Discreet Personalized Pricing

Benjamin Reed Shiler, Brandeis University

Working Paper Series

2022| 133

Discreet Personalized Pricing

Benjamin Reed Shiller*

Brandeis University

June 1, 2022

Abstract

Emerging tracking data allow precise predictions of individuals' reservation values. However, firms are reluctant to conspicuously implement personalized pricing because of concerns about consumer and regulatory reprisals. This paper proposes and applies a method which disguises personalized pricing as dynamic pricing. Specifically, a firm can sometimes tailor the "posted" price for the arriving consumer but privately commits to change price infrequently. Note this personalized pricing strategy should arise—possibly unintentionally—through algorithmic pricing when some employed variables reflect characteristics of the arriving consumer. I examine outcomes in four contexts: one empirical and three hypothetical distributions of consumer valuations. While one may expect this strategy to be most profitable for low popularity items, I find, counterintuitively, that this strategy raises profits most for medium popularity products. Moreover, typically observable measures of price discrimination suggest it is most intense for these products. Furthermore, improvements in the precision of individual-level demand estimates raise both the popularity-level where absolute profit gains peak and the range of popularities this strategy can be profitably applied to. I conclude that this is an auspicious strategy for online platforms, if not already secretly in use.

*shiller@brandeis.edu

1 Introduction

Firms undoubtedly search for ways to extract more surplus from consumers. One strategy, personalized pricing, is quite old but has gained renewed attention as consumer tracking technologies have yielded large datasets with detailed information about individual consumers' habits and tastes. Such data have made personalized pricing more profitable, potentially yielding profit gains of around 10% to 50% (Dong et al., 2009; Dube and Misra, 2017; Kehoe et al., 2018; Shiller, 2020; Smith et al., 2021; Shiller and Waldfogel, 2011; Zhang et al., 2014).

However, implementing personalized pricing too overtly risks severe backlash from consumers and policymakers. Personalized pricing is viewed as unfair (Campbell, 1999; Ho and Su, 2009; Kahneman et al., 1986; Li and Jain, 2016; Xia and Monroe, 2010); its use reduces disgruntled consumers' purchase intentions (Anderson and Simester, 2008; Darke and Dahl, 2003; Leibbrandt, 2020) and risks negative publicity. Furthermore, straightforward personalized pricing has been scrutinized by policy makers in the U.S. (e.g., Executive Office of the President 2015), is challenging and perhaps illegal to use in Europe, and is explicitly prohibited by new antitrust laws in China implemented on March 1st, 2022 (Conrad and Knight, 2022).[1,2] Furthermore, the risk of detection is substantial, as a literature has developed with the stated intent of searching for such pricing (Hannak et al., 2014; Hupperich et al., 2018; Iordanou et al., 2017; Mikians et al., 2012).

Although firms are concerned about a possible backlash, they have not abandoned personalized pricing altogether.[3] Instead, firms have focused on how to reframe or obfuscate personalized prices, a theme which emerged as a major topic of discussion at the 2014 National Retail Federation Annual Conference (Khan, 2016). A few years later, 64% of retail survey panelists agreed that "it will become important to deliver personalized prices

[1]In Europe, personalized pricing might violate Article 22 of the General Data Protection Regulation (GDPR) (Wong, 2021). If not, Article 7 still leaves firms with an untenable choice: disclose intent to use collected data to personalize prices or risk detection violating the GDPR.

[2]Cornerstone Research, an economic consulting firm, lists personalized pricing as one of the emerging competition issues for the digital economy: `https://www.cornerstone.com/practices/expertise/digital-economy-antitrust-competition/`, accessed May 31, 2022.

[3]While uniform pricing is still common at brick-and-mortar stores (DellaVigna and Gentzkow, 2019), sophisticated pricing is increasingly used online (Aparicio et al., 2021). Successful startups, such as Precima, have emerged to assist firms implementing personalized pricing: `https://precima.com/insights/prepare-for-the-future-of-personalized-pricing/`.

to shoppers in the next three years" (Baird and Kilcourse, 2017). Documented obfuscation strategies include: (1) reframing personalized pricing as (effortless) customized coupons or discounts (Reimers and Shiller, 2019; Rossi et al., 1996; Shiller, 2020), (2) steering, i.e., personalizing rank-sorting algorithms by promoting more expensive items to price-insensitive consumers (Hannak et al., 2014; Mikians et al., 2012), (3) eliminating posted prices and instead show prices only on smartphones (Allender et al., 2021), and (4) personalizing ad loads (Goli, 2020), thus personalizing the implied "nuisance" cost of visiting a website. However, these methods are either not well disguised or not as effective as hoped at increasing profits, and thus are not widely adopted.

An alternative and thus far understudied strategy is to disguise personalized pricing as dynamic pricing. The basic premise is that the firm can observe the arriving consumer's type before the webpage loads on the consumer's web browser. The firm can decide to raise the "posted price(s)"—not only to that consumer but to other consumers as well—at that exact moment.[4] The price is designed to extract surplus from the consumer that has just arrived. However, by privately committing to keeping the new price for some length of time, making prices sticky, the firm substantially complicates consumers' efforts to detect finely targeted pricing. I call this strategy *sticky targeted pricing.*

Note that this pricing strategy may arise under algorithmic pricing, possibly unintentionally. Suppose a firm maximizes profits by dynamically adjusting the posted price based on many variables, some of which include information reflecting the identity of the arriving consumer. Suppose also that the firm explicitly prevents the posted price from changing too often, perhaps in an attempt to avoid personalized pricing or suspicion of it. This should deliver outcomes similar to sticky targeted pricing, and thus is a form of personalized pricing rather than merely dynamic pricing.[5]

Would this sticky form of personalized pricing effectively avoid detection? Consider how consumers might try to verify personalized pricing. A consumer offered a high price might check whether an acquaintance is offered the same price. They would. Any two

[4] Appendix Section A details the stages of interaction between a consumer and website and the types of information at the website's disposal to personalize prices.

[5] Note that inputting many variables into a machine learning model may be less effective at raising profits compared to a model which formally incorporates dynamics.

consumers checking the price at the same moment would observe the same price because the firm has privately committed to maintaining the new price for some interval.[6] It would be difficult, if not impossible, for consumers to distinguish whether price changes arise from personalized pricing or traditional dynamic pricing, the latter of which is common and tolerated by consumers.[7] The same reasoning implies that researchers and regulators searching for personalized pricing would fail to detect sticky personalized pricing.

This paper examines sticky targeted pricing. First, it presents a dynamic pricing model, which characterizes optimal price(s) to offer an arriving consumer under the constraint that price remains locked for some interval following a change. The model shows that firms face a tradeoff between extracting surplus from the arriving consumer and profiting from later arrivals who must be charged the same price.

The model is then applied to four contexts: one empirical context (Netflix) and three theoretical distributions of consumer valuations (uniform, normal, and exponential). For the empirical context, the distribution across consumers of individual-level demand functions follows Shiller (2020), which estimates individual-level expected demand based on web-browsing habits. Profits and price paths are then simulated for each of the four considered distributions of valuations, under various assumed consumer arrival rates.

Similar patterns appear for all considered distributions of valuations. Counterfactual simulations show that optimized sticky targeted pricing meaningfully raises profits for products of low and medium popularity. Intuitively, the percentage increase in profits is largest for unpopular products, the long tail of products. However, I find that the absolute change in profits is largest for medium popularity products, for the following reason. Initially, as popularity rises, the profit gains from applying the strategy to a larger customer base outweighs the reduction in per-consumer profit gains. For very popular products, however, the firm forgoes targeted pricing and instead uses uniform pricing; profit gains from raising price to a high-value arriving consumer are offset by reduced profits from setting the same

[6]The price may change just before an acquaintance checks price. But then the first consumer would likely recheck price, this time observing the same price as their acquaintance, thus inferring that the previous price change was due to dynamic pricing.

[7]Traditional dynamic pricing arises, for example, from responses to changes in market conditions, such as aggregate demand shocks, competitor actions, and changes in inventory or costs, as well as from intertemporal (second-degree) price discrimination.

high price to many subsequent arrivals who, in expectation, have lower willingness to pay. Similarly, I find that a sensible outward measure of the extent of price discrimination—the range of prices over a reasonable length interval—is largest for medium-popularity items. These findings imply that the proposed strategy is not confined to unpopular items.

Finally, I find that the the popularity level where this strategy is most profitable and where price discrimination is most intense rises with the precision of estimated willingness to pay. We should thus expect it to be applied to a wider range of products as firms gain access to more detailed customer data.

Increasing use of inconspicuous but sophisticated pricing methods has substantive implications beyond the price discrimination literature. For example, overlooking its use yields biased estimates of consumer demand (D'Haultfœuille et al., 2019) and misleading inflation measurements (Chevalier and Kashyap, 2019). Moreover, increasingly intense price discrimination has meaningful implications for the effects of competition on profits (Thisse and Vives, 1988). However, because its use is not readily apparent, it may be used widely without academics and regulators being aware. These issues become more problematic as use of sophisticated pricing intensifies. Thus, it may be necessary to incorporate sticky personalized pricing in economic models, even when the research question is not directly related.

The remainder of the paper is organized as follows. Section 2 discusses how effectively sticky personalized pricing conceals use of price discrimination. Section 3 introduces a model of sticky targeted pricing, and Section 4 simulates counterfactual outcomes under this pricing strategy. A brief conclusion follows.

2 Background

For this strategy to be useful, it must successfully disguise personalized pricing from consumers, regulators, researchers, and competing firms.[8] Searching for straightforward (non-sticky) personalized pricing is simple; one can examine whether two individuals are offered

[8]Appendix Section A details how online prices can be personalized in practice.

different prices for the same product at the same point in time. However, it is more challenging to verify use of sticky targeted pricing.

Unless they leave long time lags—longer than the price-commitment period—between checking prices for different spoofed consumers, regulators, researchers, and competitors would infer that the same price is offered to nearly all consumers. However, researchers have typically checked prices offered to different spoofed consumers in rapid succession to distinguish personalized pricing from traditional dynamic pricing. Of the studies searching for personalized pricing online (Cavallo, 2017; Hannak et al., 2014; Hupperich et al., 2018; Iordanou et al., 2017; Mikians et al., 2012), none explicitly stated that they incorporated long lags between arrivals of different spoofed consumers. Thus, sticky personalized pricing would effectively avoid detection, at least from methods that have previously been used to search for personalized pricing.

If one instead searches for sticky personalized pricing by comparing prices offered to various consumers at different points in time, it is no longer sufficient to show that the prices differed. One must distinguish whether those price differences are attributed to sticky targeted pricing, or the host of other factors that are known to cause prices to change over time. One might try relating price differences to consumer traits that are perceived to be useful for personalized pricing (e.g., income).[9] However, even if one can show that prices tend to be higher when price insensitive consumers are arriving, one must distinguish whether these differences reflect predictable fluctuations in demand or are truly personalized and thus depend on the realized identity of the arriving customer(s).[10] For example, restaurants' early bird specials set lower prices for afternoon patrons in attempts to target price sensitive seniors, but the discount is time-based and hence would apply even in instances when many price insensitive arrive by chance. Concluding that early bird specials are a form of personalized pricing is therefore inaccurate. Note that this distinction not merely semantic. Early bird specials do not elicit the same sort of consumer resentment concerns

[9]This test requires observing quoted prices rather than transaction prices, as average completed purchase prices will be higher for groups with more high value consumers who remain willing to buy when (time-varying) prices are high.

[10]Furthermore, one must also show that price differences are not merely attributed to differences in the costs of servicing difference consumers (e.g., as for insurance) in order to classify it as price discrimination.

that personalized prices do.

Even if overcoming these concerns, testing for personalized pricing by linking offered prices to consumer characteristics suffices only if one can identify and has access to the variables the firm is using to personalize prices. Often, the variables most useful for personalizing prices are not immediately apparent. For example, Shiller (2020) found that income and other demographics revealed relatively little about a consumer's valuation for Netflix's products. Instead, it was use of websites that deliver products by mail (e.g., Amazon) that most strongly indicated high valuations for Netflix's products. Furthermore, if firms intend to evade detection—which their use of sticky targeting pricing would imply—they may exclude obvious and widely available variables from their pricing algorithm, and avoid targeting prices when a consumer arrives by a direct url (as they mostly likely would in a survey) rather than by searching on the platform.

It thus appears that intentionally making personalized pricing sticky would at least substantially complicate others' efforts to verify its use and arguably avoid being classified as personalized pricing under definitions proposed by regulators. If simultaneously effective at extracting surplus, sticky targeted pricing may be an enticing strategy for online retailers.

3 A Model of Sticky Targeted Pricing

This section characterizes the profit-maximizing sticky targeted pricing strategy. The model reflects the interaction between the firm and an arriving consumer, which evolves in several stages. First, the arriving consumer's type is revealed to the firm. Then the firm chooses the posted price price, if able to. If the time elapsed since the last price change is less than the length s of the price-commitment period, the posted price remains the same. If, instead, sufficient time has elapsed since the last price change, then the firm may change the posted price. The arriving consumer then observes the posted price, and makes a myopic decision of whether and how much to purchase.

The firm chooses price to maximize discounted expected profits. It faces a decision over the price to offer when two conditions are met: (i) the firm is able to change price (the price-commitment period has elapsed), and (ii) a new consumer is arriving. The firm's discounted profits are:

$$V(P, \psi, t) = \max_{P'} \begin{cases} W^{P'=P}(P, \psi, t) & \text{if } P' = P \\ W^{P' \neq P}(P', \psi, t) & \text{if } P' \neq P \end{cases}, \quad (1)$$

where P, ψ, and t are state variables denoting the price last offered by the firm, the consumer's type, and time, respectively. P' is the firm's decision variable: the (new) posted price to offer to the arriving consumer of then-known type ψ. $W^{P'=P}(P, \psi, t)$ denotes the expected discounted profits when the newly posted price is the same as the last, and $W^{P' \neq P}(P', \psi, t)$ denotes discounted profits when price changes.

If the posted price remains the same, then the firm's expected discounted profits are:

$$W^{P'=P}(P, \psi, t) = \pi(P, \psi, t) + \int_{\tau=0}^{\infty} \int_{\psi'} \exp(-r\tau) V(P, \psi', t+\tau) g(\psi'; t+\tau) f(\tau; \lambda, t) d\psi' d\tau, \quad (2)$$

where $\pi(P, \psi, t)$ represents expected static profits from the arriving consumer and the integral denotes the expected discounted future profits, given that two relevant factors are random: the time τ between consumer arrivals and the next consumer's type ψ'. Within the integrals, $\exp(-r\tau)$ is the continuous-time analogue to the discount factor at the next consumer arrival epoch, and $V(P, \psi', t + \tau)$ is the value function. The density functions for the arriving consumer's type and interarrival time are denoted by $g(\psi'; t)$ and $f(\tau; \lambda, t)$, respectively, where the latter depends on an arrival rate parameter: λ.

If the offer price changes, then expected discounted profits instead equal:

$$
\begin{aligned}
W^{P' \neq P}(P', \psi, t) = &\overbrace{\pi(P', \psi, t)}^{\text{A}} \\
&+ \overbrace{\left(\int_{\tau=0}^{s} h(\lambda, t+\tau) \exp(-r\tau) \int_{\psi'} \pi(P', \psi', t+\tau) g(\psi'; t+\tau) d\psi' d\tau \right)}^{\text{B}} \\
&+ \underbrace{\int_{\tau=s}^{\infty} \int_{\psi''} \exp\left(-r\tau\right) V(P', \psi'', t+\tau) g(\psi''; t+\tau) f(\tau; \lambda, t+s) d\psi'' d\tau}_{C} .
\end{aligned} \tag{3}
$$

Component A represents expected static profits from the arriving consumer at price P'. Component B represents expected discounted profits at that same price during the price-commitment period. Within Component B, $h(\lambda, t+\tau)$ denotes the expected density of arrivals at time $t+\tau$, $\exp(-r\tau)$ corresponds to the discount factor, and $\int_{\psi'} \pi(P', \psi', t+\tau) g(\psi'; t+\tau) d\psi'$ is the expected profit at price P' from a randomly drawn consumer. Finally, Component C equals expected discounted profits earned after the price-commitment period ends; it is analogous to the integral in Equation 2.

Finally, the firm's policy function equals:

$$
P'(P, \psi, t) = \arg\max_{P'} \begin{cases} W^{P'=P}(P, \psi, t) & \text{if } P' = P \\ W^{P' \neq P}(P', \psi, t) & \text{if } P' \neq P \end{cases} . \tag{4}
$$

Note this model can be applied, using ex-ante estimates of individual-level profit functions ($\pi(P', \gamma, t)$) and expected arrival rates. Specifically, one can approximate the function in Equation 1 via value function iteration, then calculate the policy function in Equation 4.

3.1 Observations

Equation 3 shows the firm faces a trade-off when changing price. By tailoring the "posted price" to the arriving consumer, the firm can usually raise expected profits earned from that consumer. However, there is an implied cost. The firm must offer the same price to consumers who arrive shortly thereafter (if it intends to use sticky targeted pricing), which may lower expected profits from these later arrivals.

For example, suppose the firm raises the posted price to extract surplus from a high-value arriving consumer. Then it must offer the same high price—higher than the optimal uniform price—to subsequent consumers arriving soon thereafter. However, the price that maximizes expected profits from later arrivals—whose types are not yet known—is the optimal uniform price. Any higher (or lower) price reduces expected profits from these later arrivals.

The importance of this observation depends on the rate of customer arrivals. If arrivals are infrequent, then the gains from exploiting a high-value consumer likely outweigh forgone profits from later arrivals, who are expected to be few in number. However, for popular products—with many expected customer arrivals during the fixed-price period—forgone profits from later arrivals are large.[11] For extremely popular products, firms forgo sticky targeted pricing altogether, even if the static gains from tailoring the posted price to the arriving consumer are large.

The intertemporal trade-off also depends on the precision of the firm's estimate of the arriving consumers' willingness to pay (WTP). Ceteris paribus, the static profit gains from tailoring price to the arriving consumer are higher when the consumer's WTP is estimated with greater precision. With perfect precision, the firm can raise the price to the WTP of the high-value arriving consumer without risking losing the sale. Similarly, lowering the price to a low-value consumer's WTP increases the probability of a completed transaction by 1. By contrast, with imprecise estimates of consumers' WTP, lowering the price to a prospective low-value consumer may be unnecessary—the transaction might still occur at a higher price if their WTP is higher than expected—or may fail to result in a completed transaction if the consumer's valuation is lower than expected.

These points can be illustrated with simple examples. Suppose consumers' valuations are drawn from the standard uniform distribution and firm costs are zero. Then, the profit-maximizing uniform price is 0.5. Now suppose a consumer arrives, and the firm knows the

[11] The expected number of arrivals logically depends on the length of the fixed-price period. This paper assumes that the length of the fixed-price period is chosen to successfully obfuscate targeted pricing, and abstracts from the question of how long it should last, which would require data not typically available to researchers.

consumer's willingness to pay is exactly 0.75. Then changing the price to 0.75 increases profits from this consumer by 0.25, relative to the profits from this consumer under uniform pricing (0.5). Now, suppose a similar consumer arrives, who the firm estimates has the same willingness to pay, 0.75, but the firm's estimate is imprecise. Specifically, suppose the firm knows only that the arriving consumer's value is drawn from a three point distribution (0.65, 0.75, 0.85) with each point occurring with equal probability. Then, the firm maximizes static profits with a price of 0.65, earning static profits only 0.15 higher than under uniform pricing.[12] Similarly constructed examples can show that lowering price to a prospective low-value consumer is less profitable when consumers' valuations are imprecisely estimated.

Hence, increasing the precision of WTP estimates expands the set of scenarios where the static gains from personalizing the posted price exceeds the opportunity cost of setting a suboptimal price to later arrivals during the fixed-price period. One should thus expect sticky targeted pricing to be applied more aggressively and to a wider range of products when estimates of individuals' willingness to pay are more precise.

The above trade-offs imply that it is not obvious whether sticky targeted pricing is a promising general strategy, or whether it is profitable only in limited contexts. It is therefore an empirical question, one which is investigated next.

4 Applications

This section simulates outcomes under various assumed distributions of consumer heterogeneity—one empirical and three theoretical—to explore the impacts of sticky targeted pricing. The first assumed distribution of consumer types ($g(\psi)$) and individual-specific profit functions ($\pi(P, \psi)$) follows Shiller (2020). The paper uses the 2006 ComScore panel of approximately 60,000 internet users to relate a consumer's decision of whether to subscribe to Netflix to the frequency of their visits to nearly 5,000 other websites. Combining predicted individual-level subscription probabilities with an optimal pricing condition allows estimation of individual-

[12] Setting price to 0.75 or 0.85 yields lower profit. For example, a price of 0.75 leads to a sale 2/3 of the time, i.e., when the consumer's realized valuation is not 0.65, implying $2/3 \times 0.75 = 0.50$ in expected profits, the same as under uniform pricing.

level demand functions, which individual-level profit functions follow from. Note that while observed demand is discrete (1[buy]), from the firm's perspective ex-ante individual-level demand is continuous: demand equals the predicted probability of purchase at the specified price. An overview of the methods used in Shiller (2020) to generate the individual-level profit functions, utilized in this paper, are provided in Appendix Section B.[13]

Three theoretical distributions of consumers' valuations are also considered: uniform, normal, and exponential. Each is normalized so they all have the same mean and standard deviation, both equal to one. For these theoretical distributions, it is initially assumed that the firm observes each arriving consumer's valuation perfectly: There is no error in the firm's estimate of the consumer's willingness to pay (WTP). Hence, assuming costs are zero, the profit earned from an individual at price P' equals: $\pi(P', \psi) = P' \times 1[P' \leq \text{WTP}]$.

The firm's pricing policy function also depends on the rate of consumer arrivals. To isolate the impact of product popularity on market outcomes, an array of different consumer arrival rates (λ) are considered [0.01; 0.05; 0.25; 1.25; 5; 25; 125; 625; 3,125; 15,625; 78,125; 390,625; 1,953,125]. Each denotes a different expected count of consumer arrivals during the interval while price is fixed.

Interarrival times are assumed to be independent and identically distributed. This implies that that consumer interarrival times follow the exponential distribution and the count of consumer arrivals during a specified interval follows the Poisson distribution.

For each assumed distribution of individual-level demand functions and each assumed consumer arrival rate, pricing simulations proceed in several steps. First, the firm's value function is approximated using value function iteration, by iterating on the Bellman equation in Equation 1. Then, the value function is used to determine the firm's policy function. Some outcomes directly follow. For example, expected discounted variable profits equal the value function: $V(P, \psi)$. Other outcomes are simulated from the policy function and a long randomly drawn path of consumer arrivals.[14]

[13]The requisite data and information used to generate the distribution of individual-level demand functions are publicly available at: `https://benjaminshiller.com/Research`.

[14]Firm price choices and consumer purchase decisions are simulated for a length of time lasting until the firm has had a chance to change the posted price 10^8 times; i.e., 10^8 consumers have arrived at times when the firm has not privately committed to keeping price the same. The actual number of consumers arriving

In these simulations, the per-period (length s) interest rate is assumed to be 0.1/365. If the price-commitment period s lasts one day, this corresponds to a 10% yearly interest rate.

The next subsection utilizes these simulations to examine outcomes under sticky targeted pricing. Specifically, it explores how profitability and pricing patterns relate to the consumer arrival rate (popularity), and examines whether patterns are generalizable across the various assumed distributions.

4.1 Comparative Statics

4.1.1 Profits

Figure 1 compares discounted profits from sticky targeted pricing and from status quo uniform pricing, for an array of different consumer arrival rates. Note that some similar patterns emerge across the different distributions of consumer valuations. Profit gains expressed in percentage terms, shown in the top panel, are largest when the consumer arrival rate is low (i.e., for relatively unpopular products). This finding is intuitive. However, the relationship between the absolute change in profits and the arrival rate, shown in the bottom panel, is more nuanced. Initially, as the consumer arrival rate increases, the gains from applying this strategy to more consumers outweighs the lower per consumer gain: Profits from sticky targeted pricing initially rise with the arrival rate. However, eventually the latter effect dominates, reversing this pattern.[15] In the extreme, for very high consumer arrival rates, the firm forgoes the opportunity to personalize prices altogether: Personalizing prices for one consumer locks in a suboptimal price for the many consumers arriving shortly thereafter.

In contrast to basic intuition, this pattern for absolute profits implies that this strategy is not merely applicable to the least popular products. In fact, with a per product fixed cost, one might expect a firm to determine that benefits exceed the cost of implementation only for medium popularity products; not for less popular products. Thus, viewing this strategy

over the interval is larger, as it includes additional consumers arriving while price is fixed.

[15]The local maximum for the exponential distribution in Figure 1d remains after increasing the tolerance in value functions computations and after doubling the count of simulated consumers. However, with computational constraints of available hardware—Nvidia Tesla K40c GPU—I cannot exclude the possibility that the local maximum would disappear with more simulated consumers.

as applicable only for products of little relevance is irrational.

Note in Figure fig.Profits that the profit gains from sticky targeted pricing are both larger and extend to higher popularities for the theoretical distributions of valuations, relative to the empirical distribution. One prominent reason for this finding is that predicted valuations are imprecise for the empirical distribution, but assumed to be exact for the theoretical distribution. As explained in Section 3.1, the gains from personalization are larger when estimates of consumer valuations are more precise.

To examine the magnitude of the impact of precision, analogous simulations are conducted with uncertain valuations. Specifically, firm estimates of consumer valuations are assumed to follow the same normalized uniform distribution as before, but in this exercise a consumer's true valuation equals the firm's estimate plus a private draw. The private draw is assumed to follow a mean-zero uniform distribution; the range is varied to explore its impact. The firm then chooses the posted price to maximize expected discounted profits given this uncertainty.

Figure 2 shows the absolute change in profits against the consumer arrival rate, separately for various extents of uncertainty in the firm's estimate of a consumer's valuation. Note the theorized pattern is observed in this targeted exercise: As uncertainty decreases, this pricing strategy becomes more profitable generally, and we see a rise in the popularity-level where profit gains peak. Moreover, observed differences are dramatic. Consider the case where the uncertainty range equals 1, implying a relatively large private random draw; with range equaling the standard deviation in the observable component of WTP across consumers. In this case, sticky targeted pricing yields the largest absolute profit gains when applied to products with relatively low arrival rates; profit gains peak when roughly five consumers are expected to arrive during the fixed-price period. When the range of private random draws is a tenth as large, sticky targeted pricing is most profitable for products popular enough to expect about 25 consumers to arrive while price remains fixed. For the case with no uncertainty—consumers' exact valuations are observed by the firm—this pricing strategy is most profitable for products viewed and considered by 3,125 consumers during the period while price is fixed.

These findings imply an advantage for large multi-product online retailers. They have larger consumer tracking datasets, allowing more precise predictions of individuals' WTPs. Large online platforms thus benefit more from sticky targeted pricing, relative smaller rivals and brick-and-mortar competitors who lack equivalent information.

To illustrate the feasibility of this pricing strategy in practice, consider book sales. According to NPD's Bookscan dataset, in the week ending February 19^{th}, 2022, the top-selling book sold 118,800 copies ($118,800/7 \approx 17,000$ per day) in the U.S., the second highest-selling book sold 34,897 copies ($\approx 5,000$ per day), and the 10th-ranked book sold 16,266 copies ($\approx 2,300$ per day).[16] If using results from the empirical context, Netflix in 2006, and conservatively assuming all sales flow though a single retailer, then sticky targeted pricing is profitable for best-selling books only if a relatively short fixed-price period is sufficient to obfuscate targeted pricing. If the fixed-price interval lasts longer, say a full day, then even the 10th-ranked book's daily sales (about 2,300) exceed the count of consumer arrivals in which sticky targeted pricing can be profitably applied. However, if the requisite fixed-price period lasts only an hour, and a retailer only accounts for, say, a fifth of sales, then the 10th-ranked book sells only about 20 ($2,300/(24 \times 5)$) copies during the fixed-price period. If sales are somewhat close to the count of arriving consumers, then sticky targeted pricing would yield profit gains even among best-selling books. Note in Figure 1c that while profit gains peak when five consumers are expected to arrive during the fixed-price period, profit gains when 25 consumers are expected to arrive are still quite large, approximately 90% of peak.

Furthermore, note that the empirical application uses data from a relatively early in the Internet era, 2006, and uses a relatively small sample of consumers, about 60,000. Consumers now spend much more time shopping online—generating more useful data to segment consumers—and large online platforms like Amazon have hundreds of millions of active consumers to train their models on.[17] Presumably, individual-level demand estimates are much more precise now than the somewhat dated empirical application suggests. Figure

[16]See `https://www.publishersweekly.com/pw/by-topic/industry-news/bookselling/article/88638-this-week-s-bestsellers-february-25-2022.html`

[17]Reisinger (2019) notes that Amazon surpassed 100 million "Prime" subscribers in early 2019.

1d shows that when consumer valuations are sufficiently precise, profit gains peak when about 3,125 consumers are expected to arrive during the fixed-price period. Note that this popularity level exceeds the sales of the 10th-ranked book even under conservative assumptions that the fixed-price period is relatively long, a day, and all sales flow through a single retailer. Thus, it seems that firms with better consumer-level data can apply this strategy to nearly all products, including some best sellers; sticky targeted pricing is not merely a strategy confined to fringe products, but rather may be profitably applied widely.

4.1.2 Pricing Patterns

This subsection investigates the impact of sticky targeted pricing, and the role of popularity, on the extent of price discrimination. The impact across consumers in the empirical context is illustrated in the left panel of Figure 3. The graph shows the consumer's type on the horizontal axis against the markup they would be offered on the vertical axis.[18] Each line shows the relationship between these two variables for a separate assumed rate of consumers arrivals (popularity). First, note that price discrimination has large impacts on consumers at the extremes. Second, note that the range across consumers diminishes with product popularity: For higher arrival rates, the line becomes flatter, implying the firm reduces the extent of price targeting to avoid locking in an extreme price for the many customers expected to arrive shortly thereafter during the fixed-price period. When popularity rises to the point where 625 consumer arrivals are expected during the fixed-price period, the line becomes perfectly flat, implying that the firm forgoes the opportunity to change the posted price to extract profit from the arriving consumer, even for extreme customers.

Figure 4 shows similar patterns for the three theoretical distributions of consumer valuation, in a succinct manner. The figure plots the interdecile price range across different consumer types against the consumer arrival rate. Note that the pattern is similar to the empirical case: As product popularity rises, there is a decline in the extent of price discrimination, at least when measured as the range of prices offered across different consumer types. However, the firm continues to tailor the price somewhat even for popular products:

[18] The markup to offer a given consumer depends on the state variable: the last offered price. In the graph, it was assumed that the last price equals the price the firm would choose if price could never change.

The interdecile price range is still above zero for popular products expecting at least 1,000 consumer arrivals during the fixed-price period.

The preceding analysis suggests that the extent of price discrimination declines in product popularity, a finding which matches basic intuition. However, this conclusion relies on a particular measure of price discrimination, one infrequently used in empirical applications because one typically cannot observe the price each consumer is offered.[19] Instead, price discrimination intensity may be measured by the range of observed posted prices.

Figure 5 shows the expected range of prices offered over a medium-length period against product popularity. Note the inverted U-shaped pattern, which implies the largest price variation for medium-popularity products. This pattern arises due to two competing forces. First, as the arrival rate increases, the firm chooses a smaller range of prices across different types of arriving consumers. However, frequent arrivals imply more opportunities to personalize prices during a specified length of time: more type/price draws. Initially, the latter effect dominates, and the range of prices offered over an interval increases with the rate of consumer arrivals. However, the impact of the latter effect—additional price draws—is diminishing. More existing price draws implies a larger initial price range (in expectation), both (1) reducing the chance that the next offered price falls outside this range, and (2) implying that a particular price draw outside the range expands the range by less.[20] Eventually the former effect—smaller price ranges across different consumer types—dominates, giving rise to an inverted U-shaped relationship between the price range over a medium-length interval and the consumer arrival rate.[21].

Overall, this latter measure of price discrimination intensity, which is analogous to how price discrimination would presumably be measured using widely available datasets, implies that the extent of price discrimination is highest for medium popularity items.[22] This

[19]Relying on purchase prices, rather than posted pricing, yields misleading findings. If prices merely vary randomly and not due to price discrimination, one would still observe some groups paying more than others. The average completed purchase prices will be higher for groups with high value consumers who remain willing to buy even when (time-varying) prices are high.

[20]The probability the next draw increases the range equals $2/n$, where n is the number of prior draws.

[21]Appendix Figure C1 shows that this pattern is robust to the length of period used to calculate the price range.

[22]The frequency of price changes over time is also largest for medium-popularity products. See Appendix Figure C2b

implies that sticky targeted pricing yields the most intense price discrimination for products of medium popularity, rather than for fringe products with limited appeal.

5 Conclusion

Recent anecdotes suggest a trend towards more sophisticated pricing strategies that simultaneously allow finer targeting while assuaging consumer and regulatory concerns. This paper presents a novel pricing strategy to discreetly implement personalized pricing so as to avoid consumer resentment and regulatory scrutiny. I find that the pricing strategy substantially raises profits over non-personalized pricing. Perhaps surprisingly, the overall profits gained from this strategy and measures of discrimination intensity are largest for products with medium popularity, rather than for products with the lowest popularity. If firms' estimates of individual-consumers' willingness to pay are rather precise, then this strategy is most profitable for products viewed rather frequently, on the order or 3,000 consumer arrivals during the interval the firm privately commits to keeping the posted price fixed (e.g., a day or an hour). This pricing strategy should therefore become more enticing and be applied to a wider range of products as increasing amounts of consumer data become available. It also adds to the list of market features that favor large online platforms.

Still, common wisdom in some academic circles suggests that firms remain reluctant to implement any form of personalized pricing due to concerns of reprisals from consumers, policy-makers, and regulators. However, if firms can raise profits via more sophisticated targeted pricing methods, without their strategies being realized by those opposed to them, then why would firms not do so? This reasoning further suggests that academics' tendency to overlook use of sophisticated pricing techniques may soon become (or may already be) untenable.

References

Allender, W. J., J. Liaukonyte, S. Nasser, and T. J. Richards (2021). Price fairness and strategic obfuscation. *Marketing Science 40*(1), 122–146.

Anderson, E. T. and D. I. Simester (2008). Research note—does demand fall when customers perceive that prices are unfair? the case of premium pricing for large sizes. *Marketing Science 27*(3), 492–500.

Aparicio, D., Z. Metzman, and R. Rigobon (2021). The pricing strategies of online grocery retailers. *Working Paper*.

Baird, N. and B. Kilcourse (2017). Retail pricing 2017: The dawn of personalized prices. Technical report, Retail Systems Research.

Campbell, M. C. (1999). Perceptions of price unfairness: antecedents and consequences. *Journal of Marketing Research 36*(2), 187–199.

Cavallo, A. (2017). Are online and offline prices similar? Evidence from large multi-channel retailers. *American Economic Review 107*(1), 283–303.

Chevalier, J. A. and A. K. Kashyap (2019). Best prices: Price discrimination and consumer substitution. *American Economic Journal: Economic Policy 11*(1), 126–59.

Conrad, J. and W. Knight (2022, February 22). China is about to regulate AI—and the world is watching. *Wired*.

Darke, P. R. and D. W. Dahl (2003). Fairness and discounts: The subjective value of a bargain. *Journal of Consumer psychology 13*(3), 328–338.

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in US retail chains. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Dong, X., P. Manchanda, and P. K. Chintagunta (2009). Quantifying the benefits of individual-level targeting in the presence of firm strategic behavior. *Journal of Marketing Research 46*(2), 207–221.

Dube, J.-P. and S. Misra (2017). Scalable price targeting. *Working Paper*.

D'Haultfœuille, X., I. Durrmeyer, and P. Février (2019). Automobile prices in market equilibrium with unobserved price discrimination. *The Review of Economic Studies 86*(5), 1973–1998.

Executive Office of the President (2015). Big data and differential pricing. *Washington DC*.

Goli, A. (2020). *Personalized Versioning: Product Strategies Constructed from Experiments on Pandora*. Ph. D. thesis, The University of Chicago.

Hannak, A., G. Soeller, D. Lazer, A. Mislove, and C. Wilson (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 305–318. ACM.

Ho, T.-H. and X. Su (2009). Peer-induced fairness in games. *American Economic Review 99*(5), 2022–49.

Hupperich, T., D. Tatang, N. Wilkop, and T. Holz (2018). An empirical study on online price differentiation. In *Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy*, pp. 76–83.

Iordanou, C., C. Soriente, M. Sirivianos, and N. Laoutaris (2017). Who is fiddling with prices? Building and deploying a watchdog service for e-commerce. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pp. 376–389.

Kahneman, D., J. L. Knetsch, and R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review 76*(4), 728–741.

Kehoe, P. J., B. J. Larsen, and E. Pastorino (2018). Dynamic competition in the era of big data. *Working paper, Stanford University and Federal Reserve Bank of Minneapolis*.

Khan, L. M. (2016). Amazon's antitrust paradox. *Yale Law Journal 126*, 710.

Leibbrandt, A. (2020). Behavioral constraints on price discrimination: Experimental evidence on pricing and customer antagonism. *European Economic Review 121*, 103303.

Li, K. J. and S. Jain (2016). Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Science 62*(9), 2705–2721.

Mikians, J., L. Gyarmati, V. Erramilli, and N. Laoutaris (2012). Detecting price and search discrimination on the internet. *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, 79–84.

Reimers, I. and B. R. Shiller (2019). The impacts of telematics on competition and consumer behavior in insurance. *The Journal of Law and Economics 62*(4), 613–632.

Reisinger, D. (2019, Jan). Amazon prime has more than 100 million u.s. subscribers. *Forbes*.

Rossi, P. E., R. E. McCulloch, and G. M. Allenby (1996). The value of purchase history data in target marketing. *Marketing Science 15*(4), 321–340.

Shiller, B. and J. Waldfogel (2011). Music for a song: An empirical look at uniform pricing and its alternatives. *The Journal of Industrial Economics 59*(4), 630–660.

Shiller, B. R. (2020). Approximating purchase propensities and reservation prices from broad consumer tracking. *International Economic Review 61*(2), 847–870.

Smith, A., I. Aggarwal, and S. Seiler (2021). Optimal price targeting.

Thisse, J.-F. and X. Vives (1988). On the strategic choice of spatial price policy. *The American Economic Review 78*(1), 122–137.

Wong, B. (2021). Online personalised pricing as prohibited automated decision-making under article 22 gdpr: a sceptical view. *Information & Communications Technology Law 30*(2), 193–207.

Xia, L. and K. B. Monroe (2010). Is a good deal always fair? examining the concepts of transaction value and price fairness. *Journal of Economic Psychology 31*(6), 884–894.

Zhang, J. Z., O. Netzer, and A. Ansari (2014). Dynamic targeted pricing in B2B relationships. *Marketing Science 33*(3), 317–337.

Figure 1: Counterfactual Profit Gain v. Consumer Arrival Rate

Percentage

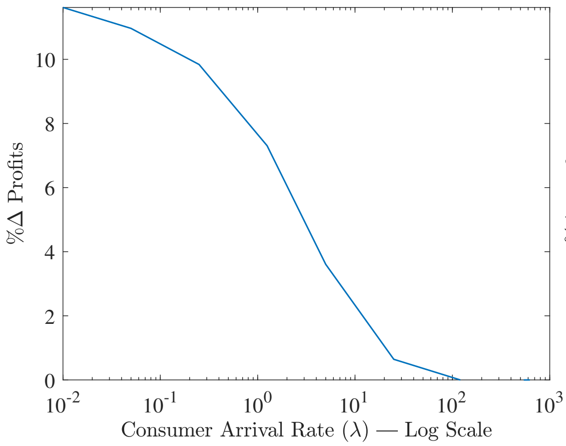

(a) Empirical: Netflix

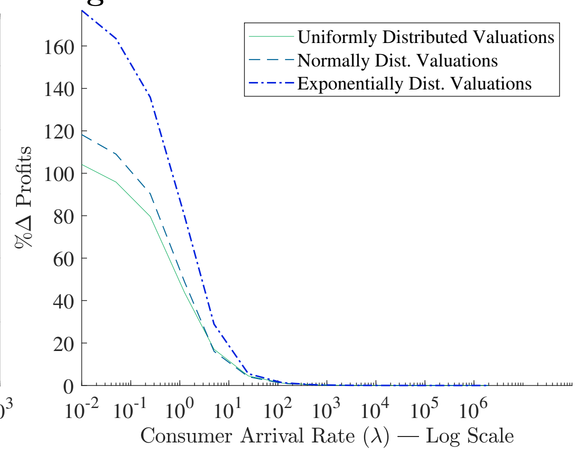

(b) Theoretical Distributions

Absolute (Normalized)

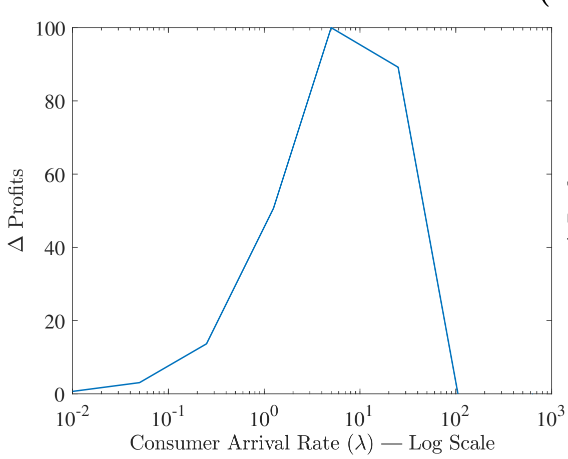

(c) Empirical: Netflix

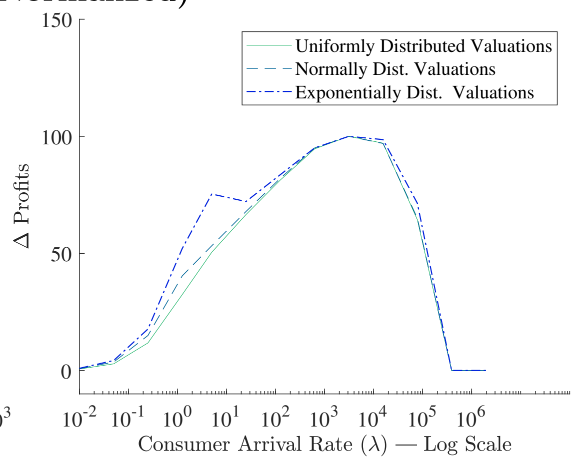

(d) Theoretical Distributions

Notes: The top panel shows the percentage change in profits from implementing sticky targeted pricing relative to uniform pricing for the empirical application (on the left) and theoretical distributions (on the right). The bottom panel shows similar graphs for the absolute increase in profits, normalized such that the highest value across the various arrival rates equals 100. Note that expected discounted profits under status quo uniform pricing equal: $\max_{\hat{P}} \frac{\lambda}{r} \int_{\psi} \pi(\hat{P}, \psi) g(\psi) d\psi$, and under sticky targeted pricing equal $\int_{\psi} V(\hat{P}, \psi) d\psi$, where $\hat{P}$ denotes the optimal uniform price. The consumer arrival rate (λ) denotes the expected number of consumer arrivals during the fixed-price interval.

Figure 2: Counterfactual Profits and the Impact of Uncertainty

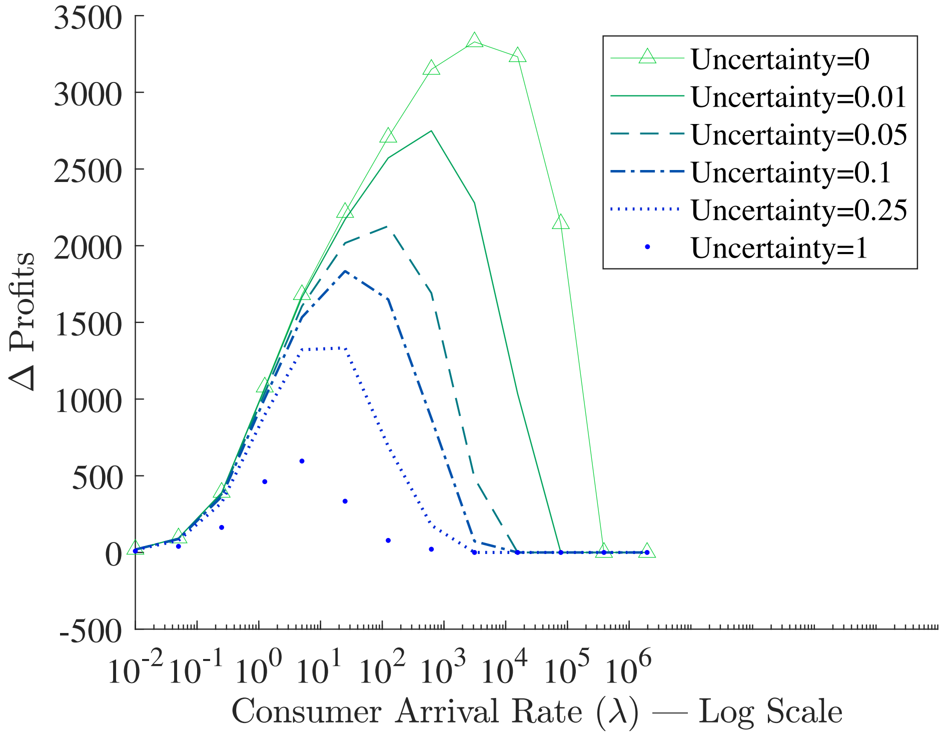

This figure depicts the absolute profit gained from sticky targeted pricing as a function of the consumer arrival rate, separately for select levels of uncertainty. A consumer's true valuation is assumed to equal the sum of the firm's estimate of willingness to pay—following the same uniform distribution as before—and random draw that is private information of the consumer. The latter (private) random draw is assumed to follow the mean zero uniform distribution, with range equal to listed uncertainty.

Figure 3: Simulated Price Range: Across Consumers

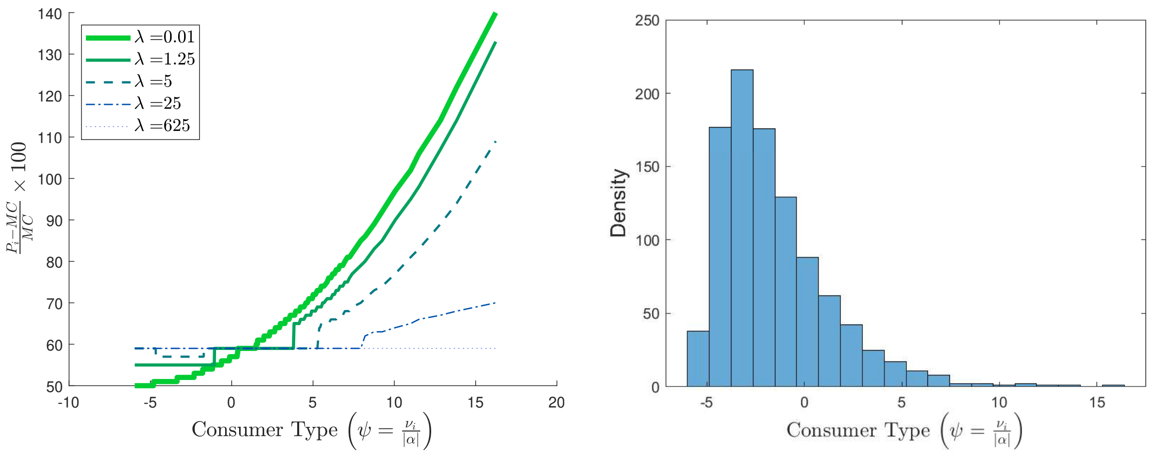

(a) Empirical: Netflix

(b) Empirical Type Density

Notes: The left panel shows the relationship between the percent markup offered and the arriving consumer's type, assuming the previous markup was the markup the firm would choose if price could never change. Each line on the graph shows the range of markups across consumers for a specific arrival rate (λ). The right panel shows the density of consumer types.

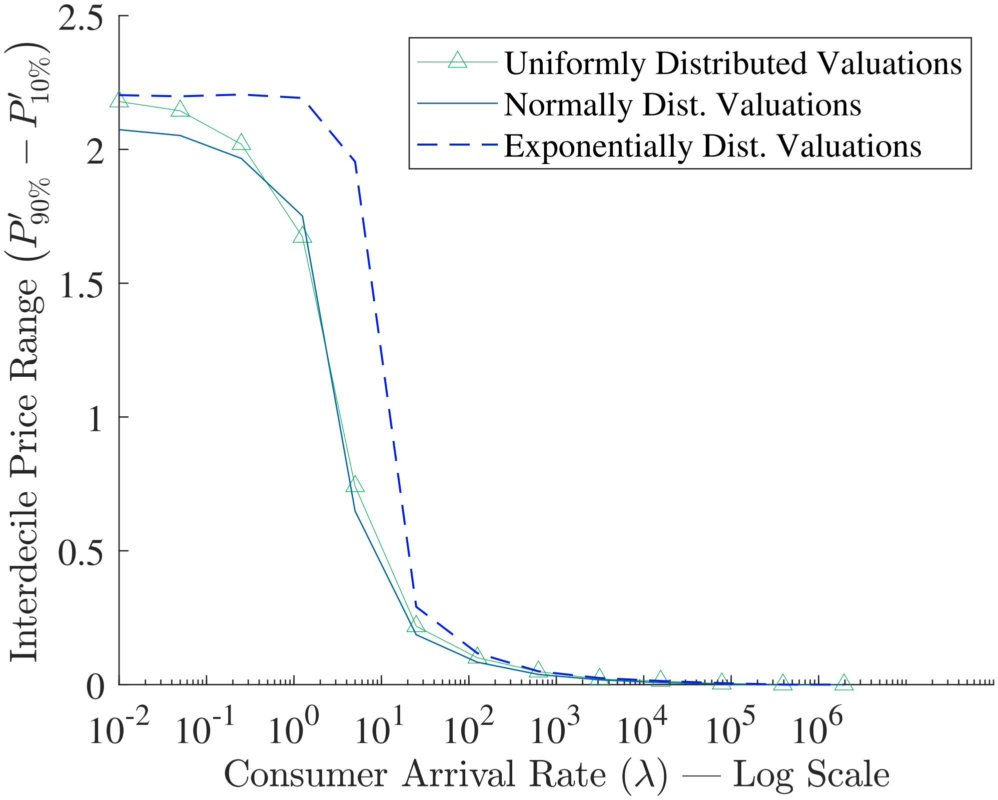

Figure 4: Price Range Across Consumers: Theoretical Distributions

Notes: This figure shows the interdecile range of simulated prices offered across different consumer types—when the firm can freely change price—against the consumer arrival rate, for the three theoretical distributions.

Figure 5: Simulated Price Range: Time-Normalized

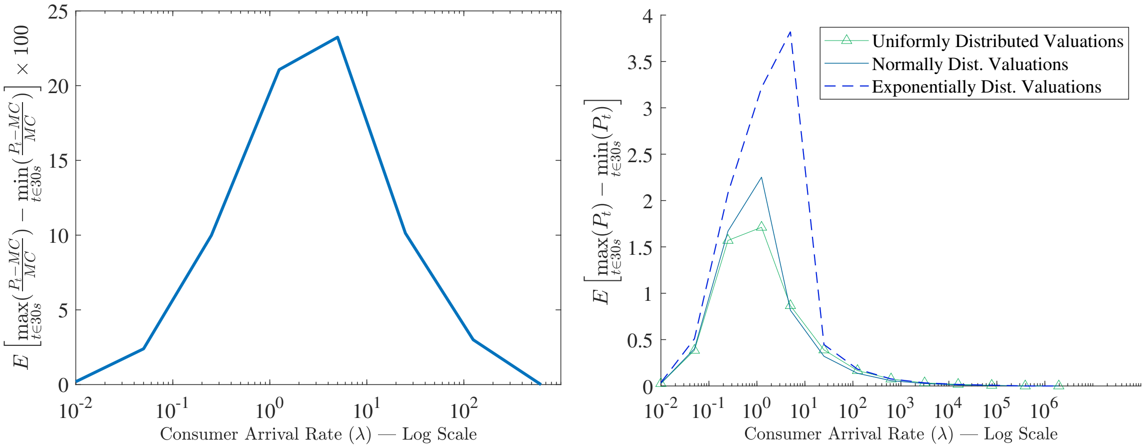

(a) Empirical: Netflix

(b) Theoretical Distributions

Notes: This figure shows the expected range of markups and prices offered over a time interval of length $30 \times s$ against the consumer arrival rate, for the empirical distribution (on the left) and the theoretical distributions (on the right).

Online Appendix

A Feasibility of Personalized Pricing Online

The process of visiting a website involves two steps. First, the client (e.g., a consumer's computer or phone) sends a request to the server to send packets (code and files) that comprise the requested website. The client's request includes information about the requester, including cookies and IP address.[23] Consumers can also be required to provide login credentials to access the requested domain.

Thus, before the server sends the client (consumer) the packets constituting the website, it already knows a lot about the consumer. The server knows the consumer's IP address, which reveals the consumer's location, allowing the server to infer local demographics such as average income and to respond to local demand shocks. The server retrieves cookies, which can reveal prior interactions on the client device and login information that reveals prior interactions with the same consumer on other devices. Gleaned information from login credentials includes browsing histories on the site and linked data from third parties (e.g., from Acxiom). Additionally, the server might access third-party cookies, revealing information about the consumer's activities at other websites. All of this information can then be used to create finely targeted prices.

B Individual-Level Demand Estimation in the Context of Netflix

This section summarizes the methods in Shiller (2020) used to form estimates of individual-level demand and profit functions. Publicly available outputs from Shiller (2020) are used to explore the profitability of sticky targeted pricing in Section 4.

[23] See `https://developer.mozilla.org/en-US/docs/Learn/Getting_started_with_the_web/How_the_Web_works`, and `https://developer.mozilla.org/en-US/docs/Web/HTTP/Cookies`

B.1 Data

Data were obtained from ComScore via the Wharton Research Data Service (WRDS). The dataset contains demographics and browsing histories during 2006 for a large representative sample ($\approx 60,000$) of computer users. I collapsed the data to a cross-section, yielding one observation per panelist.

The browsing data are used to form a set of variables that reveal consumers' habits and tastes: (1) the count of visits the user had to each of the 4,600 most popular websites during 2006, (2) total visits to all websites, and (3) the fraction of visits during select time periods and each day of the week.[24]

Additionally, Netflix subscription status is inferred. For a sample of panelists, the panelist's chosen subscription tier (1, 2, or 3 DVDs at a time) is observed directly.[25] For remaining panelists, browsing histories are used to impute whether the user subscribed to any tier of Netflix's services.[26] See Shiller (2020) for a detailed description of the dataset.

B.2 Empirical Model

The estimation procedure includes demand- and supply-side models. Typically, a supply-side model is included to identify marginal costs. In this context, marginal costs are known a priori. The supply-side model is instead used to estimate consumers' mean price sensitivity, which is not identified from the demand-side model alone because Netflix did not change its prices during the observed period.

B.2.1 Demand

Each consumer makes a discrete choice, selecting from the outside good and three tiered Netflix plans: a 1 DVD at-a-time plan for \$9.99, a 2 DVDs at-a-time plan for \$14.99, and a

[24] Initially, the 5,000 most popular websites were selected. Then, some categories of websites were excluded: movie rental chains, pornography, and sites known to host malware.

[25] Netflix did not offer a streaming service during the observed period.

[26] It is assumed that a user subscribed if the user viewed more than 2 subpages per visit to the Netflix domain, on average. A non-subscriber would be unlikely to do so, because a non-subscriber is unable to log in and view subpages available only to subscribers.

3 DVDs at-a-time plan for \$17.99. The conditional indirect utility consumer i receives from tier j of Netflix's services is:

$$u_{ij} = \alpha P_j + \nu_i + \delta_j + \epsilon_{ij}, \tag{5}$$

where P_j denotes tier j's price, and α and $\nu_i + \delta_j$ denote individual i's price sensitivity and intrinsic utility for product tier j, respectively. I normalize δ_1 to zero, because otherwise there are infinite combinations of $\delta_1, ..., \delta_J$ and $\nu_1, ..., \nu_N$ which imply the same intrinsic utilities, implying the model would not be identified. The error term (ϵ_{ij}) is assumed to follow the type 1 extreme value distribution.

The probability consumer i selects tier j equals:

$$s_{ij}\left(\nu_i, \alpha, \delta, P\right) = \frac{\exp\left(\alpha P_j + \nu_i + \delta_j\right)}{1 + \sum_{k \in J} \exp\left(\alpha P_k + \nu_i + \delta_k\right)}. \tag{6}$$

The probability consumer i chooses any inside tier of service, as opposed to the outside good, equals:

$$s_{ij\neq 0}\left(\nu_i, \alpha, \delta, P\right) = 1 - s_{i0}\left(\nu_i, \alpha, \delta, P\right) = 1 - \frac{1}{1 + \sum_{k \in J} \exp\left(\alpha P_k + \nu_i + \delta_k\right)}. \tag{7}$$

The demand-side model is used to construct two sets of moment conditions: (1) ex-ante estimates of subscription probabilities ($\hat{s}_{ij\neq 0}\left(X_i\right)$; described in Section B.2.4) less the corresponding model predictions from Equation 7 ($s_{ij\neq 0}\left(\nu_i, \alpha, \delta, P\right)$), and (2) the aggregate share of consumers choosing each tier ($\hat{s}_j$) less the model's prediction $\left(\int_\nu s_{ij}\left(\nu, \alpha, \delta, P\right) f\left(\nu\right) d\nu\right)$.

B.2.2 Supply

Firm profits are:

$$\pi = \sum_{j \in J} \left(P_j(\theta) - c_j\right) M s_j - \Gamma = \sum_{j \in J} \theta c_j M s_j - \Gamma, \tag{8}$$

where c_j is the marginal cost of tier j, θ is a markup parameter, $P_j(\theta) = (1 + \theta)\, c_j$ is the price of tier j, s_j is the aggregate share of consumers selecting tier j, M is the market size, and Γ denotes fixed cost. Note the fraction markup over cost is the same for all tiers. This assumption follows a conversation with a former Vice President of Marketing at Netflix, who indicated that this was approximately true, by design.

The corresponding first-order condition is:

$$\frac{d\pi}{d\theta} = \sum_{j \in J} c_j \left(s_j + \theta \frac{ds_j}{d\theta} \right) = 0. \tag{9}$$

This first-order condition comprises the final moment condition.

B.2.3 Objective Function and Identification

Consumer preference parameters are estimated by minimizing an objective function comprised of the demand- and supply-side moment conditions. Specifically, the objective function is:

$$G(\alpha, \delta_1, ..., \delta_J, \nu_1, ..., \nu_N) = \left(\begin{array}{l} \sum_{i=1}^{N} \left(\hat{s}_{ij \neq 0}(X_i) - s_{ij \neq 0}(\nu_i, \alpha, \delta, P) \right)^2 \\ + \sum\limits_{j \in J} \left(\hat{s}_j - \int_\nu s_{ij} \left(\nu, \alpha, \delta, P \right) f \left(\nu \right) d\nu \right)^2 \\ + \left(\sum\limits_{j \in J} c_j \left(s_j + \theta \frac{ds_j}{d\theta} \right) \right)^2 \end{array} \right). \tag{10}$$

The first component of the objective function is the squared difference between the ex-ante probability consumer i subscribes to Netflix and the corresponding model prediction, summed across consumers. This first component identifies ν_i: It is apparent from Equation 7 that consumer i's probability of selecting the inside good monotonically rises with ν_i. The second component is the squared difference between the aggregate share known to choose tier j and the corresponding model prediction, summed across tiers. It identifies δ_j: The implied share choosing tier j rises monotonically with δ_j. The last component is the squared first-order condition, from the supply-side model. As explained next, it identifies the mean price sensitivity α.[27]

[27] Note that the model is exactly identified.

Note that there are four sets of terms in the last component of the objective function: θ, c_j, s_j, and $\frac{ds_j}{d\theta}$. Three of these four are fixed: θ, c_j, and s_j are known ex-ante. The markup is estimated from annual financial reports: $\theta = 0.59$.[28] Given the prices of the three tiers [9.99, 14.99, 17.99], this markup implies marginal costs are \$6.28, \$9.43, and \$11.32, respectively. Finally, the aggregate share choosing each tier (s_j) is inferred from the data.[29]

The only remaining terms are $\frac{ds_j}{d\theta}, \forall j$. They depend on consumer preference parameters from the demand-side model, in particular on price sensitivity α. Note that $\frac{ds_j}{d\theta}$ monotonically increases with α, implying that α is identified.[30]

Finally, note that as the scale of α, ν_i, and δ_j jointly increase, the error draws (ϵ) become less likely to impact a consumer's choice. Hence, the scale of α, ν_i, and δ_j reflects the precision of estimated demand. Thus, when one has access to data that allow precise predictions of individuals' choices (i.e., ex-ante individual subscription probabilities ($s_{ij\neq0}(X_i)$) are close to either 0 or 1) then this will be reflected in the model by larger estimates of α, ν_i, and δ_j.

B.2.4 Ex-ante Estimates of Individual Subscription Probabilities

The first component of the objective function is the sum of squared differences between ex-ante estimates ($\hat{s}_{ij\neq0}(X_i)$) and model predictions ($s_{ij\neq0}(\nu_i, \alpha, \delta, P)$) of individual subscription probabilities. Hence, before estimating the model, one must first estimate the probability each consumer subscribes.

The probability that each individual subscribes to Netflix is estimated using a lasso-penalized logit model. Specifically, the penalized log-likelihood function equals:

$$\ell(\phi, \beta) = \sum_{i=1}^{N} \ln\Big(s_{ij\neq0}(X_i) \times 1\,(buy) + \big(1 - s_{ij\neq0}(X_i)\big) \times \big(1 - 1(buy)\big)\Big) - \omega \sum_{k=1}^{K} |\beta_k|, \quad (11)$$

[28] According to Netflix's 2006 financial statement, the costs of subscription and fulfillment were 62.9% of revenues, implying the (constant marginal cost) markup equals $\frac{1}{0.629} - 1 = 0.59$.

[29] Tier choice is observed for a sample of panelists.

[30] The explanation relies on the point that $s_{ij}(\nu_i, \alpha, \delta, P(\theta))$ is determined by the moment conditions from the demand-side model. As α changes, other parameters in the model adjust to keep these two moment conditions satisfied, leaving $s_{ij}(\nu_i, \alpha, \delta, P(\theta))$ unchanged. With s_{ij} fixed, $\frac{ds_j}{d\theta}$ is monotonic in α when there is a common percent markup over costs. See Shiller (2020) for details.

where $1(buy)$ is an indicator for subscription, and $s_{ij\neq 0}(X_i)$ denotes the predicted probability of subscribing:

$$s_{ij\neq 0}(X_i) = \frac{\exp(\phi + X_i\beta)}{1 + \exp(\phi + X_i\beta)}. \tag{12}$$

Parameters to estimate include ϕ, β, and the lasso penalty parameter ω. ϕ, β are estimated by maximizing the in-sample penalized likelihood, and ω is estimated by maximizing the out-of-sample likelihood, using two-fold cross-validation.

The lasso model is estimated on a set of 4,633 normalized variables, including individual's web-browsing and demographic variables. See Shiller (2020) for a detailed analysis of variable importance and additional estimation details.

B.3 Individual Demands

After estimating the model in Section B.2, the next step is to calculate expected static profits from each individual type ($\psi = \nu_i/|\alpha|$) as a function of markup (θ):

$$\pi\left(P(\theta), \psi = \frac{\nu_i}{|\alpha|}\right) = \sum_{j\in J} s_{ij}(\nu_i, \alpha, \delta, P(\theta)) \times (P_j(\theta) - c_j). \tag{13}$$

Figure B1 shows expected profits from each consumer type, both when the firm personalizes the markup for each consumer, and under uniform pricing. Note that the profit gains from personalizing the markup are large for captive consumers (with large $\psi = \nu_i/\ |\alpha|$). The density of consumer types is shown in Figure B1b. Overall, personalizing markups raises profits by 12.99% relative to status-quo uniform pricing, if ignoring impacts of personalized pricing on consumer backlash.[31]

[31]The percent profit increase, 12.99%, incorporates Netflix's fixed cost. Variable costs are assumed to equal the "cost of revenues" from Netflix's 2006 Annual Report, about \$627 million. Fixed costs are assumed to equal "operating expenses," about \$305 million. Revenues were \$997 million. Thus, variable profits were \$370 million and total profits were \$65 million. Multiplying the percent chanSge in variable profits by 370/65 yields the percent change in total profits.

Figure B1: Profits by Type: Static Personalized Pricing

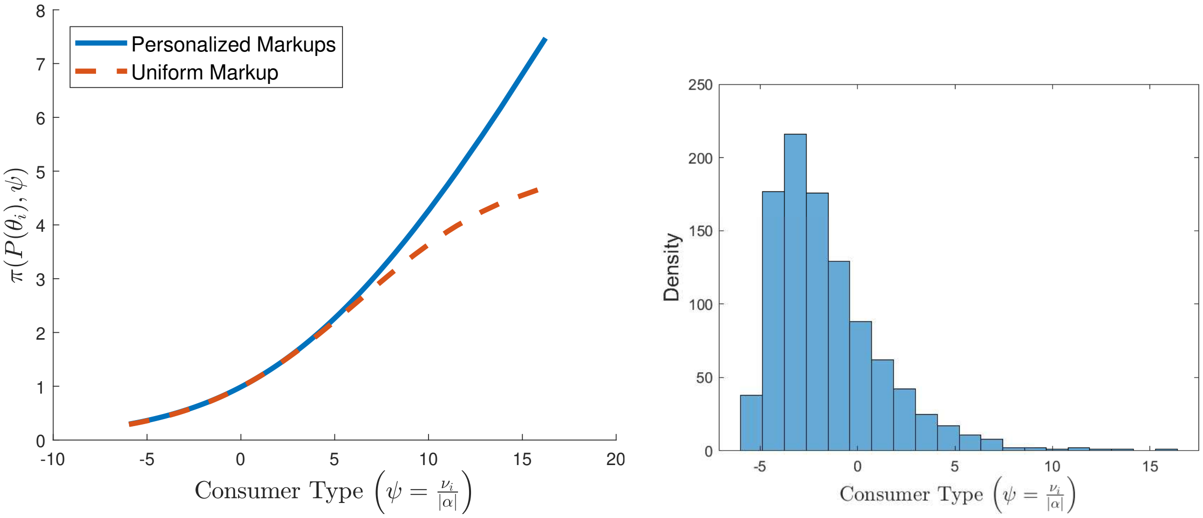

(a) Profits by Consumer Type

(b) Distribution of Consumer Types

Notes: Figure B1a shows the expected profits earned from each consumer type under static personalized pricing (solid line) and uniform pricing (dashed line). Figure B1b shows a histogram of consumer types.

C Supplementary Figures

Figure C1: Simulated Price Range: Time-Normalized—Various Period Lengths

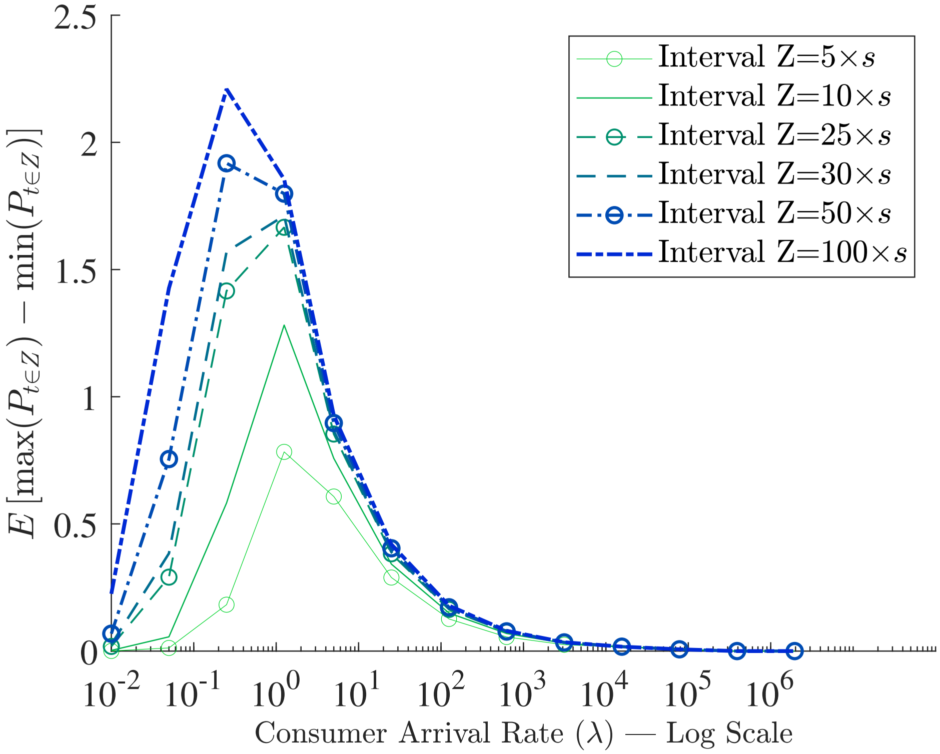

Notes: This figure shows the expected range of prices offered over a a variety of different time interval lengths when the distribution of willingness to pay across consumers follows the uniform distribution.

Figure C2: Simulated Price Change Frequency—Time-Normalized

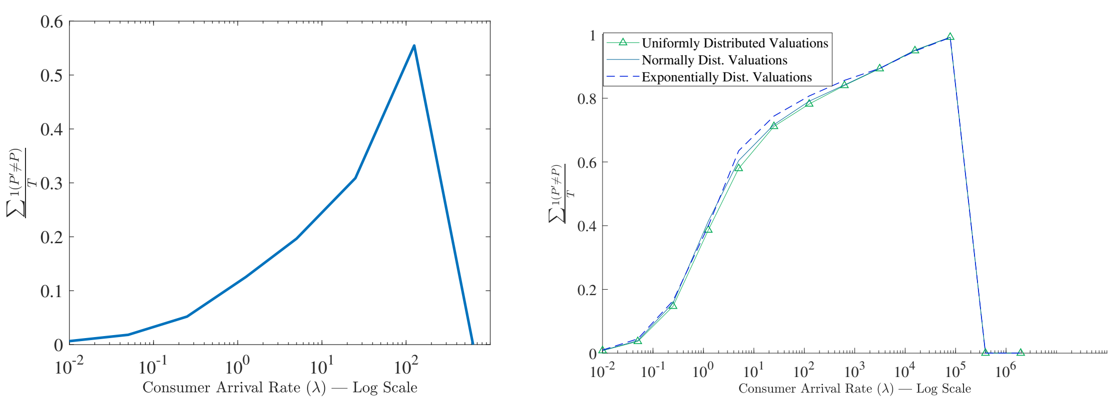

(a) Empirical: Netflix

(b) Theoretical Distributions

Notes: This figure shows the expected number of price changes occurring during an interval of length s (length of the fixed-price interval) for the empirical distribution (on the left) and the theoretical distributions (on the right).