patterns support the existence and soundness of the model and justify a structural estimation.

2.4. *Structural estimation of left-digit bias*

I now turn to connect the reduced-form evidence with the model by estimating the left-digit bias parameter θ. This is a parameter of interest for multiple reasons. First, it provides a clear interpretation by quantifying the fraction of the cents-component consumers ignore. Second, structural estimation allows to describe the demand curve with a functional form and find

17. These instruments are strong because chains implement uniform pricing across stores, making them similar to OLS as uniform pricing means that pricing is not a response to demand at the local level.

18. Call these floor and ceiling dummies D^f_{ist} and D^c_{ist}, *e.g.*, $D^f_{ist} = \mathbf{1}(p_{ist} \geq \lfloor p_{i\bar{c}(s)} \rfloor)$. Then the specification I run is

$$\begin{aligned}\log(Q+1)_{ist} = {} & \epsilon_i \log P_{ist} + \alpha_{i0} \log P_{c(i)st} + \alpha_{i1} \log P_{1(i)st} + \alpha_{i2} \log P_{2(i)st} \\ & + \rho^f D^f_{ist} + \rho^c D^c_{ist} \\ & + \beta_{is} + \gamma_{i,year(t)} + \delta_{i,month(t)} + \psi_{i,week(t)} + \mu_i \text{Sale}_{ist} + \phi_i \text{Spell-length}_{ist} + e_{ist} \qquad (18)\end{aligned}$$

TABLE 2
Estimates of drops and left-digit bias parameter

		Mean estimate	Number of estimates	Percentile 25th	Percentile Median	Percentile 75th
Panel A: Drops in demand (% change)						
(1)	Chain	−6.61	25	−8.68	−6.83	−3.34
(2)	Module	−6.36	41	−8.27	−6.40	−4.73
(3)	Module-chain	−6.24	872	−12.21	−6.10	−0.62
(4)	UPC	−7.28	3206	−14.79	−6.43	1.33
(5)	UPC-chain	−6.62	13,760	−16.85	−5.74	4.49
Panel B: Left-digit bias estimates						
(6)	Chain	0.22	25	0.17	0.22	0.25
(7)	Module	0.21	41	0.17	0.21	0.28
(8)	Module-chain	0.20	869	0.11	0.21	0.31
(9)	UPC	0.22	3196	0.06	0.20	0.36
(10)	UPC-chain	0.22	7900	0.04	0.21	0.40

Notes: Panel A shows distributions of changes-in-demand at dollar crossing, and panel B shows the distribution of left-digit bias estimates. Changes in demand are estimated as discrete breaks in demand at the UPC-chain median price's dollar digit and the digit above it (*e.g.*, if the median price is \$3.42, it will have drops at \$3 and \$4). The numbers in the column titled "Number of estimates" may be larger due to the doubling effect (*e.g.*, the number of estimated drops is larger than the number of UPC-chains). The estimates describe the discontinuous change in demand at these thresholds. For example, a mean estimate of −6.61 means that demand drops on average by 6.61% when the dollar digit changes, above and beyond the price sensitivity. All rows display shrunk estimated drops at the UPC-chain-digit level, aggregated at different levels. Row (1) shows estimated drops in demand averaged at the chain level; row (2) is averaged at the module level; row (3) is averaged at the module-chain level; row (4) shows product-level estimates; and row (5) shows the underlying UPC-chain level estimates.
Panel B shows estimates of left-digit bias and their distributions. For example, a mean estimate of 0.22 means that the difference between a ninety-nine-ending price and the round price above is perceived as about twenty-two cents difference. Rows display averaged left-digit bias based on UPC-chain level estimates. Row (6) averages estimates at the chain-level; row (7) averages at the module level; row (8) is averaged to module-chain; row (9) is at the UPC level (across chains); and finally, row (10) shows the underlying UPC-chain level shrunk estimates.
Supplementary Material, Appendix Table A-1, displays the equivalent estimates estimated directly at the various levels (chain, module, etc.)—both raw and shrunk.

(analytical) optimal pricing. Finally, resulting quantification enables to conduct counterfactual exercises.

The left-digit bias parameter is estimated at the product-chain level—I estimate multiple θ_{ic}s where i is UPC and c is chain—for two main reasons. First, an alternative aggregated estimation can lead to a form of statistical aggregation bias by ignoring different price elasticities (and it indeed leads to larger estimates). Second, it allows to explore heterogeneity of the left-digit bias estimates across modules, chains, and products.

I estimate θ by substituting the perceived price $\hat{p}$ as the explanatory variable instead of price. That is, I run a similar specification to specification[19] (3) and replace log-price with the price linearly, resulting with the following specification:

$$\log(Q+1)_{ist} = \beta_i^p P_{ist} + \beta_i^{\lfloor p \rfloor} \lfloor P \rfloor_{ist} + \alpha_{i0} P_{c(i)st} + \alpha_{i1} P_{1(i)st} + \alpha_{i2} P_{2(i)st} + \beta_{is} + \gamma_{i,year(t)} + \delta_{i,month(t)} + \psi_{i,week(t)} + \mu_i \text{Sale}_{ist} + \phi_i \text{Spell-length}_{ist} + e_{ist} \quad (5)$$

19. Notice that Δ is not identified because it is absorbed by the constant.

Since I define the perceived price as $\hat{p} = (1-\theta)p + \theta \cdot (\lfloor p \rfloor + \Delta)$, doing so effectively equates to regressing log-quantity on $\hat{p}$, with the coefficient on $\hat{p}$ being the semi-elasticity.[20] Therefore, the estimated left-digit bias parameter is

$$\hat{\theta} = \frac{\hat{\beta}^{\lfloor p \rfloor}}{\hat{\beta}^{p} + \hat{\beta}^{\lfloor p \rfloor}}$$

where $\hat{\beta}^{p}$ is the coefficient on the exact price and $\hat{\beta}^{\lfloor p \rfloor}$ is the coefficient on the floor of the price. Standard errors, which are clustered at the store level, are obtained using the Delta-method.[21]

Because of the high number of estimands, estimation is noisy and calls for some regularization. This a common issue with scanner data estimation, and it has a common solution of Bayesian shrinkage of the estimates (see Butters *et al.*, 2019; DellaVigna and Gentzkow, 2019). Specifically, I take the most precise estimates—those with below-median standard errors—and calculate their mean and variance, denoted as $\bar{\theta}$ and $\mathrm{Var}(\theta)$, respectively. Then, I recalculate each estimated θ_{ic} as $\tilde{\theta_{ic}} = (\mathrm{Var}(\theta) \cdot \theta_{ic} + \sigma^2(\theta_{ic}) \cdot \bar{\theta})/(\mathrm{Var}(\theta) + \sigma^2(\theta_{ic}))$. In other words, the estimates are being "shrunk" to a prior with mean $\bar{\theta}$ and variance $Var(\theta)$, a procedure in which the more precise the estimate (the lower its standard error $\sigma(\theta_{ic})$), the lower the weight on the prior.[22]

Panel B of Table 2 shows the results. The procedure generates estimates at the product-chain level with an average left-digit bias of 0.22 and median of 0.21, with interquartile range of 0.04–0.40. The full distribution can be seen in Figure 4. An advantage of the UPC-chain level estimation is the ability to calculate left-digit bias at more aggregated levels—such as across-chains or product-modules. As evident in Table 2 and Figure 4, a significant heterogeneity between chains and modules remains. For example, the interquartile range between chains is 0.17–0.25. In the next section, I explore the relationship between the above-mentioned heterogeneity and pricing.[23]

2.5. *Exploring left-digit bias heterogeneity*

Product-chain level estimates allow to explore the bias' correlation with the characteristics of each product and its clients. While the demand-side model makes no assumptions on the source of the bias, correlated factors with left-digit bias can inform us on possible mechanisms that might generate it, or shed light on its source (see also Macé, 2012; Jiang, 2020; Sokolova *et al.*, 2020). As this is not at the core of the paper, I defer the analysis and discussion to Supplementary Material, Appendix F. There, using the characteristics data from the NielsenIQ Consumer Panel, I find that, unlike factors such as buyers' education, income, or purchase frequency, a higher

20. In a previous version of the paper (Strulov-Shlain, 2019b), I used two additional methods to estimate left-digit bias—using the drops in demand coupled with separately estimated elasticities, and non-linear least squares. Since all methods generated nearly identical results, for brevity's sake, I opted for the simplest one here.

21. The raw estimates without Bayesian shrinkage are presented in Supplementary Material, Appendix Table A-1.

22. I have also estimated the parameters at those more aggregated levels directly, and they are consistently larger. For example, while in the main approach (averaging across UPC-chain by chain) the interquartile range is 0.17–0.25, in the aggregated version in which I directly estimate the bias at the chain level, the interquartile range is 0.36–0.58. Similarly, estimating a single parameter, even with very flexible controls, gives an estimate of 0.570 (0.009)—a less credible estimate due to the aforementioned statistical aggregation bias.

23. In Supplementary Material, Appendix B, I solve a case of firms pricing under Bertrand-Nash price-competition with left-digit biased consumers. The case of a multi-product monopolist is analysed in Supplementary Material, Appendix C.

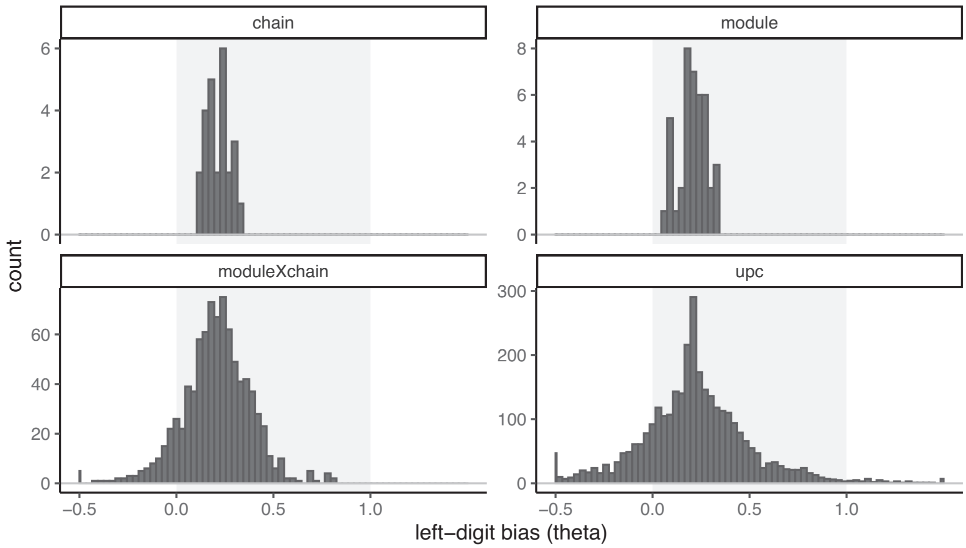

FIGURE 4

Histogram of product-level left-digit bias estimates from drops in demand at the dollar-digit

Notes: The figure shows histograms of left-digit bias shrunk estimates from regressing log(quantity) on the perceived price and controls (details in Section 3.4). The light-grey background highlights thetas (LDB) of 0–1. Inside textbox displays the level at which left-digit bias is aggregated—chain, module, module-by-chain, and product.

average price of an item is associated with a stronger bias. For example, over and above the price elasticity, ninety-nine cents seem to get lower weight in the perceived price for a $7.99 item versus a $1.99 item.

3. FIRM PRICING

Building on the findings of the demand structure in Section 2, I investigate firms' pricing response—what it should be and what can describe it in practice. I estimate firms' beliefs about demand as revealed by their pricing behaviour, and also entertain alternative pricing procedures.

3.1. *Optimal pricing*

To study pricing, I analyse a model of monopolistic pricing facing left-digit biased demand since it is estimable and allows me to incorporate the demand analysis. In the existing literature, there are only a few models of firm pricing that relate directly to nine-ending prices (Basu, 1997, 2006; Stiving, 2000), and for those that do, their predictions for optimal pricing are either only ninety-nine-ending prices or none. Another approach, of bounded-rationality models, generates discrete price-setting behaviour, but these discrete prices need not be nine-ending (Chen *et al.*, 2010; Gabaix, 2014; Matějka, 2015). The model in this paper consists of left-digit-biased consumers and monopolistic firms.[24] Consumers choose a quantity to purchase to maximize utility, but may

24. This is a common feature of models with inattention (*e.g.*, Farhi and Gabaix, 2015; Gabaix, 2019), in which there is a discrepancy between the true utility and maximized utility on the consumer side, causing distortions.

misperceive the price. Firms price to maximize profits according to their beliefs of consumer misperception.

Consider a monopolist who faces a demand function as in Section 2, $D(p;\theta,\Delta)=A\hat{p}^{\epsilon}$, a function of prices and left-digit bias. The firm earns p per unit, and pays a fixed unit cost c. Consumers perceive the price as $\hat{p}(p;\theta,\Delta)$ and choose quantities accordingly. Then, the monopolist's gross profits (absent fixed costs) are

$$\Pi(p;\theta,\Delta)=D(p;\theta,\Delta)\cdot(p-c)=A\hat{p}^{\epsilon}\cdot(p-c)$$

i.e., demand is driven by the *perceived* price while per unit profits are governed by the *true* price.[25]

For ease of exposition, I decompose a price p into its decimal basis components, such that, $p=p_1+p_{0.1}+p_{0.01}$. For example, $p=3.49=3+0.40+0.09$, so $p_1=3$.

The discontinuity in demand when $p\to p_1$ is easy to see. In the limit from below, demand is $A(p_1+\theta\Delta-\theta)^{\epsilon}$ versus $A(p_1+\theta\Delta)^{\epsilon}$ from above. To simplify the solution, assume that prices can be chosen from all real numbers between a natural number and a ninety-nine-ending number—$\underset{q_1}{\cup}[q_1,q]\subset\mathbb{R}_+$ where $q_1\in\mathbb{N}$ and $q=q_1+0.9+0.09$. Proofs to the following propositions and corollaries appear in Supplementary Material, Appendix A.

Solving for optimal pricing yields the following result:

Proposition 1 (optimal pricing formula, for small θ)**.** *For any cost c and parameters ϵ and θ, find the appropriate .99 ending number $q=(q_1,.9,.09)$ such that $c\in[\underline{c}_q,\underline{c}_{q+1}]\triangleq[q\cdot(1+1/\epsilon)+\theta/(1-\theta)\cdot(q_1+\Delta)/\epsilon,\ (q+1)\cdot(1+1/\epsilon)+\theta/(1-\theta)\cdot(q_1+1+\Delta)/\epsilon]$. Then, the optimal price for that cost c is*

$$p(c;\theta,\Delta,\epsilon)=\begin{cases} q & \text{if } c\in\left[\underline{c}_q,\overline{c}_q\right] \\ \left(c-\dfrac{\theta}{1-\theta}\cdot\dfrac{q_1+1+\Delta}{\epsilon}\right)\cdot\dfrac{\epsilon}{1+\epsilon} & \text{if } c\in\left[\overline{c}_q,\underline{c}_{q+1}\right) \end{cases} \tag{6}$$

where $\underline{c}_q$ and $\underline{c}_{q+1}$ are defined above, and $\overline{c}_q$ is defined below as the minimal cost for which it is profitable to price strictly above q ($\overline{c}_q$ solves Equations (7) and (8)).

The pricing behaviour in Equation (6) is different from standard pricing with no-bias in two ways. First, the top case in Equation (6) represents that a region of costs maps to ninety-nine-ending prices, and some of these costs would have otherwise been mapped to prices just above ninety-nine-ending prices.[26] Second, the bottom case describes interior solutions as a modified markup rule with an added component driven by the bias. The price is slightly higher than it would have been absent the bias because left-digit bias makes demand less elastic for changes within the same first digit.

To further elucidate Equation (6), I elaborate on the three threshold costs $\underline{c}_q$, $\underline{c}_{q+1}$, and $\overline{c}_q$. The former two, $\underline{c}_q$ and $\underline{c}_{q+1}$, are the costs for which the monopolist's profit is maximized as an *interior* solution at q and $q+1$, respectively, *i.e.*, where the first-order condition is satisfied at ninety-nine-ending prices. The image of costs in $[\underline{c}_q,\underline{c}_{q+1}]$ is hence prices in $[q,q+1]$, which allows to plug in q and q_1 in the two cases in Equation (6). The third threshold, $\overline{c}_q$, is the cost for which prices will switch from q to a price in the segment $[q_1+1,q+1]$. At $\overline{c}_q$, two

25. The top equation shows that prices are set at ninety-nine-ending prices (q) for a range of unit costs, and hence with varying markups. Thus, this model generates price stickiness that is not driven by frictions on the supplier side.

26. For a very low θ, there is no gap, and $P=q_1+1$.

conditions are met: (1) the profit is maximized (with an internal solution) on some price $P \in [q_1+1, q+1] > q$; and (2) profits are equal at that price, P, and at q:[27]

$$\bar{c}_q = \begin{cases} P + \dfrac{(1-\theta)P + \theta\,(P_1+\Delta)}{\epsilon\,(1-\theta)} \\ \dfrac{\hat{P}^{\epsilon}\cdot P - \hat{q}^{\epsilon}\cdot q}{\hat{P}^{\epsilon} - \hat{q}^{\epsilon}} \end{cases} \tag{7}$$

Definition 1 (Next-Lowest Price)**.** The *Next-Lowest Price* P is the lowest price used above a ninety-nine-ending price q, and is a function of the parameters θ and ϵ as defined by the following implicit equation:

$$P + \frac{(1-\theta)P + \theta\,(P_1+\Delta)}{\epsilon\,(1-\theta)} - \frac{((1-\theta)P + \theta\,(P_1+\Delta))^{\epsilon}\,P - ((1-\theta)q + \theta\,(q_1+\Delta))^{\epsilon}\,q}{((1-\theta)P + \theta\,(P_1+\Delta))^{\epsilon} - ((1-\theta)q + \theta\,(q_1+\Delta))^{\epsilon}} = 0 \tag{8}$$

This equation is not analytically solvable, but can be easily solved numerically for P. Given P, we can recover $\bar{c}_q$ from either part of Equation (7). The empirical meaning of P is that no prices should be set between q_1+1 and P.

The pricing schedule is illustrated in Figure 5, showing price as a function of cost under left-digit biased demand. The diagonal grey line are prices without bias, and the thick lines are the optimal prices with the bias. The figure demonstrates regions of costs translating to ninety-nine-ending prices (some with lower markups than absent a bias), and higher markups for non-bunching prices.

The testable predictions of the model—bunching at ninety-nine-ending prices, with asymmetry of missing prices only above the round price thresholds—are the key to taking the model to the data (see also Dube *et al.*, 2017).

Proposition 2 (Comparative statics of the Next-Lowest Price)**.** *For $\epsilon < -1$, the next-lowest price is lower when demand is more elastic, i.e., $\partial P/\partial|\epsilon| < 0$. For $q_1 > 0$, the next-lowest price is higher when there is more bias, i.e., $\partial P/\partial\theta > 0$.*

Proposition 2 shows that more elastic demand leads to a lower next-lowest price, but this is a result of counteracting forces. The first-order effect is that more elastic demand will cause a steeper decline in profits for the same price difference, and hence leads to a lower P to preserve profit indifference. On the other hand, a higher elasticity means lower markups overall and a shift in the profit curves. After some algebra, the latter can be shown to be a weaker effect.

The second part of Proposition 2 states that higher bias leads to higher next-lowest prices. The effect of higher bias is simpler to understand. Fix some optimal P and consider an increase in θ. The infra-marginal benefits from changing the price to the lower ninety-nine are now larger (because the perceived gap, $\hat{P} - \hat{q}$, has increased) while the costs of increasing the price above P are lower because of the lower sensitivity of within-left-digit demand. So P must increase.

A corollary is that the right-digits component of the next-lowest price is increasing from one threshold to the next:

Corollary 1. *For $\epsilon < -1$, conditional on θ, the lower digits of the Next-Lowest Price (i.e., $P - P_1$) are increasing with the first digits.*

27. Rule-of-thumb pricing is similar to what Dube *et al.* (2017) find in wage setting by employers. Interestingly, they do not find left-digit bias on the employees side.

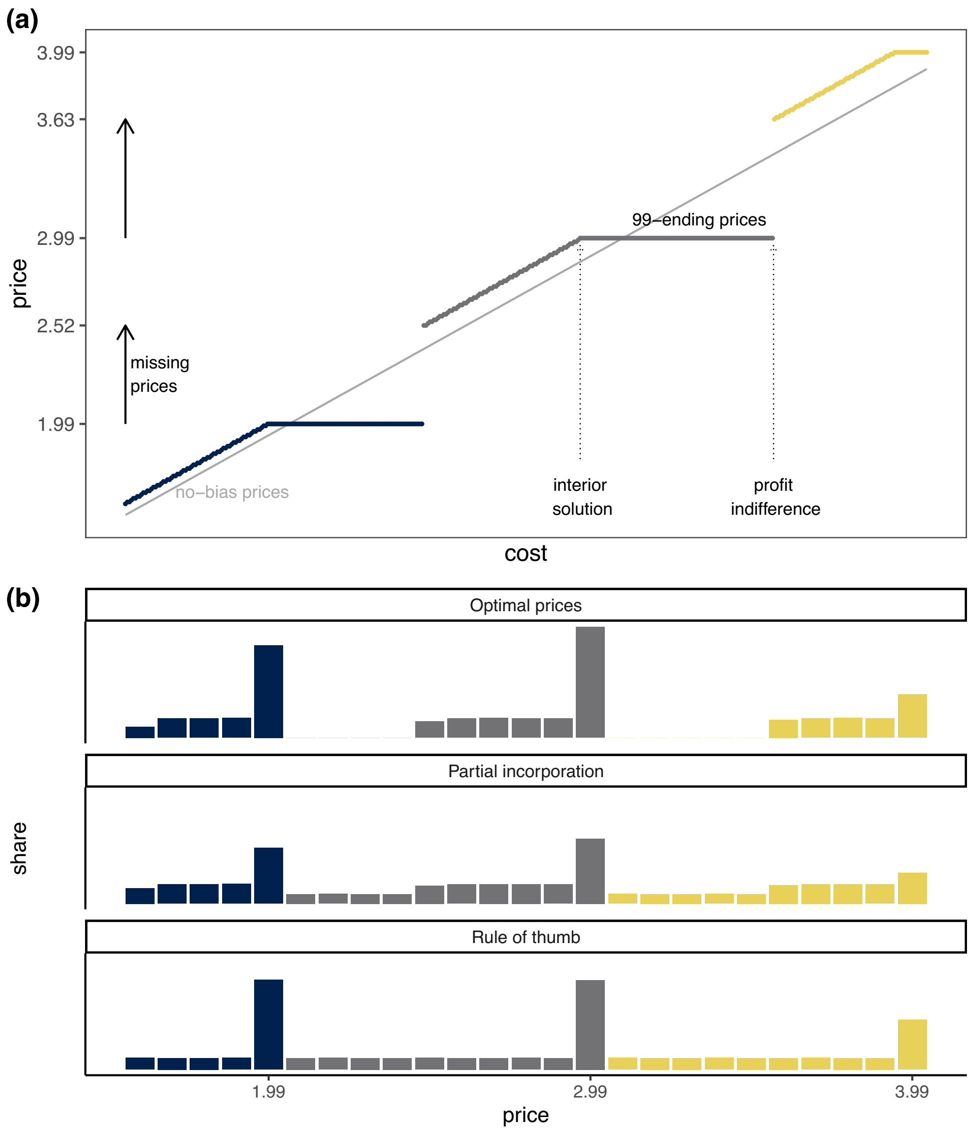

FIGURE 5

Illustration of the optimal price schedule

Notes: Panel A: The thick lines in different shades illustrate optimal prices of a monopolist facing constant elasticity demand with left-digit biased consumers, displayed as a function of cost. The thin grey line represents the no-bias counterfactual where $p = c(\epsilon/(\epsilon+1))$. The horizontal sections of the thick lines are the ninety-nine-ending prices, which are the optimal prices for regions of costs. The lowest cost that translates to a ninety-nine-ending price is when this cost is the interior solution at ninety-nine, and the highest cost is when there is profit indifference between ninety-nine and the next-lowest price. The resulting ranges of missing prices are highlighted by vertical arrows. Panel B: Illustration of the resulting price distributions with uniform costs for different pricing procedures. From top to bottom: optimal pricing (matches panel A); with partial incorporation; with rule-of-thumb pricing.

This corollary means that nominally more expensive products have larger regions of missing prices. For example, for the same parameters of left-digit bias and elasticity, the lowest observed price above \$4 should be \$4.30 and above \$5 should be \$5.32.

Proposition 2 and Corollary 1 show that left-digit bias and price elasticity map into next-lowest prices and excess mass. Therefore, point identification of left-digit bias and elasticity from these price moments is theoretically possible. The parameters have differing effects on the moments of missing prices and excess mass at ninety-nine-ending prices, which can be pinned down with some further assumptions.

3.2. *Other pricing procedures*

While the model in Section 3.1 prescribes optimal pricing, firms can follow other pricing procedures, and I consider two classes of deviations: "partial incorporation" and "rule-of-thumb rounding". The first class of deviations, *partial incorporation*, means that firms take left-digit bias into account and use the optimal pricing formula but only sometimes. The motivation for this deviation is that different people or rules set prices at different times, stores, or products. On aggregate, a firm can understand the bias but only incorporate it into pricing in some occasions. The second class of deviations is model-free occasional price-rounding behavior. It is motivated by price setters who learned that ninety-nine-ending prices are good, but are not sure when and why (see also Strulov-Shlain, 2019a). A natural pricing rule is then rounding to a ninety-nine-ending price if the price is "close".

To operationalize partial incorporation pricing, consider a price $p(c;\theta,\Delta,\epsilon)$ from Equation (6). Under partial incorporation, the actual price set is a mixture of that optimal price with probability ζ and a price that ignores left-digit bias:

$$p^{\zeta}=\begin{cases}p(c;\theta,\Delta,\epsilon) & \text{with probability } \zeta\\ p(c;0,\Delta,\epsilon) & \text{with probability } 1-\zeta\end{cases} \tag{9}$$

In contrast, *rule-of-thumb pricing*, is not based on direct optimization. Instead, I assume that companies believe that potentially "better" prices are ninety-nine-ending or forty-nine-ending without employing careful calculation when should they be used (forty-nine-ending prices are also prevalent in the data yet are not associated with excess demand, see Figure 1). To match patterns in the data, since not all prices end with ninety-nine or forty-nine, I further allow them to apply that belief only occasionally. Specifically, the price is rounded to ninety-nine or forty-nine with probability ξ, and rounded to the nearest ninety-nine-ending if it falls within a d-width segment centred at ninety-nine (otherwise rounded to nearest forty-nine); with probability $1-\xi$, it is set as in the standard case, as if there is no left-digit bias. That is,

$$p^{\xi,d}=\begin{cases}\begin{cases}round(p)-0.01 & \text{if } |p-(round(p)-0.01)|\le\dfrac{d}{2}\\ \lfloor p\rfloor+0.49 & \text{otherwise}\end{cases} & \text{with probability } \xi\\ p(c;0,\epsilon) & \text{with probability } 1-\xi\end{cases} \tag{10}$$

The three procedures—optimal, partial incorporation, and rule-of-thumb—differ in predictions about the shape of the price distribution. While all pricing procedures generate excess mass at ninety-nine-ending prices, they differ in the predictions for low versus high price endings. This is illustrated in Figure 5(b), showing the distribution of prices under each pricing procedure (assuming uniform cost distribution). Optimal pricing leads to excess mass at ninety-nine-ending prices and missing mass of low-ending prices. The partial incorporation procedure keeps the asymmetry in which there are fewer low-ending prices than high-ending, but unlike the optimal pricing procedure, it predicts a positive mass of prices everywhere, and in particular, for low-ending prices (of measure $1-\zeta$). In contrast, the rule-of-thumb procedure leaves a positive mass of prices everywhere, symmetrically pulling prices into ninety-nine-ending prices from above and below.[28]

3.3. *Pricing patterns in the data*

Given the large magnitude of left-digit bias in demand, the next step is to examine firms' pricing behaviour in response to the bias. The optimal pricing model predicts: (1) excess mass at ninety-nine, (2) missing prices with low price-endings, and (3) more missing prices with higher dollar digits. In contrast, partial incorporation predicts fewer prices with low price endings rather than none at all (changing prediction 2), and rule-of-thumb pricing means symmetric distribution of low and high price-endings. I first examine the stylized patterns in the data to see if they are indeed in-line with these predictions, and which procedure seems most aligned with the empirical patterns.

Figure 6 shows the price-endings histogram for all prices in the final sample across all products and chains. This sample is the same one used for demand estimation, thus being internally consistent. Pricing in that sample should be driven by the estimated demand-side parameters.

First, there is an excess mass at ninety-nine: 41% of prices indeed end with ninety-nine. Furthermore, almost all prices end with nine as the cent digit—with a share of 87%.[29] Furthermore, there is a noticeable spike at forty-nine-ending prices with 14% of the total prices. Note that in contrast, in demand data in Figure 1, there is no excess demand at these forty-nine-ending prices (except, perhaps, at the price bin of $2.40–$2.49).

Second, Figure 6 shows that there are marginally fewer low price endings. There are fewer prices ending with cents component lower than nineteen, especially for regular prices. However, while there are almost no 09-ending prices (0.08% compared with 4.5% for nineteen-ending prices), 3.8% of prices end with 00—of which 38% are on-sale prices. The different price ending distribution for regular prices versus on-sale prices is noticeable but for simplicity, I ignore it in this paper. Overall, it does not seem the excess mass at ninety-nine is drawn asymmetrically and exclusively from otherwise low-ending prices. This conclusion will be tested formally in the next section.

Finally, Figure 6 masks the heterogeneity of how price endings differ between low and high dollar digits. Under optimal left-digit bias facing pricing, next-lowest prices should increase with the dollar digit if the bias and elasticity are constant. Recall that Supplementary Material, Appendix F, shows correlates of the bias with covariates, finding that the bias is increasing with

28. The latter finding also implies that there might be an added bias ignoring the last digit of the price, hence leading to many nine-endings. This effect is probably true but impossible to estimate, exactly because there is a lack of variation in the last digits, and hence the drops in demand at the dime thresholds cannot be estimated. However, I show that conclusions are insensitive to assuming additional bias regarding the last digits (that is, assume that with some measure θ_2, consumers see the price as the dollar and dime digits but ignore the cents), in Supplementary Material, Appendix Table A-3.

29. In Public Economics, the idea is that a non-linear tax schedule creates incentives for tax payers to change their reported income (or labour supply) and bunch at kinks or notches of the tax schedule. The duality can be seen as follows: the tax schedule is the dual of the demand curve (elasticity); the notch given by the tax schedule is like the left-digit bias parameter; the utility function is like a profit function; and the latent ability distribution is like the cost distribution—estimated from the observed income distribution in the tax literature and from the non-bunching price distribution here.

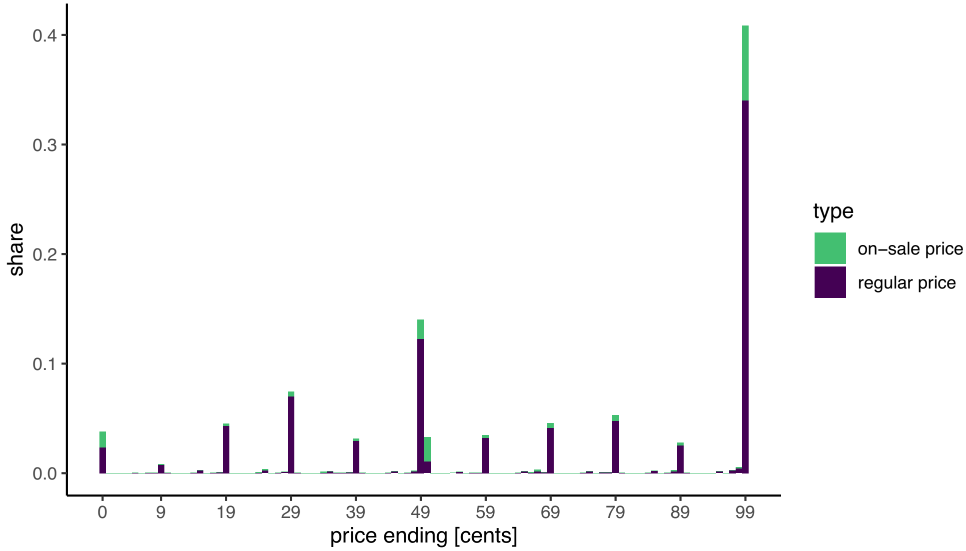

FIGURE 6
Pricing response: Price-endings distributions

Notes: The figure shows the price-endings histogram of cleaned Nielsen data across twenty-four chains (77.4 million observations, see Table 1 for more details). Light bars represent on-sale prices, while dark bars represent regular prices. Main patterns are excess mass at nine- and ninety-nine-ending prices, and some missing prices at low price endings.

the price, thus this prediction becomes more pronounced. The increase in next-lowest prices is illustrated in the by-dollar CDF of price endings in Supplementary Material, Appendix Figure A-3, where the darker lines represent lower prices (lower dollar digit). Indeed, Supplementary Material, Appendix Figure A-3, shows that the distribution of price-endings is shifting to the right, meaning that for higher dollar digits, there are fewer low price-endings.

3.4. *Structural estimation of pricing procedures*

Firms' behaviour is only partially consistent with patterns of optimal pricing facing the left-digit biased demand found in the data: prices bunch at ninety-nine; there are some missing prices with low price endings but without a pronounced asymmetry; and the price-ending distribution is shifting to the right with higher prices as dollar digits increase. I now turn to analyse more granular pricing patterns in order to learn about the pricing procedures of the different chains and in the aggregate.

Optimal pricing given the demand parameters should lead all prices to end at ninety-nine, but less than half do. Consider an optimal pricing with left-digit bias of 0.2 and an elasticity of -1.5. With these parameters, "next-lowest prices" are ninety-nine. *i.e.*, regardless of the underlying costs, *all* prices should end at ninety-nine. Since only 41% of actual prices end with ninety-nine, 59% of prices are dominated (by a lower or higher ninety-nine-ending price). Quantifying the effects of this mis-optimization requires a structural interpretation of firms' behaviour.

To quantify the deviation and study its implications, I fit each pricing procedure—optimal, partial incorporation, and rule-of-thumb—to explain the actual price distribution. For each pricing procedure I search for the parameter values that minimize distance between predicted and actual prices. The pricing procedures are "as if" models used to infer the *perceived* left-digit bias levels (and other parameters of partial incorporation and rule-of-thumb pricing) that rationalize

TABLE 3
Perceived parameters of pricing procedures

	Specification	Elasticity	Perceived LDB	Frequency of LDB incorporation	Rule-of-thumb frequency	Goodness of fit
(1)	Rule-of-thumb ninety-nine or forty-nine	−1.5			0.4 [0.29, 0.58]	0.0018
(2)	Estimate left-digit bias and its incorporation	−1.5	0.017 [0.015, 0.02]	0.61 [0.53, 0.68]		0.0023
(3)	Estimate left-digit bias and its incorporation	−3	0.07 [0.062, 0.088]	0.6 [0.53,0.68]		0.0023
(4)	Assume left-digit-bias (0.1), estimate incorporation	−3	0.1	0.5 [0.43, 0.56]		0.0026
(5)	Estimate left-digit bias	−3	0.032 [0.024, 0.039]	1		0.0033
(6)	Estimate left-digit bias	−1.5	0.0075 [0.0058, 0.0094]	1		0.0034
(7)	Assume left-digit-bias (0.2), estimate incorporation	−1.5	0.2	0.29 [0.25, 0.33]		0.0048

Notes: This table presents estimated parameters of pricing procedures. Each row represents a different pricing procedure or assumptions. Numbers in brackets represent the range of 2.5% and 97.5% of estimated parameters under different assumptions with 300 bootstraps for each; if there are no square brackets, the parameter value was assumed and not estimated. Rows are ordered by their goodness-of-fit measured as the average sum-of-squared errors (SSE) between predicted and empirical moments. Row (1) assumes the companies know the demand estimated price elasticity of −1.5, and round prices to ninety-nine-ending nearest prices in 40% of instances. Rows (2)–(7) show estimates of a pricing procedure in which firms price according to their perceived left-digit bias level, with some frequency of incorporation. Row (2), for example, shows estimates where left-digit bias is taken into account when pricing in 61% of instances, and when it does, firms price as if left-digit bias level is 0.017. Row (6) can be thought of as based on incorrect belief yet sophisticated benchmark. In row (6), firms follow the optimal pricing procedure, but with potentially wrong beliefs about left-digit bias levels. Indeed, perceived left-digit bias is estimated at 0.0075 (cf. the demand estimated bias of more than 0.2). In addition, the average SSE is 0.0034 versus 0.0018 for the rule-of-thumb nearest rounding pricing procedure.

retailers' pricing behaviour. That is, I assume that retailers are pricing according to different models, predict prices given the underlying parameters for each procedure, and examine the parameters of each pricing procedure and its goodness of fit. If in Section 2, I estimate the elasticity and left-digit bias that explain *demand* behaviour, here I estimate the elasticity, left-digit bias, and other parameters that rationalize *supply* behaviour.

The results of these estimates are shown in Table 3, but before discussing the results, I discuss the estimation procedure.

Estimation procedure. As shown in Section 3.1, excess mass and missing prices can identify perceived left-digit bias. This is similar to identification of labour supply elasticity from notches in the tax schedule (see Saez, 2010; Kleven and Waseem, 2013).[30] The similarities are that discontinuities lead to bunching. A key component in estimation is determining which part of the bunched mass is "excess" mass and which masses are missing. To learn those quantities

30. In fact, the same assumption is taken in the taxation literature, often implicitly. There, the argument is that excess mass at a tax notch is driven by the notch and not because there is a mass of people with skills who happen to earn exactly that amount.

requires creating a counterfactual distribution of the masses absent the notch. The tax literature often relies on the assumption that there are no distortions below or far enough above the discontinuities. In contrast, as Proposition 6 shows, in our setting all prices are affected. Therefore, I take a more structural approach that uses the entire price distribution rather than just the excess mass. It requires an explicit assumption of a smooth cost distribution, *i.e.*, that masses at ninety-nine-endings are not due to a particular combination of costs and elasticities.[31] Since each of the models described in Section 3.3 translates cost, elasticity, and perceived left-digit bias, or pricing rules parameters, to prices, a cost distribution predicts the price distribution. I first show in detail how the perceived left-digit bias under the optimal pricing model is identified from the price distribution, and I then discuss the latter pricing procedures.

I estimate the models using simulated minimum distance. I create the empirical density function of prices at ten-cent bins , and do so at several levels of aggregation. For example, I calculate the empirical PDF, the share S_p of each price p set by retailer X between \$1.29 and \$7.89. These price densities are the empirical moments to be matched by the predictive pricing procedures. Aggregating prices at ten-cent bins across products is done to satisfy the need for a smooth price distribution; estimation using Minimum Distance is done to treat outliers (such as 00-ending prices) as random errors under the optimal pricing model, and to identify ζ in the partial incorporation procedure.[32] I describe the technicalities of the estimation procedure, with a detailed step-by-step algorithm in Supplementary Material, Appendix E. The idea behind the estimation is to generate a theoretical distribution of prices given the parameters, $\hat{S}_p$, and search for the parameters that minimize the distance between the empirical distribution S_p and the predicted one—a simulated method of moments estimation. Predicted price densities are simulated given the cost distribution, price elasticity, and left-digit bias. The main insight is that given elasticity and left-digit bias, the cost distribution is a known monotonic transformation of the price distribution, such that non-bunching prices provide a counterfactual shape parameters. Given the shape of the cost distribution we now have all the information required to generate the predicted pricing distribution for each parameters set. This is illustrated in Supplementary Material, Appendix Figure A-4, showing an example of empirical price distribution and the smoothed polynomial fitted on a subsample of the prices. The algorithm then uses the shape of the cost distribution together with perceived left-digit bias and elasticity to create a predicted price distribution and searches for the parameters that minimize its distance from the empirical distribution.

I use the same algorithm to estimate the partial incorporation pricing procedure which assumes that a share $1-\zeta$ of prices are priced ignoring left-digit bias. That is, given the cost distribution, predicted prices are a mixture of two distributions: prices governed by perceived left-digit bias $\hat{\theta}$ with a weight ζ, and prices set as if there was no bias with weight $1-\zeta$.

Estimating the rule-of-thumb pricing procedure is straightforward. Given the predicted smooth cost (and price) distribution absent left-digit bias, a share $\xi \cdot d$ of counterfactual prices is rounded to match the excess masses at ninety-nine-ending prices (and $\xi \cdot (1-d)$ at forty-nine-ending prices).

Estimating multiple models. I estimate several restricted and unrestricted versions of the pricing procedures, and estimate those at the aggregate and at the chain level. First, I estimate

31. Kleven and Waseem (2013) face a similar issue in which they observe income in certain regions of the income distribution where there should be none. They interpret it as frictions and implicitly assume that they are random. The approach I follow takes a similar stance as I put equal weight on pricing above or below the dollar threshold.

32. Supplementary Material, Appendix E, details the estimation algorithm, and also discusses the feasibility of estimating the price elasticity from the price distribution.

the optimal pricing procedure assuming chains know the price elasticity but might have wrong beliefs on the bias, and estimate their perceived left-digit bias. Second, I estimate the partial incorporation procedure assuming that firms know the demand structure correctly, but sometimes choose to ignore it. Third, I allow for both wrong beliefs on left-digit bias and partial incorporation. Finally, I estimate the rule-of-thumb pricing.

Considering that the underlying true elasticity is a key parameter as it impacts the estimated bias, I test the implications of twice-as-elastic demand. The average elasticity is roughly -1.5 with an accompanied left-digit bias of 0.2. However, absent truly random variation in prices, there is a concern that elasticity estimates are biased upward. If that is the case, and maintaining the assumption that price endings are uncorrelated with demand shocks, an actually more elastic demand will also mean that true left-digit bias is proportionally lower. Therefore, to increase credibility of the results, I take a wide range of elasticity estimates as the "truth"—between the average estimate of -1.5 to double that number (-3) with a corresponding updated "true" bias of 0.1.[33]

I estimate the perceived parameters in multiple ways to mitigate the sensitivity to functional form assumptions and to learn about heterogeneity. To limit the sensitivity of the inferred parameters to incidental assumptions, I use a variety of sub-samples from each chain (namely, I sample different parts of the price distribution and bootstrap those), and I make various assumptions on the counterfactual smooth distribution (*i.e.*, the polynomial degree, the moments used for estimating the shape). For each pricing procedure version and assumed elasticity, I make sixteen assumption combinations,[34] times 300 product-level bootstraps of the underlying price distribution for each.

Results. The results are shown in Table 3 for the aggregate sample. Each row reports the estimated parameters, a 95% bandwidth of estimates, and mean SSE (difference between predicted and empirical moments). If a parameter is assumed rather than estimated it has no estimates bandwidth. Models are ordered according to their average goodness-of-fit.[35] An example of the empirical moments versus the predicted ones for one sample and one set of assumptions is shown in Supplementary Material, Appendix Figure A-5.

A few robust findings arise, showing that either there is substantial underestimation and partial incorporation, or that firms use rule-of-thumb pricing. First, firms' pricing behaviour reflects significant underestimation of the true left-digit bias. In Table 3, row (6), assuming firms know the correct elasticity and take the bias into account in all their pricing decisions, I find an estimated bias of 0.0075 [0.006, 0.009] instead of 0.2, which is twenty-six times lower. Even in the extreme hypothetical case assuming the true elasticity is -3, I still find a perceived left-digit bias of 0.032, which is three times lower than the hypothetical lower left-digit bias of 0.1 (row 5). However, these models fare worse in terms of fit than those that allow for partial incorporation of the bias. The better fit is driven by allowing for a positive mass with low price endings. In rows (2)–(4), I find that a good description of pricing is one in which the left-digit-bias-led pricing is incorporated in 50–60% of pricing instances and ignored in the other 50–40%. Importantly, estimates show that even in instances when the bias is taken into account, the perceived bias is

33. I take all combinations of $\{bottom - cents\} \times \{top - cents\} \times \{lowest - dollar\} \times \{highest - dollar\} \times \{polynomial - degree\} = \{0.29, 0.59\} \times \{0.89\} \times \{\$1, \$2\} \times \{\$5, \$7\} \times \{5, 7\}$

34. Alternative measures of goodness-of-fit, by ranking per estimation, are displayed in Supplementary Material, Appendix Table A-2. The ranking remains unchanged.

35. The force that pushes prices up is the lower elasticity within a dollar-digit price change, while the force that pushes prices down is the discontinuity in demand at round prices. Both of these forces exist due to θ regardless of Δ even though their magnitude is affected to some degree. See Proposition 1 for more detail.

much lower than the true one. For example, assuming an elasticity of -1.5, the incorporation frequency is 61% , but the perceived bias conditional on incorporation is 0.017, or eleven times lower than the true bias (Table 3, row 2). Finally, the pricing pattern that best fits the data is the "rule of thumb" pricing, in which firms use ninety-nine-ending prices randomly, *i.e.*, not based on profit maximization. It is the best fitting model on average (Table 3, last column), and case by case (Supplementary Material, Appendix Table A-2). It is best fitting because the price distribution is more symmetric around ninety-nine- or forty-nine-ending prices. In 40% of pricing instances, a company will round to the nearest ninety-nine or forty-nine (only counterfactual prices ending between forty-five and fifty-five are rounded to forty-nine and the rest to ninety-nine. See Supplementary Material, Appendix Figure A-6); in the rest, it will price as if demand is smooth.

Heterogeneity between chains. Interestingly, there is no substantial heterogeneity of pricing procedures and perceived parameters between the twenty-five chains. Results by chain are shown in Supplementary Material, Appendix Figures A-6 and A-7. This is perhaps surprising given the large documented between-chain differences in price levels (Hitsch *et al.*, 2017; DellaVigna and Gentzkow, 2019) and pricing style (Ellickson and Misra, 2008). While there is variation in the magnitude of the perceived bias and incorporation frequency, they do not differ in a conceptually meaningful way. All firms, except perhaps "9-874", strongly underestimate the bias and only incorporate it partially, with similar magnitudes. Figure 7 shows the average left-digit bias faced by a chain on the horizontal axis against the average perceived left-digit bias on the vertical one. Even under the assumption of highly elastic demand, and allowing for partial incorporation, the perceived bias (conditional on incorporation) is well below the 45-degree line, with a weak slope between chains. Regardless of the left-digit bias a retailer faces—and some face much stronger bias than others—they all price as if the bias is about an order of magnitude smaller than it actually is, and this behavior is similar across chains. Furthermore, at the chain level, twenty out of twenty-five chains are characterized with having "rule of thumb" as the best fitting pricing procedure. Comparing goodness-of-fit of non-nested models does not show which model is "correct", but it is informative that the "rule of thumb" is frequently selected as better fitting the data, driven by the symmetry of shares of non-ninety-nine-ending prices.

In summary, firms' pricing reveals that they under-respond to the bias and pricing mistakes are weakly reduced by incentives. These findings are driven by the propensity to use low-ending prices and forty-nine-ending prices. The analysis quantifies chains' pricing procedures, allowing to study the implications of actual pricing against the estimated structure of demand. The next section studies these implications.

4. EFFECTS ON PROFITS AND WELFARE

Left-digit bias and firms' response to left-digit bias jointly affect welfare. Left-digit bias of consumers has direct effects on *demand* conditional on prices, and the perceived left-digit bias of firms has further effects on *pricing* itself. Together, these forces affect firm profits, consumer surplus, and deadweight loss.

I examine two counterfactuals to study these effects: the effects of left-digit bias levels, and the effects of the *sub-optimal* pricing procedures. First, I study the effects of the magnitude of left-digit bias itself, assuming firms optimize. Second, I fix the bias level and study the consequences of underestimation of the bias by the pricing firm. The main findings from these exercises are that, (1) through its effect of lowering prices, left-digit bias has ambiguous effects on welfare, and might even increase consumer surplus and lead to a more efficient outcome;

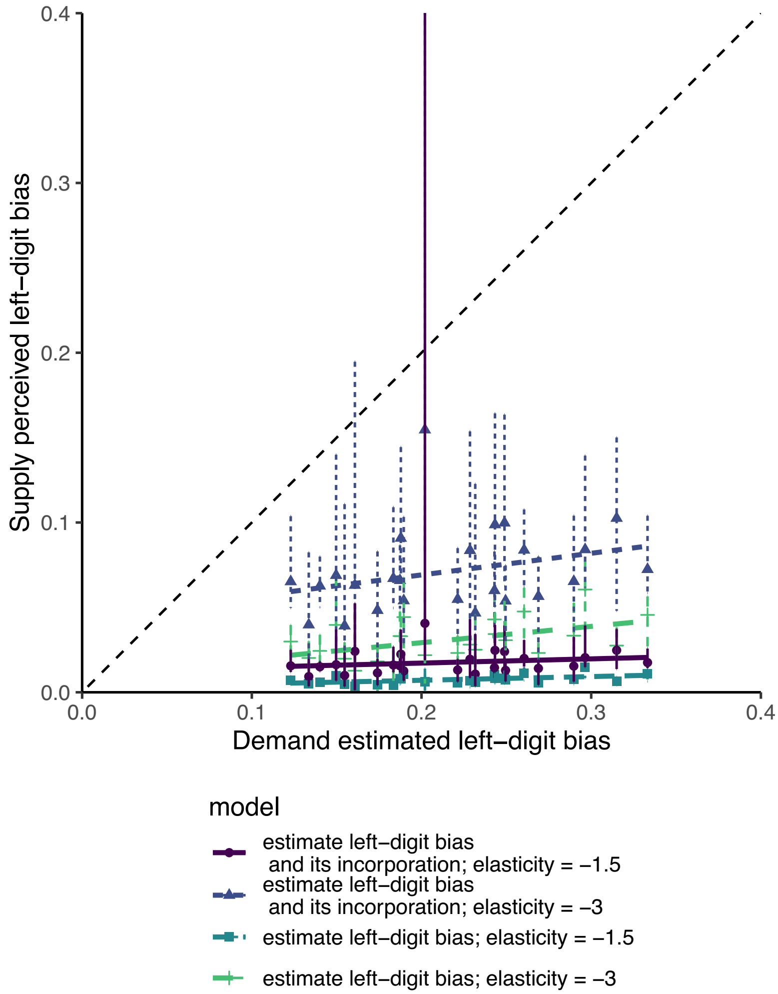

FIGURE 7

Estimated supply-perceived LDB versus estimated LDB

Notes: The figure plots firms' estimated perceived left-digit bias against the estimated left-digit bias from demand. The dotted line represents a 45-degree line. The plot displays estimations using four different assumptions regarding the pricing procedure, each represented by a different shape and colour. Each point is the estimated perceived left-digit bias for each chain by specification. Vertical lines span the 2.5% and 97.5% estimates from 300 cluster-bootstraps of the moments at the UPC-chain level.

(2) firms lose a few percent of gross profits by underestimating the bias, with small effects on consumer surplus.

The first-order effects of left-digit bias on demand levels are driven by the non-identified Δ (recall, the perceived price is $\hat{p} = (1-\theta)p + \theta(\lfloor p \rfloor + \Delta)$). To illustrate why, consider the cases of $\Delta = 0$ and $\Delta = 0.99$. If $\Delta = 0$, then perceived prices are lower, and hence demand and profits are higher everywhere; while if $\Delta = 0.99$, demand and profits are lower. However, Δ cannot be inferred from the demand nor supply estimations. Since this is an untestable assumption, I focus on comparative statistics conditional on various levels of Δ, and highlight those that hold regardless of its level.

Unlike the demand effects, *pricing* decisions are insensitive to Δ, and are a function of what the firm perceives left-digit bias to be, *i.e.*, $\hat{\theta}$.[36] Conditional on the perceived left-digit bias $\hat{\theta}$, the true bias, θ, does not affect prices. A price governed by $\hat{\theta} > 0$ versus $\hat{\theta} = 0$ can be higher (*e.g.*, 4.29 instead of 4.20) or lower (*e.g.*, 3.99 instead of 4.20). These mixed movements of prices are key to understanding the overall effects.

Together, the demand and pricing effects lead to distortions in consumers' choice and firms' performance.

4.1. *The effects of left-digit bias* ($\theta > 0$)

Consumers' left-digit bias creates opportunities (or constraints if $\Delta = 0.99$) for the firm since it distorts perceived prices, but comes at a cost for consumers as consumption distortions. The net effect, on both consumers and deadweight loss, is inherently undetermined due to the often opposing effects of excessive consumption at lower prices. Supplementary Material, Appendix G, derives these quantities.

The effects on prices, consumer surplus, and deadweight loss are illustrated in Figure 8. The figure shows equilibrium prices and quantities comparing $\theta > 0$ to $\theta = 0$. The key idea in interpreting the figure is that welfare is governed by the *true* demand curve, prices, and costs even if prices are perceived with distortion. Assume first that $\Delta = 0$ as in Figure 8, panel A. The no-bias price point is x, while under a bias of $\theta > 0$, the equilibrium price can either be the ninety-nine-ending price below, y, or a higher price, y'. In the former, consumers enjoy the lower price $p(y)$ and purchase more of it, leading to an increase in consumer surplus represented by the grey trapezoid; but consumers over-consume because $\hat{p}(y) < p(y)$, leading to a transfer of surplus from consumers to firms, represented by the triangle. Therefore the overall consumer surplus effect of a lower price depends on which effect dominates, which is a function of parameters and cost distribution. However, the price may actually go up, as in y', and in which case, consumers are unequivocally worse off. Supplementary Material, Appendix Figure A-8, shows that the average price goes down with stronger left-digit bias. Deadweight loss is also ambiguous in sign since the overall quantity sold increases for prices at ninety-nine-ending prices (as in y) but decreases if prices increase due to the bias (as in y').[37]

Importantly, Δ interacts with the effects of price distortions. To see why, note that a stronger θ leads to more ninety-nine-ending prices. But if $\Delta = 0$ then a ninety-nine-ending price is the most distorted perceived price, while if $\Delta = 0.99$ then a ninety-nine-ending price is perceived correctly.

That is, two forces create ambiguity about the effects of stronger left-digit bias. First, the effect of left-digit bias reducing the price level may counteract the increased distortion. Second, Δ interacts with the price-level effects. Therefore, the overall effects of an increase in left-digit bias are undetermined.

4.2. *The effects of firms underestimating the bias* ($\hat{\theta} < \theta$)

Now consider the findings of Section 3, in which firms underestimate the bias and follow other pricing procedures. Quantifying these effects is crucial because a plausible explanation of the

36. The effects in Figure 8, panel B, when $\Delta = 0.99$, are similar. Moving from x to y leads to an increase in consumer surplus unambiguously because the ninety-nine-ending perceived price is closest to the truth. Still, a change from x to y' leads to a decrease.

37. As Butters *et al.* (2019) argue for lack of seasonal adjustments.

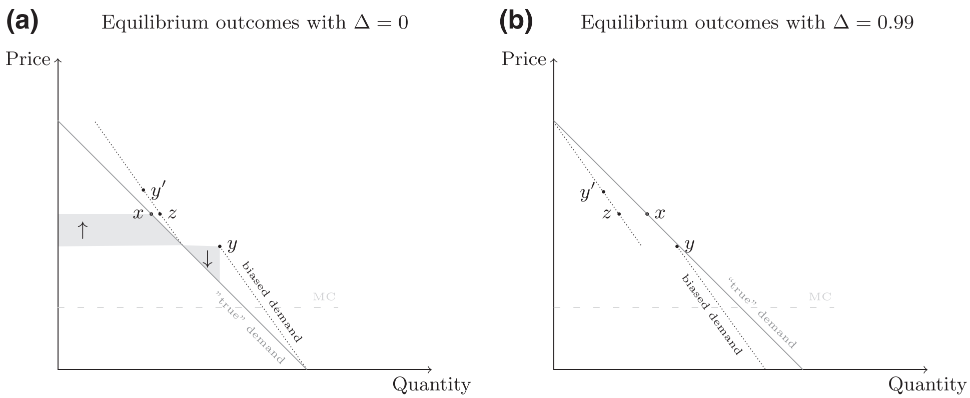

FIGURE 8

Illustration of welfare effects. (a) Equilibrium outcomes with $\Delta = 0$. (b) Equilibrium outcomes with $\Delta = 0.99$

Notes: The figures describe the optional equilibrium quantities and prices under different levels of left-digit-bias and firm perception of the bias. Point x is the solution in the standard no-bias case, on the solid grey no-bias demand line. The no-bias demand line captures the true underlying valuation of the product. In contrast, the dotted lines show the demand under the perceived price $\hat{p} = (1-\theta)p + \theta(\lfloor p \rfloor + \Delta)$. Panel A shows the biased demand when the focal price-ending is 0, and hence demand is higher everywhere; while panel B shows the biased demand when the focal price ending is ninety-nine. Points y and y' describe the new possible equilibrium outcomes when there is left-digit bias $\theta = \hat{\theta} > 0$. Point y is the outcome when the optimal price is the lower ninety-nine-ending price, and point y' is the outcome for interior price updating. Finally, point z is the outcome when there is bias $\theta > 0$, but the firm prices as if there is no bias, $\hat{\theta} = 0$. Consumer surplus is the area between true demand and the price for all quantity sold. Over-consumption (as in point y of panel A), where the price is higher than the true demand, implies a negative consumer surplus. Producer surplus, or firm profits, is the area between the price and marginal cost for all quantity sold. Deadweight loss is the area between true demand and marginal cost for all unsold quantity in which demand is higher than cost.

existence of these pricing procedures is that they are not consequential to the bottom line.[38] Fundamentally, if the firm ignores the bias, then there are no price effects but only consumption distortions for consumers. Pure distortions, of course, unambiguously harm consumers (point z in Figure 8). In contrast, an underestimated bias that is not zero also leads to price changes. As in the previous subsection, the effects on consumer surplus and deadweight loss are mixed.

By definition, pricing at the non-optimal price lowers profits. I simulate two sets of prices—optimal and actual. For optimal prices I simulate prices based on the model with estimates from the demand side as "the truth". For the actual pricing simulations I use the descriptive pricing procedures estimates from Section 3.4, with $\hat{\Theta}$ as the set of pricing procedures parameters of perceived left-digit bias, partial incorporation, or rule-of-thumb parameters. I compare profits between the optimal and actual price distributions. Namely, I calculate firm profits under the assumptions that price setting $p(\ldots, \hat{\Theta})$ is governed by pricing procedures $\hat{\Theta}$, and demand $D(\ldots, \theta)$ is determined by the true bias θ (I further assume that the firm knows the true price elasticity ϵ and focal price ending Δ). Profits are then demand times price minus costs, integrated over the cost and parameters distributions:

$$\Pi = \int D\left(p\left(c; \epsilon, \Delta, \hat{\Theta}\right); \epsilon, \Delta, \theta\right) \cdot \left(p\left(c; \epsilon, \Delta, \hat{\Theta}\right) - c\right) \mathrm{d}F\left(c, \epsilon, \theta, \Delta, \hat{\Theta}\right)$$

38. For example, consider two products with unit costs of \$1 and \$1.25, price elasticity of −2, and a bias of 0.15. Absent the bias, the products should be priced at \$2.00 and \$2.50, respectively, but with the bias both should be priced at \$1.99. The gains of changing a naive \$2 to an optimal \$1.99 are 15.5% in profits, while the gains from changing a naive \$2.50 to an optimal \$1.99 are much lower, at 2.6%.

Table 4 shows gross profits lost by pricing according to $\hat{\Theta}$ relative to the profits generated by pricing according to θ. The results are shown for different distributions of elasticity, left-digit bias, incorporation, rule-of-thumb pricing, and perceived left-digit bias parameters. Panel A shows profits lost under different pricing procedures corresponding to the estimated scenarios in Table 3 (and assuming $\Delta = 0$). Profits are 0.6–4.1% lower relative to the case in which the firm prices according to the true parameters. In panel B, I include additional scenarios, allowing for heterogeneity of the true or perceived left-digit bias and elasticity. I also consider the maximal loss in which the firm were to ignore left-digit bias in its pricing. Finally, I consider the effects when a firm sells two substitute products with a cross-elasticity of 0.1. These scenarios lead to similar magnitudes of relative lost profits. In Supplementary Material, Appendix Table A-3, panel A, I repeat the same exercise for the case in which there is additional left-digit bias on the dime digit versus the cent digit (and assume the firm knows that); in panel B, I repeat the exercise for the case of $\Delta = 0.99$, meaning that prices are perceived as higher than they are. The effects on profits are largely unaffected due to additional bias, and they are somewhat weaker if consumers round up but are still on the same magnitude (losses of 0.74–2.2%). That is, under a broad range of assumptions, chains forgo a few percent of profit because their pricing does not match the demand structure they face.

The effects on profits, of a few percent of loss, are substantial in magnitude. If, as in Montgomery (1997), operating margins are about 12%, it means that the firm is losing 5–34% of operating profits by underestimating the bias. While these numbers are large, the conclusions are supported in other papers. Following circulation of a previous version of this paper, Hilger (2018) conducted the same empirical exercise on proprietary online subscriptions of private vendors and finds similar effect sizes, and (List *et al.*, 2021) find similar results for a large marketplace firm.

Table 4 and Supplementary Material, Appendix Table A-3, also show that effects on consumer surplus are potentially large but depend on the unobservable Δ. Therefore, the sign of the effects on consumer surplus is ambiguous. It seems that for the assumed parameters and costs, the effects are negative and large in the most part. However, when consumers round up (panel B of Supplementary Material, Appendix Table A-3), the sign indeed flips. Since Δ is not identifiable, the effects on consumer surplus remain unknown. Nonetheless, it is interesting that consumers are not necessarily harmed by firms responding to their biases, nor necessarily helped by firms under-responding to their biases.

While the model of monopolistic pricing is simple, the exercise is internally consistent and shows low sensitivity to many assumptions. The magnitude of the effects is robust and economically meaningful.

4.3. *Discussion*

How can it be that firms concurrently respond to the bias and yet stop short of fully optimizing?

One alternative is that there is no discrepancy, and the model is wrong. For example, if the bias is endogenous to pricing, or if the model is missing something fundamental about market conduct. Another alternative is that firms are making a consistent mistake even though they try to optimize.

One limitation of the model is the assumption of exogeneity of parameters, assuming that consumers will not be differently biased if the price distribution changes. Instead, perhaps, θ and Δ are endogenous to the price distribution such that if all prices were to end at ninety-nine as the estimation suggests, the bias would have become much lower. Note, however, that from a rational inattention standpoint the opposite should hold for θ—if all prices end with ninety-nine, there is no point in paying attention to the rightmost digits and $\theta \rightarrow 1$. In contrast, the focal price

TABLE 4
Profits lost when the firm underestimates left-digit bias

		Assumed scenario Elasticity (true and perceived)	Perceived LDB incorporation	Frequency of LDB	True LDB	Rule-of-thumb frequency	Relative profits	Relative CS
Panel A: Matched to estimated models								
(1)	Rule-of-thumb pricing at ninety-nine or forty-nine	−1.5			0.20	0.4	−4.1	−5.30
(2)	Partial incorporation of underestimated LDB	−1.5	0.017	0.61	0.20		−3.7	−5.40
(3)	Partial incorporation of underestimated LDB	−3.0	0.070	0.60	0.10		−1.2	−0.70
(4)	Partial incorporation of true LDB	−3.0	0.100	0.50	0.10		−1.4	−0.32
(5)	Underestimated LDB	−3.0	0.032	1.00	0.10		−0.6	−0.59
(6)	Underestimated LDB	−1.5	0.008	1.00	0.20		−3.2	−5.40
(7)	Partial incorporation of true LDB	−1.5	0.200	0.29	0.20		−3.9	−3.90
Panel B: Additional scenarios								
(8)	Heterogenous true LDB	−1.5	0.017	0.60	N (0.2,0.05)		−4.1	−5.40
(9)	Heterogenous underestimated LDB	−1.5	N (0.017,0.005)	0.60	0.20		−3.7	−5.40
(10)	Heterogenous elasticity	U [−1.9, −1.1]	0.017	0.60	0.20		−3.7	−6.00
(11)	Ignoring LDB	−1.5	0	1.00	0.20		−5.7	−5.50
(12)	Ignoring LDB	−3.0	0	1.00	0.10		−3.1	−1.10
(13)	Partial incorporation of underestimated LDB (multiproduct)	−1.5	0.017	0.61	0.20		−3	
(14)	Underestimated LDB (multiproduct)	−1.5	0.0075	1.00	0.20		−2.5	
(15)	Ignoring LDB (multiproduct)	−1.5	0	1.00	0.20		−4.9	

Notes: This table presents percent of profits and consumer surplus relative to optimal pricing procedure (the two right-most columns). Analysis is done assuming $\Delta = 0$, meaning that prices are always perceived as lower than they are. Each row describes a different assumed scenario as detailed by the left 6 columns. For example, row (1) shows that if the true and perceived elasticity is −1.5 and true left-digit bias is 0.2, but a firm prices according to a rounding rule-of-thumb for a random 40% of prices, it loses 4.1% of potential gross profits with 5.3% lower consumer surplus relative to if it were to price optimally given the true price elasticity and left-digit bias. Panel A rows (1)–(7) match rows (1)–(7) of Table 3. Panel B provides additional scenarios: rows (8)–(10) assume that the true left-digit bias, perceived left-digit bias, and elasticity are heterogeneous. Rows (11) and (12) provide the consequences of completely ignoring left-digit bias in pricing, when the truth is either elasticity of −1.5 and left-digit bias of 0.2 (as in the data, row 11), or −3 and 0.1 (extreme case, row 12). Finally, rows (13)–(15) provide similar numbers for a monopolist that prices two products with a cross-elasticity of 0.1. Supplementary Material, Appendix Table A-3, shows these effects for the same set of scenarios under different assumptions of multiplicative left-digit bias and for $\Delta = 0.99$.

ending Δ is more likely to change in this alternative model. That is, if the price distribution shifts toward high price-endings, the focal price-ending Δ will increase and in turn suppress demand. Yet, it is reasonable to assume that the firm is a "Δ-taker", meaning that the focal price-ending is not store-contingent. In this case, and even in the case in which firm's own pricing has a low effect on Δ, the optimal pricing is as stated above. Indeed, even large sophisticated firms selling only a handful of products, such as smartphones (Samsung, Apple), or subscriptions (Netflix, Amazon Prime), use ninety-nine-ending prices for their small set of products. If there were cases of endogenous Δ, these will be it, which should have led to a different distribution of price endings. Finally, left-digit bias might be driven by long-term human learning processes of mental number representation and encoding, which are also unlikely to be affected by a few firms' pricing decisions. However, I do not attempt to predict what will happen if all prices everywhere were to end in ninety-nine. That is, I do not argue that if all firms were to follow the paper's recommendations left-digit bias would remain the same, but I do argue that in the current situation individual firms are making a costly mistake.

Another explanation is that the optimal pricing model is simplistic. First, the model simplifies competitive forces to enter through the price elasticity. A model of pure price competition with left-digit biased consumers also leads to missing prices as shown in Supplementary Material, Appendix B, but it is hard to take to the data. Insofar as elasticity is a good approximation for the strength of competition between retailers, it does reduce the ranges of missing prices. However, as shown in Section 3.4, we need to believe that chains think demand is extremely elastic (many times more than reflected in demand estimation) in order to think that firms price according to the correct level of left-digit bias. Furthermore, since optimal response to the bias lowers the average price level (see Supplementary Material, Appendix Figure A-8), the optimal response does not hinder between-chain competition over prices. A second limitation regards substitution between products within a store. Consider the case of a multi-product monopolist, as retailers are. Keeping same-price elasticity and left-digit bias fixed, substitution between products leads to *higher* next-lowest prices. This result holds whether the other product is more or less elastic, and whether its cost is higher or lower (see Supplementary Material, Appendix C, which solves a two-product monopolist problem and simulates price distributions). The conclusion is that adding cross elasticities to the counterfactual analysis is likely to make the discrepancy between the predicted and the observed retailers' behaviour even *stronger*, though possibly making the effects on profits somewhat smaller (rows 13–15 in Table 4).

The third explanation, which I find more plausible, is that firms follow heuristics rather than purely optimizing. I was able to interview three chief executives in various retail chains. They were in charge of pricing and pricing strategy in several of the United States' largest chains. Although their pricing strategies were quite different, their reasoning of ninety-nine-pricing was similar: There is a belief in the industry that \$2.99 is perceived as less than \$3.00, and although they themselves are not sure if that is correct, there is a lot of "inertia"/"legacy"/"copycat behaviour". In terms of the model, they do believe there may be some left-digit bias, but think it is not meaningful and hence do not incorporate it quantitatively into their pricing decisions. Their descriptions of what they actually do are more heuristic in nature: when the "absent-bias" desired price point is close to ninety-nine, they sometimes round to ninety-nine, or if their own past price or competitors' prices are two-something, they will not want to cross to three-something. These lines of reasoning imply a heuristic approach to pricing.

These pricing heuristics can create bunching and missing prices while being consistently wrong. Indeed, prices are more easily explained by the partial incorporation or rule-of-thumb pricing procedures. To see how heuristics are not enough, note that it makes intuitive sense that an absent-bias price of \$2.00 should be adjusted to \$1.99, but it might be less clear how to adjust \$2.50. Furthermore, not only is it more intuitive to make small changes in prices, but the impact on profits is larger for the small changes which makes complete experimentation and learning less likely. So, given the counter-intuitive nature of big adjustments in price, and considering the relatively lower gains of big adjustments (masked in fluctuations in demand in real data), it seems reasonable that firms would stop shy of full adjustment. Making the full adjustment requires either "brave" and large scale experimentation with prices, or a rigorous analysis with a quantitative model-based decision making.

A companion paper (Strulov-Shlain, 2019a) analyses a policy reform and finds support for fundamental misunderstanding of left-digit bias. Following a reform, firms had to change prices and became increasingly likely to use low-ending prices. The main argument is that firms try to optimize but stop short of full optimization, probably due to partial learning. Partial learning can sustain either sub-optimal model-free decision making or beliefs in a wrong model used for decision making in the long-run.

5. CONCLUSIONS

In this paper I argue that a model of left-digit biased consumers and partially-optimizing firms offers a better description of retail data than the standard assumptions over unbiased consumers and sophisticated firms. Consumers exhibit substantial left-digit bias in everyday choices while shopping at supermarkets. The supermarkets, in response, are using ninety-nine-ending prices which indeed lead to higher profits. However, retailers also seem to underestimate the extent of left-digit bias and misunderstand its source and structure. I find that firms do better than if they were to price as if there is no bias at all, but behave far from optimally. I estimate that the chains lose 0.6–4% of gross profits from such deviations with ambiguous effects on consumer surplus.

One limitation of the study is that it is confined to a specific setting and price level—single dollar-digit items in U.S. supermarket chains. However, I cover a large and arguably representative sample of products and chains. Left-digit bias was found in other domains and price levels (*e.g.*, Repetto and Solis, 2018; Jones and Strulov-Shlain, 2021, for tens- and hundreds-of thousands and millions), but it is indeed rare to observe high prices ending with ninety-nine cents. A study that allows for finer, less parametric, treatment of digits can be interesting for future research.

Another limitation is that the data used have clear issues—no direct measures of prices, no random variation, and no information on what prices were displayed to consumers among other marketing techniques. I believe that the data cleaning procedures deal with the systematic concerns, and that noise is absorbed by the large number of estimates. Furthermore, a previous version (Strulov-Shlain, 2019b) used completely different data and reached similar estimates of left-digit bias, and so did a separate group of authors using data from an online marketplace (List *et al.*, 2021).

Left-digit bias is an acknowledged yet underestimated force in the academic and public discourse. Its effects go beyond prices and profits documented here and elsewhere. Left-digit bias represents a distortion of number processing, and since numbers are ubiquitous to decision making in multiple domains, it has far and wide effects. For example, because left-digit bias leads to price discretization, it can cause price stickiness (Nakamura and Steinsson, 2008; Eichenbaum *et al.*, 2014) and affect tax pass-through (Conlon and Rao, 2016). Furthermore, since it affects number perception in general, left-digit bias distorts negotiation outcomes (Jiang, 2020), affects education attainment via tests retakes (Goodman *et al.*, 2020), and affects credit scores via credit monitoring (Fong and Hunter, 2022).

The existence and magnitude of left-digit bias, found in this and other recent papers, is now well documented but there are some promising avenues for future research. First, Δ is an unknown parameter, with first-order effects on welfare. Second, future research can dive deeper into what might affect the bias magnitude (*e.g.*, Sokolova *et al.*, 2020), or use left-digit bias as a tool to study other phenomena (*e.g.*, the examples above, or Strulov-Shlain, 2019a on how firms learn and make decisions).

This paper serves as an example of large and experienced firms making a persistent mistake. Chains set prices that are potentially driven by wrong beliefs about the parameters of demand, but more likely because they have the wrong model in mind, if they have any model at all. The assumption of firms' optimization should be further scrutinized.

Acknowledgments. I am thankful and grateful to Stefano DellaVigna for endless encouragement and support. Ned Augenblick, Zarek Brot-Goldberg, Ben Handel, Ganesh Iyer, Ulrike Malmendier, and Dmitry Taubinsky provided many thoughtful comments. Christy Kang and Anthony Adam Sy provided superb research assistance. I have also immensely benefited from discussions with Eric Avis, Bryan Chu, Natalie Cox, Josh Dean, Jean-Pierre Dube, Daniel Haanwinkel, Yoram Halevy, Johannes Hermle, Peter Jones, Shachar Kariv, Kei Kawai, Jennifer Kwok, Julien Lafortune, Nick Li, John Loeser, Filip Matejka, Sanjog Misra, Ted O'Donoghue, Waldo Ojeda, Devin Pope, Matthew Rabin, Gautam Rao, Alex Rees-Jones, Emmanuel Saez, Jon Schellenberg, Yotam Shem-Tov, David Sraer, Alex Steiny, Justin Sydnor, Jonas Tungodden, J. Miguel Villas-Boas, from the comments of referees, and probably many others who I hope will forgive my imperfect memory or forgot themselves. I am thankful for Liran Einav, Nir Jaimovich, Art Middlebrooks, and Sofia Villas-Boas in assisting me with data collection. Researcher's own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analysing and preparing the results reported herein. The code underlying this research is available on Zenodo at https://doi.org/10.5281/zenodo.7075310.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

Researcher's own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business (NielsenIQ, 2020). The code underlying this research is available on Zenodo at https://doi.org/10.5281/zenodo.7075310 and provides all figures, numbers, and tables in this manuscript from the raw data. Raw data are widely available through the Kilts Center for Marketing, via institutional agreements.

REFERENCES

AALTO-SETALA, V. (2005), "How Do Markets Behave? The Adjustment of Price Endings", *Journal of Product & Brand Management*, **14**, 455–459.

ANDERSON, E., JAIMOVICH, N. and SIMESTER, D. (2015), "Price Stickiness: Empirical Evidence of the Menu Cost Channel", *Review of Economics and Statistics*, **97**, 813–826.

ANDERSON, E. T. and SIMESTER, D. I. (2003), "Effects of $9 Price Endings on Retail Sales: Evidence from Field Experiments", *Quantitative Marketing and Economics*, **1**, 93–110.

ASHTON, L. (2014), "Left-digit Bias and Inattention in Retail Purchases: Evidence from a Field Experiment" (Available at SSRN: https://ssrn.com/abstract=2538816).

ATER, I. and GERLITZ, O. (2017), "Round Prices and Price Rigidity: Evidence from Outlawing Odd Prices", *Journal of Economic Behavior & Organization*, **144**, 188–203.

BACKUS, M., BLAKE, T. and TADELIS, S. (2019), "On the Empirical Content of Cheap-talk Signaling: An Application to Bargaining", *Journal of Political Economy*, **127**, 1599–1628.

BASU, K. (1997), "Why Are So Many Goods Priced to End in Nine? And Why This Practice Hurts the Producers", *Economics Letters*, **54**, 41–44.

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

BIZER, G. Y. and SCHINDLER, R. M. (2005), "Direct Evidence of Ending-digit Drop-off in Price Information Processing", *Psychology and Marketing*, **22**, 771–783.

BLATTBERG, R. C. and WISNIEWSKI, K. J. (1988), *Issues in Modeling Store Level Scanner Data*, University of Chicago, Center for Research in Marketing.

BRAY, J. P. and HARRIS, C. (2006), "The Effect of 9-ending Prices on Retail Sales: A Quantitative UK based Field Study", *Journal of Marketing Management*, **22**, 601–617.
BUSSE, M. R., LACETERA, N., POPE, D. G., SILVA-RISSO, J. and SYDNOR, J. R. (2013), "Estimating the Effect of Salience in Wholesale and Retail Car Markets", *American Economic Review*, **103**, 575–579.
BUTTERS, R., SACKS, D. W. and SEO, B. (2019), "Why Don't Retail Prices Vary Seasonally with Demand?" (Kelley School of Business Research Paper).
CARVER, J. R. and PADGETT, D. T. (2012), "Product Category Pricing and Future Price Attractiveness: 99-ending Pricing in a Memory-based Context", *Journal of Retailing*, **88**, 497–511.
CHEN, Y., IYER, G. and PAZGAL, A. (2010), "Limited Memory, Categorization, and Competition", *Marketing Science*, **29**, 650–670.
CHO, S. and RUST, J. (2008), "Is Econometrics Useful for Private Policy Making? A Case Study of Replacement Policy at an Auto Rental Company", *Journal of Econometrics*, **145**, 243–257.
CHO, S. and RUST, J. (2010), "The Flat Rental Puzzle", *The Review of Economic Studies*, **77**, 560–594.
CONLON, C. T. and RAO, N. (2020), "Discrete Prices and the Incidence and Efficiency of Excise Taxes", *American Economic Journal: Economic Policy*. 12(4), 111–143.
DALRYMPLE, D. J. and HAINES JR, G. H. (1970), "A Study of the Predictive Ability of Market Period Demand–Supply Relations for a Firm Selling Fashion Products", *Applied Economics*, **1**, 277–285.
DELLAVIGNA, S. (2009), "Psychology and Economics: Evidence from the Field", *Journal of Economic Literature*, **47**, 315–372.
DELLAVIGNA, S. and GENTZKOW, M. (2019), "Uniform Pricing in US Retail Chains", *The Quarterly Journal of Economics*. 134(4), 2011-2084.
DUBE, A., MANNING, A. and NAIDU, S. (2017), "Monopsony and Employer Mis-optimization Account for Round Number Bunching in the Wage Distribution" (Unpublished Manuscript).
EICHENBAUM, M., JAIMOVICH, N., REBELO, S. and SMITH, J. (2014), "How Frequent are Small Price Changes?", *American Economic Journal: Macroeconomics*, **6**, 137–155.
ELLICKSON, P. B. and MISRA, S. (2008), "Supermarket Pricing Strategies", *Marketing Science*, **27**, 811–828.
EL SEHITY, T., HOELZL, E. and KIRCHLER, E. (2005), "Price Developments After a Nominal Shock: Benford's Law and Psychological Pricing After the Euro Introduction", *International Journal of Research in Marketing*, **22**, 471–480.
FARHI, E. and GABAIX, X. (2015), "Optimal Taxation with Behavioral Agents" (Technical Report, National Bureau of Economic Research).
FONG, J. and HUNTER, M. (2022), "Can Facing the Truth Improve Outcomes? Effects of Information in Consumer Finance", *Marketing Science*. 41(1), 33-50.
GABAIX, X. (2014), "A Sparsity-based Model of Bounded Rationality", *The Quarterly Journal of Economics*, **129**, 1661–1710.
GABAIX, X. (2019), "Behavioral Inattention", in Bernheim, B. D., DellaVigna, S. and Laibson, D. (eds) *Handbook of Behavioral Economics – Foundations and Applications 2* (Vol. 2, Chap. 4, North-Holland) 261–343. doi:10.1016/bs.hesbe.2018.11.001.
GINZBERG, E. (1936), "Customary Prices", *The American Economic Review*, **26**, 296–296.
GOLDFARB, A. and XIAO, M. (2011), "Who Thinks About the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets", *American Economic Review*, **101**, 3130–3161.
GOODMAN, J., GURANTZ, O. and SMITH, J. (2020), "Take Two! SAT Retaking and College Enrollment Gaps", *American Economic Journal: Economic Policy*, **12**, 115–158.
HACKL, F., KUMMER, M. E. and WINTER-EBMER, R. (2014), "99 Cent: Price Points in e-commerce", *Information Economics and Policy*, **26**, 12–27.
HANNA, R., MULLAINATHAN, S. and SCHWARTZSTEIN, J. (2014), "Learning Through Noticing: Theory and Evidence from a Field Experiment", *The Quarterly Journal of Economics*, **129**, 1311–1353.
HASTINGS, J. and SHAPIRO, J. M. (2018), "How are SNAP Benefits Spent? Evidence from a Retail Panel", *American Economic Review*, **108**, 3493–3540.
HAUSMAN, J. A. (1996), "Valuation of New Goods Under Perfect and Imperfect Competition", in *The Economics of New Goods* (University of Chicago Press) 207–248.
HEIDHUES, P. and KÖSZEGI, B. (2018), "Behavioral Industrial Organization", in Bernheim, B. D., DellaVigna, S. and Laibson, D. (eds), *Handbook of Behavioral Economics – Foundations and Applications 1* (Vol. 1, Chap. 6, North-Holland) 517–612. doi:10.1016/bs.hesbe.2018.07.006.
HILGER, N. (2018), "Heuristic Thinking in the Market for Online Subscriptions" (Available at SSRN 3296698).
HITSCH, G. J., HORTACSU, A. and LIN, X. (2017), "Prices and Promotions in US Retail Markets: Evidence from Big Data", *Chicago Booth Research Paper* 17–18.
JIANG, Z. (2020), "An Empirical Bargaining Model with Left-digit Bias—A Study on Auto Loan Monthly Payments" (Available at SSRN 3445171).
JONES, P. and STRULOV-SHLAIN, A. (2021), "Is Left-digit Bias Intentional?". mimeo.
KALYANAM, K. and SHIVELY, T. S. (1998), "Estimating Irregular Pricing Effects: A Stochastic Spline Regression Approach", *Journal of Marketing Research*, **35**, 16–29.
KLEVEN, H. J. and WASEEM, M. (2013), "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan", *The Quarterly Journal of Economics*, **128**, 669–723.

LACETERA, N., POPE, D. G. and SYDNOR, J. R. (2012), “Heuristic Thinking and Limited Attention in the Car Market”, *American Economic Review*, **102**, 2206–2236.
LEVY, D., LEE, D., CHEN, H., KAUFMANN, R. J. and BERGEN, M. (2011), “Price Points and Price Rigidity”, *Review of Economics and Statistics*, **93**, 1417–1431.
LIST, J., MUIR, I., POPE, D. and SUN, G. (2022), “Unveiling the Law of Demand using a Large-scale Natural Field Experiment”. *Review of Economic Studies*.
MACÉ, S. (2012), “The Impact and Determinants of Nine-ending Pricing in Grocery Retailing”, *Journal of Retailing*, **88**, 115–130.
MATĚJKA, F. (2015), “Rationally Inattentive Seller: Sales and Discrete Pricing”, *The Review of Economic Studies*, **83**, 1125–1155.
MONTGOMERY, A. L. (1997), “Creating Micro-marketing Pricing Strategies using Supermarket Scanner Data”, *Marketing Science*, **16**, 315–337.
MORRISON, W. and TAUBINSKY, D. (2019), “Rules of Thumb and Attention Elasticities: Evidence from Under-and Overreaction to Taxes” (Technical Report, National Bureau of Economic Research).
NAKAMURA, E. and STEINSSON, J. (2008), “Five Facts About Prices: A Reevaluation of Menu Cost Models”, *The Quarterly Journal of Economics*, **123**, 1415–1464.
NielsenIQ (2020), “Nielseniq Data Provided by the Kilts Center for Marketing 2006–2019”. Accessed on 08-05-2022.
REPETTO, L. and SOLIS, A. (2020), “The Price of Inattention: Evidence from the Swedish Housing Market”. *Journal of the European Economic Association*, 18(6), 3261-3304.
ROSSI, P. E. (2014), “Even the Rich Can Make Themselves Poor: A Critical Examination of IV Methods in Marketing Applications”, *Marketing Science*, **33**, 655–672.
SAEZ, E. (2010), “Do Taxpayers Bunch at Kink Points?”, *American Economic Journal: Economic Policy*, **2**, 180–212.
SCHINDLER, R. M. and CHANDRASHEKARAN, R. (2004), “Influence of Price Endings on Price Recall: A By-digit Analysis”, *Journal of Product & Brand Management*, **13**, 514–524.
SCHINDLER, R. M. and KIBARIAN, T. M. (1996), “Increased Consumer Sales Response Though Use of 99-ending Prices”, *Journal of Retailing*, **72**, 187–199.
SCHINDLER, R. M. and KIRBY, P. N. (1997), “Patterns of Rightmost Digits Used in Advertised Prices: Implications for Nine-ending Effects”, *Journal of Consumer Research*, **24**, 192–201.
SCHINDLER, R. M. and WIMAN, A. R. (1989), “Effects of Odd Pricing on Price Recall”, *Journal of Business Research*, **19**, 165–177.
SNIR, A., LEVY, D. and CHEN, H. A. (2017), “End of 9-endings, Price Recall, and Price Perceptions”, *Economics Letters*, **155**, 157–163.
SOKOLOVA, T., SEENIVASAN, S. and THOMAS, M. (2020), “The Left-digit Bias: When and Why are Consumers Penny Wise and Pound Foolish?”, *Journal of Marketing Research*, **57**, 771–788.
STIVING, M. (2000), “Price-endings When Prices Signal Quality”, *Management Science*, **46**, 1617–1629.
STIVING, M. and WINER, R. S. (1997), “An Empirical Analysis of Price Endings with Scanner Data”, *Journal of Consumer Research*, **24**, 57–67.
STRULOV-SHLAIN, A. (2019a), “Firm Sophistication: A Natural Experiment on Price Endings” (Unpublished Manuscript).
STRULOV-SHLAIN, A. (2019b), “More Than a Penny’s Worth: Left-digit Bias and Firm Pricing” (Chicago Booth Research Paper No. 19-22).
THOMAS, M. and MORWITZ, V. (2005), “Penny Wise and Pound Foolish: The Left-digit Effect in Price Cognition”, *Journal of Consumer Research*, **32**, 54–64.