



Dynamic Competition in the Era of Big Data

Patrick J. Kehoe, Bradley J. Larsen, and Elena Pastorino*

Economics Working Paper 22102

HOOVER INSTITUTION
434 GALVEZ MALL
STANFORD UNIVERSITY
STANFORD, CA 94305-6010

February 14, 2022

The advent of rich and highly-detailed information on individual web-browsing and purchase histories—an instance of so-called Big Data—has begun to make feasible sophisticated forms of personalized pricing, heretofore considered too informationally demanding to implement. We argue these pricing strategies are especially relevant in markets for differentiated experience goods. Taking the view that this ability to price discriminate both intertemporally and interpersonally will become increasingly relevant in the future, here we investigate its implications on the dynamics of prices and on efficiency in such markets. In particular, we derive a simple characterization of the equilibrium pricing rule that shows how prices contain a variety-specific dynamic component that depends on the relative informativeness of competing varieties about consumers' tastes. Over time, this pricing rule leads to discontinuous price changes that take the form of fluctuating price discounts for a given consumer. We provide evidence on the gains associated with these sophisticated forms of price discrimination using data on individual consumers' purchases of Apple and Samsung products over time. We estimate primitives in the setting in which firms use the uniform pricing rule we observe in the data and then simulate the counterfactual world with first-degree price discrimination. We find that a significant fraction of consumers are better off under price discrimination relative to uniform pricing, as price discrimination intensifies competition for each individual consumer. Consumers worse off under price discrimination are those who are a good match for only Apple or Samsung. Firm profits from these consumers are correspondingly higher under price discrimination.

Keywords: Dynamic pricing; Oligopoly; Experience goods; Strategic experimentation; Big Data; Price discrimination

JEL Classification: D21, D43, D81, D83

The Hoover Institution Economics Working Paper Series allows authors to distribute research for discussion and comment among other researchers. Working papers reflect the views of the authors and not the views of the Hoover Institution.

* Stanford University. We are grateful to Dirk Bergemann, Aureo de Paula, April Franco, Jan Eeckhout, and Matt Mitchell, as well as seminar and conference participants at Berkeley, Carnegie Mellon, UCL, the Federal Communications Commission, and the 2018 ASSA Meetings for many useful comments and suggestions. We thank Alessandro Dovis, Enoch Hill, and Nitish Vaidyanathan for expert research assistance.

“Historically, first-degree price discrimination has been very difficult to implement, mostly for logistical reasons. With advances in technology and collecting of big data, then it may be that it will become easier to do.” [John Gourville, Harvard Business School Professor of Business Administration, quoted by Forbes (April 14, 2014)]

Until recently, first-degree price discrimination in the form of individual and intertemporal price discrimination has been considered to be rarely used in practice due to the detailed information on individual characteristics and purchasing behavior required to implement it. Times may be changing. Various technologies exist today that allow firms to identify and track individual customers. The corresponding advent of large datasets on individual characteristics, purchasing, and web-browsing behavior—so-called “Big data”—together with the increasing role of customer feedback and social media in disseminating information about consumers’ experiences, has dramatically expanded the possibility for firms to engage in personalized pricing.

For example, in *Forbes* Quentin Gallivan, the CEO of a provider of business analytics software, Pentaho Corp, recently explained how retailers are “*using [B]ig data to analyze tweets, reviews, and Facebook likes and matching this data against customer lists, transactions, and loyalty club memberships*” to target price and product combinations to individual consumers. (See Gallivan (2012).) Forms of personalized deals, in-store and on-line, are now common at major U.S. grocers, pharmacies, department stores, and for magazine and newspaper subscriptions, telecommunications, banking, and credit card services.¹ Indeed, a growing industry is engaged in gathering the big data sets necessary for such personalized pricing, developing the relevant algorithms to process data in real time and implementing such sophisticated pricing strategies. (See Fudenberg and Villas-Boas (2007, 2012).)

The growth in Big data has led empirical researchers to begin addressing forward-looking questions about the effects of this form of price discrimination. One finding in the literature is that as personalized pricing becomes more feasible, potentially large increases in profits are possible. For example, in their study of the pricing of digital music based on survey data on consumers’ valuations, Shiller and Waldfogel (2011) find that personalized pricing can increase revenues by over 50%. Using data linking detailed information about consumers’ web-browsing histories and demographic characteristics, Shiller et al. (2014) shows that even simple personalized pricing schemes can raise revenue by over 12%.²

¹Founded by Sam Odio, the former product manager of photos for Facebook, the startup Freshplum uses machine learning algorithms to develop software for sellers that allows to promote targeted online discounts to narrowly defined categories such as specific geographic areas, repeat customers, or those predicted as unlikely to buy a given product or service.

²While the ability to personalize prices has become increasingly relevant in the digital age, Shapiro (1998) already argued that the online data provider Lexis-Nexis sells to virtually every user at a different price.

This increased ability to price discriminate across consumers and over a consumer's purchases over time have become especially relevant in the context of the rising phenomenon of umbrella branding. As firms market different varieties of goods, personalized pricing allows them to engage in sophisticated forms of price discrimination that take into account consumers' correlated taste for the varieties of a firm's brand. For instance, it is well documented that consumers' experiences about the match of their tastes with one variety in a line of products are positively correlated with their perceptions of their match with other varieties in the same line or related lines. See Erdem (1998) and Erdem and Chang (2012) for evidence of this phenomenon.

Here we take an extreme forward-looking view by supposing that personalized pricing is feasible and analyze the resulting equilibrium pricing patterns and its efficiency properties. We are well aware that at the present date such strategies are only slowly becoming feasible. But our contention is that given the potential profit margins associated with personalized pricing and the availability of new technologies that allow firms to practice it, over time many markets are likely to move closer to these pricing schemes. In this sense, this paper analyses a world towards which we are headed rather than the one we are living in today.

Specifically, we focus on optimal personalized pricing in a market for branded experience goods. The *experience* quality of these goods is due to consumers gaining information about their tastes for these goods only by consuming them. In these markets, a consumer's taste for the products of a given firm may in large part reflect the quality of the idiosyncratic match between the consumer and the firm's products. The *branded* quality of these goods is due to consumers' experiences about the match of their tastes with one variety in a firm's line of products being correlated with their perceptions of their match with other varieties in the same line or related lines.

We are interested in the following questions: First, in these markets, which pricing strategies Big data give rise to? Second, when does personalized pricing lead firms to generate the efficient amount of information about demand in the market? Does the availability of Big data intensify firms' competition and, if so, under which circumstances? Finally, how does the intensity of competition affect pricing and efficiency in such markets? To date, little is known about the answers to any of these questions.

In our model, each firm can produce multiple goods, which we interpret as differentiated varieties of a brand. Consumers have prior beliefs about the quality of the match of their preferences with each firm's brand. If a consumer purchases a variety of a certain brand, then the experience with that variety gives the consumer information about the quality of her match with all the varieties of the firm's brand.

Moreover, each variety provides possibly different amounts of information about the quality of the match of a consumer's taste and a firm's products: some varieties may be very informative about match quality, whereas others may not. (For example, the purchase of a Samsung TV set, one of Samsung's core products, may be more informative about a consumer's tastes for Samsung's products in general than the purchase of a Samsung Android phone, one extension of Samsung's traditional product line.) We begin with a situation in which consumers' tastes are uncorrelated across brands and then discuss (in Appendix B) the case in which tastes are correlated across brands, so that the experience with one brand affects consumers' perceptions of other brands.

In the model, firms compete in a Markovian fashion in a dynamic Bertrand game—we primarily focus on economies with a finite horizon. We follow the classic study of Bergemann and Välimäki (1996), henceforth BV, in assuming that both the purchase history and the experiences of a consumer are public information. In the resulting game, the strategic interaction among firms is complex: firms not only compete directly to attract a consumer in the current period, but also strategically manage the information flow to the consumer. Specifically, by appropriately choosing the prices for its product varieties, a firm can make a certain variety the most attractive and hence control how much is learned about a consumer's taste for its products.³ By so doing, firms can then influence their strategic positions in the continuation game that follows their interactions in any period.

In terms of pricing, in the duopoly context, we show that prices are the sum of the standard static Bertrand pricing terms, namely, the utility difference between the purchased variety and the second-best variety in the market, and of a *compensating price differential*, which compensates the consumer for the value of the forgone information about the match with the selling firm's competitor that the consumer would have acquired by purchasing from the competitor.

Given this pricing rule, our model can generate rich patterns of seemingly random price increases and decreases for a given variety. Suppose, for instance, that as beliefs about a consumer's taste evolve, the selling firm finds it optimal to offer a given variety but its best competitor switches from offering a less informative variety to offering a more informative one. Then, in this case, the price of the selling firm's variety will discontinuously *decrease*. If at a later point, the competitor switches from offering a more informative variety to offering a less informative one, then the price of the selling firm's variety will discontinuously *increase*. Of course, if the competitor has only one variety then no such pattern can arise. More generally, the frequency and size of these price changes depend on the number of varieties of the

³ That firms offer different product and price combinations to different consumers based on consumers' purchase histories has been documented in the literature; see the reviews by Fudenberg and Villas-Boas (2007, 2012).

best competitor and their informativeness.⁴

In terms of efficiency, in the duopoly context, we show that two restricted efficiency properties hold. These properties help make clear when efficiency occurs and when it does not, the particular type of inefficiency that prevails. The first property, *match efficiency*, is that the variety *offered* by each firm—that is, the variety that each firm would induce the consumer to purchase if it was the selling firm—maximizes the sum of the value of the match between that firm and the consumer. The second property, *conditional efficiency*, is that the equilibrium solves a restricted planning problem, in which the planner can choose which firm produces but is restricted to choosing one of the two varieties preferred by each firm. Hence, inefficiencies can arise only if the match-efficient choices of varieties by firms do not coincide with those of the planner.

We rely on these results to clarify existing efficiency results in the literature and examine the extent to which they can be extended. First, we note that if each firm sells a single variety, as in BV, then conditional efficiency implies efficiency and we obtain the BV’s result. Second, we show that if all varieties within a brand have the same informativeness, then the firms offer the same varieties as the planner would, so that conditional efficiency again implies efficiency. This latter result shows how the BV’s efficiency result can be extended to the case of many varieties. In this latter case, when the informativeness of varieties differs across brands, (noncompetitive) dynamic pricing is necessary to support the efficient outcome.

These efficiency results, however, are sensitive to the information content of the varieties of a given brand. Indeed, when a firm’s varieties are differentially informative, the equilibrium may induce firms to either underprovide or overprovide information. The reason why in equilibrium information can be underprovided is straightforward: since information is valued by all firms but is produced only by the selling firm, a standard free-rider problem arises. Interestingly, in equilibrium information may also be *overprovided*. For instance, a firm may strategically induce a consumer to purchase a more informative variety to protect its selling position (and profits) relative to an efficient outcome in which a less informative variety is selected. This outcome occurs when the experience of a less informative variety is likely to induce a consumer to switch to the product of a competitor.

We also analyze efficiency in the oligopoly context (in Appendix C). Here, we find that match efficiency still holds but conditional efficiency may fail. In particular, we construct an example in which three firms produce one variety each and equilibrium is inefficient. Intuitively, in equilibrium the three firms

⁴For instance, Oh and Lucas Jr (2006) provide evidence that prices in online markets for computer components and peripherals do not always monotonically decline or increase over time but rather exhibit rich patterns of increases and decreases. In general, significant temporal price variation has been empirically detected for several product categories and attributed to consumers learning about their tastes for a brand (see the references in Villas-Boas and Villas-Boas (2008) among others).

can be ranked according to the (expected present discounted value of) utility that their offers entail for the consumer. Although at any node the third–best firm can express its desire to be the winning firm through its price offer, the third–best firm has no means of expressing its preferences over the *remaining* firms. The crux of this example is that at the inefficient equilibrium, the third–best firm could profitably compensate the best firm for the cost of losing to the second–best firm, but has no way to do so in equilibrium. The third–best firm has an incentive to do so when selling by the second–best firms makes it more likely for the third–best firm to become the selling firm in the future.

Finally, we show that if competition is made sufficiently intense, efficiency is restored. An example of how new internet technologies lead to more intense competition is given by Athey and Mobius (2012), who argue that web browsers, search engines, aggregators, and social network enable consumers to virtually costlessly move between firms and, hence, dramatically increase consumer switching and thus competition between firms. One natural way to intensify competition is to allow for the free adoption of each technology so that all firms face at least one competitor with an identical technology. Another way is to allow each firm to sell a possibly unique collection of varieties but to let each variety be produced by at least two firms in the market. We can interpret this second scenario as corresponding to retail stores that offer unique portfolios of various branded consumer products. In both scenarios, efficiency is restored and pricing is static. In the second scenario, efficiency occurs despite the heterogeneity of firms in the market. Hence, competition among identical firms is not necessary for efficiency. Moreover, we show that competitive pricing is not sufficient for efficiency. Lastly, we argue that in a precise sense, these scenarios entail the minimal intensity of competition among firms for equilibrium to be efficient. If either type of competition fails, by only one firm or one variety, equilibrium is typically inefficient.

We consider an application of our model to the market for smartphones and tablets. To this purpose, we have collected data on all new–in–box sales of these products by Apple and Samsung from posted–price transactions on eBay over the period between 2014 and 2017 for several hundred thousand buyers. This data allows us to track individuals’ purchase histories on this platform over time and, for a subset of the data, consumers’ ratings of the products purchased. Based on these data, we recover consumer preference parameters and the distribution of utility signals that consumers face upon experiencing these products. We then use the estimates of these parameters, together with known estimates of the cost structures and manufacturer prices of Apple and Samsung, to compare consumer surplus, producer surplus (variable profits), and total welfare under the current no–discrimination (uniform pricing) regime and under a counterfactual scenario in which firms are allowed to engage in the forms of price discrimination characterized

here. We find that a significant fraction of consumers benefit from the introduction of price discrimination. Specifically, consumers who benefit are those with relatively similar beliefs about their tastes for Apple’s products or Samsung’s products, whereas consumers who are harmed are those who have a high taste for the products of only one firm. This latter group is more “captive” and, correspondingly, firms’ profits from these consumers are higher under discriminatory pricing than under uniform pricing.

Related Literature. Our model is a direct extension of the classic model by BV on pricing and efficiency with experience goods. As noted, we follow BV in assuming that firms price discriminate across consumers by charging prices that depend not only on a consumer’s past purchases but also on the consumer’s experiences with purchased products. Thus, firms and consumer have the same information. BV focus on a duopoly setting in which each firm produces only one variety. Their key results are that equilibrium is efficient and the prices of purchased varieties are static, namely, that the equilibrium pricing rule coincides with the familiar one from static Bertrand games. The winning firm sells at a price equal to the difference in current expected utility from its product and from that of its competitor.

Our model extends BV in several relevant directions. First, in the duopoly context, we allow each firm to produce many varieties of a brand and each variety to convey information about the other varieties of a brand. Second, we also analyze an oligopoly context. In both contexts, we highlight how results differ in two cases: when a consumer’s tastes are uncorrelated across brands and when tastes are correlated across brands. Below we discuss the relation of our work to Bergemann and Välimäki (2006a), henceforth BV2, who consider a deterministic economy with time-varying payoffs meant to capture habit formation or learning-by-doing.

We focus on an economy with experience goods in which firms compete for consumers. For studies of the optimal pricing of experience goods in a monopoly context, see Bergemann and Välimäki (2006b) and Bonatti (2011). Bergemann and Välimäki (2006b) find that, depending on market characteristics, one of two types of pricing paths are optimal: either prices are declining over time or increasing over time. Bonatti (2011) analyzes optimal menu pricing when consumers have partly private and partly common valuations. He shows that optimal pricing is monotone: the monopolist initially charges low prices, sacrificing short-term revenues to increase sales, and subsequently, as more information is revealed, increases its prices. Our model of competition between firms produces richer patterns.

In our paper, we assume that the matches between firms and consumers are idiosyncratic, so that observing the experience of another consumer has no information value. Bergemann and Välimäki (2000) consider an environment in which different consumers’ preferences for firms have a common component.

In such an environment, the information externalities that arise across consumers lead to an inefficiency of a type very different from the one we consider here. Briefly, individual consumers does not internalize the informational benefits that their actions have on other consumers.

Eeckhout and Weng (2015) also consider a model of strategic pricing in which consumers' preferences have a common component. They extend Bergemann and Välimäki (2000) by showing that the Markov perfect equilibrium with cautious strategies is efficient if and only if the signal-to-noise ratios of the products of the two firms whereas if the ratios are different, there will be excessive experimentation. More generally, they provide a novel condition on dynamic payoffs that must be satisfied whenever the common value experimentation problem has a continuous increment component as it does with Brownian motion.

Our paper is also related to the literature on strategic experimentation in many-player common-value extensions of a standard experimentation problem, the two-armed bandit problem. Bolton and Harris (1999) show that in these settings, when players can learn from the experiments of others, an information externality arises that leads to an inefficiently low amount of information being acquired in equilibrium. Further work by Keller et al. (2005), Keller and Rady (2010), and Klein and Rady (2011) shows how different versions of this environment and different strategies can lead to outcomes in which this inefficiency is ameliorated. The fundamental difference between our model and those in this literature is that in our model, firms endogenously choose prices for varieties and consumers optimally purchase varieties that yield the highest utility, taking into consideration both the price and the information value of each variety. As we show, these forces can sometimes completely restore efficiency. Moreover, when they do not, these forces may lead to overprovision of information in equilibrium rather than just underprovision.

In our model, a firm can produce a variety of goods all of which are informationally related. Hence, information learned about the match with a firm is portable to some extent to other varieties of products. Moreover, as consumers and firms learn about their matches, firms' optimal strategies push consumers towards different varieties based on the information about their tastes learned through experience. In interesting related work, Mitchell (2000) analyzes the dynamics of firm size and scope in a setting in which firms perform a number of informationally related tasks and knowledge in one task is partially portable to the operation of other tasks. Moreover, a firm chooses which tasks to undertake at a point in time, learns from that experience, and then chooses a new set of tasks to operate the following period. The information relation across the tasks to produce products in a firm in Mitchell's model has parallels with the information content of varieties that firms produce in our model.

Here, we have abstracted from consumers' private information about their tastes and firms' private signals about their demand. For a model of behavior-based price discrimination in which consumers have private information about their tastes, see Calzolari and Pavan (2006). The issues addressed in their work, such as whether committing to ignore a consumer's record of past purchases is beneficial for a firm, are complementary to those analyzed here.

Other papers allow firms to have private information about consumer demand. For example, Hellwig and Veldkamp (2009) connect these results to the complementarity of the actions of firms. Their general setup implies that information sharing is optimal for Bertrand competitors whose prices are strategic complements. Angeletos and Pavan (2007) analyze equilibrium and efficiency in a class of large economies with externalities and heterogeneous information across agents. They show that in a Bertrand context, firms are better off when the precision of both public and private information increases. They interpret their result as implying that information sharing is optimal under Bertrand competition.

Finally, in this work, including our own, the initial distribution of types is taken as fixed. For recent work that endogenizes the entry of different types in a market, see Atkeson et al. (2014).

1 Background on Pricing Based on Big Data

The technological assumptions we make in the paper on firms' abilities to personalize prices and on their information about consumers are those of a world we are moving towards rather than of a world that currently exists. Here we briefly argue how even now, given the potential for increases in profits, some firms are starting to make progress in adopting such technologies and using sophisticated pricing strategies. We discuss evidence in three areas: the ability of firms to personalize prices, the ability of firms to acquire information about their own consumers' tastes, and the ability of firms to acquire information about their own consumers' tastes for their competitors' products.

Technological Ability to Personalize Prices. Various technologies exist today that allow firms to identify and track individual customers and, in so doing, offer them personalized prices. For example, the online data provider Lexis-Nexis sells to virtually every user at a different price (Shapiro and Varian (1999)). Choudhary et al. (2005) explain how Amazon experimented with offering different prices to different consumers on its popular DVD titles (Morneau (2000)). Although this experiment was short-lived due to a consumer backlash, Amazon has since found alternative innovative ways of implementing personalized pricing, through the use of the "Gold Box." Each consumer is provided access to a prominently

displayed Gold Box with their name (e.g., John Doe's Gold Box) on web pages at Amazon. Opening the Gold Box provides access to a limited number of products with special discounts that are not available outside the Gold Box. Since the items offered in the Gold Box are different for different consumers, this allows Amazon to charge personalized prices. This is an example of the continuing evolution of personalized pricing and an indication of the likely use of such pricing by other online retailers. Chen and Iyer (2002) already mention several other examples of personalized pricing, including by major providers of long-distance telephone service (such as AT&T, MCI, and Sprint), direct marketing companies like Land's End and L. L. Bean, who have individual specific catalog prices, and financial services and banks, who engage in personalized pricing through personalized discounts on card fees. (Zhang (2003) mentions Wells Fargo and MBNA.)

More recently, the retail grocery store chain Stop & Shop introduced a mobile application, run by the personalized digital media company Catalina, that allows shoppers to scan products. When they do, Catalina identifies them through their frequent shopper number or phone number, and locates where in the store consumers are. Special e-coupons are then created on the spot to induce consumers to buy related products for which consumers may have a similar taste. *"If someone is in the baby aisle and they just purchased diapers,"* said Todd Morris, president of Catalina, *"we might present to them at that point a baby formula or baby food that might be based on the age of their baby and what food the baby might be ready for."* (NYT August 10, 2012, "Shopper Alert: Price May Drop for You Alone.")

Technological Ability to Acquire Information About Consumers' Tastes for Own Products. The ability of firms to acquire information about consumers' tastes is rapidly advancing. Almost every major retailer, from grocery chains to investment banks to the U.S. Postal Service, now has a "predictive analytics" department devoted to understanding not just consumers' shopping habits but also their personal habits. (NYT Magazine February 19, 2012, "How Companies Learn Your Secrets"). A prime example is Target, which has collected vast amounts of data on every person who regularly walks into one of its stores. Once a consumer is assigned a Guest ID, a very rich set of individual-specific characteristics is linked to it. More importantly, through a combination of in-house collection and purchasing from outside digital media vendors, Target routinely obtains data on the brands a consumer prefers, say, of coffee, paper towels, cereal or applesauce.

The ability for firms to track individuals' web-browsing and purchasing behavior, to the detail of every web page visited and purchase made over a one to two year period, has originated with the widespread use of "traveling cookies" based on consumers' logins at popular sites, for example, airline sites or Face-

book. Once a customer logs in, the cookie follows the customer wherever he or she surfaces on the web. BlueKai, a startup well-known for collecting and sorting data about consumers, is one of the best known companies providing the software that enables web sites to track their users so as to assign them to specific market segments. BlueKai's customers—which have included travel sites, like Kayak and Expedia, that personalize advertising to individual consumers—track more than 80 percent of the U.S. online population and have created hundreds of millions individual profiles based on what consumers browse and buy online.

Technological Ability to Acquire Information About Consumers' Tastes for Competitors' Product.

The ability to assess consumers' preferences for the products of competitors is beginning to grow. For instance, the Executive Office of the President's 2015 report "*Big Data and Differential Pricing* " recognizes two trends associated with the ever increasing use of big data for targeted marketing and personalized pricing. One trend is related to the application of big data to develop secondary markets in consumer information for personalized ads. The other trend is the widespread adoption of new information technology platforms, of which the internet and smartphones are the most important, which have made possible to track users' location via mapping software, their browser and search history, whom and what they like on social networks like Facebook, the songs and videos they stream, their retail purchase history as well as the contents of their online reviews and blog posts. From these data, it is becoming increasingly easy for firms to gather information about a consumer's experiences not just with their own brands and products but also with the brands and products of their competitors.

2 A Duopoly Model

We consider here a duopoly model for consistency with our empirical application. See the Appendix for the multi-product multi-firm case and the case in which consumers' tastes for firms' products are correlated rather than independent across firms. Specifically, we consider a market in which each of two firms has a vector of products, which we interpret as different varieties of the same *brand*. These firms compete in each period of an infinite or a finite horizon for a consumer who has unknown tastes for the two brands. The model is designed to capture the interaction between dynamic Bertrand competition in prices and varieties among firms and the learning process generated by a consumer's experience with the underlying products. In particular, both firms and the consumer evaluate products based not only on their current payoff—profit for the firms and utility for the consumer—but also on their information content.

Here, the state of the world corresponds to the unknown taste of the consumer for the two brands, which summarizes the quality of the match between the consumer’s preferences and the attributes of the two firms’ products. We assume that the quality of the match of the consumer’s preferences with each brand is independent across brands, and we later discuss generalizations. As the consumer experiences the products of the two firms, both the consumer and the firms symmetrically learn about the consumer’s tastes. Thus, the consumer is learning about her taste for the two brands at the same time as the two firms are learning about the consumer’s tastes.⁵ As discussed in the introduction, we think of this model as capturing personalized pricing based on purchase history in a simple way. The idea of personalized pricing is that, given the newly available information–gathering technologies, firms can intertemporally price discriminate across consumers by keeping track of the history of a consumer’s purchases as well as the history of the consumer’s experiences with products, for instance, through consumer reports or direct feedback after purchase. Thus, a firm can charge prices tailored to each individual consumer.

Our analysis immediately extends to allowing for many consumers with idiosyncratic matches with brands as long as firms can price discriminate across consumers. In particular, since the attributes of a firm’s products are known, but the quality of the idiosyncratic match between a consumer’s taste and the firm’s products is unknown, there are no externalities across consumers.

We establish three results. The first is *match efficiency*, namely, that the variety that each firm f desires to induce a consumer to buy, the *offered variety*, maximizes the sum of the values of the firm’s profits and the consumer’s utility. The second is *conditional efficiency*, namely, that conditional on the varieties offered by firms, the equilibrium is efficient. The third is *dynamic pricing*, namely, that the price of a purchased variety not only reflects the quality difference between the purchased variety and the next–best (or second–most–preferred by the consumer) variety, as in a standard static Bertrand model of price competition, but also compensates the consumer for the lost opportunity to learn about her taste for the unpurchased brand.

2.1 Setup

We consider a market in which consumers and two firms $f \in \{A, B\}$ interact over a finite or infinite horizon with periods $t = 1, 2, \dots, T$ with T possibly infinite. Firms and consumers discount the future by the factor $\delta \in (0, 1)$ and normalize the period payoffs of the firms and the consumer by $1 - \delta$. Each consumer can have two levels of taste for the products of each firm f , labeled by $\theta_f \in \{\bar{\theta}_f, \underline{\theta}_f\}$ and

⁵Symmetric observability is commonly assumed in the literature (see BV among many others). We interpret this assumption as an extreme that is becoming increasingly relevant as we enter a world with richer and richer data.

referred to as a *good* match and a *bad* match for the brand of firm f , respectively. Thus, the unknown state of the world is the quality of the matches of the consumer with the two brands, $\theta = (\theta_A, \theta_B)$. We imagine that there is an arbitrary number of consumers but that each consumer independently draws an idiosyncratic taste for each firms' products. Given that firms can charge each consumer a personalized price, the game with two firms and many consumers decomposes into separate games with two firms and a single consumer in each. For ease of notation only, we focus on one of these component games with two firms and a single consumer throughout.

Firms differ in their products. Firm f has K_f varieties, indexed by $k \in \{1, 2, \dots, K_f\}$. Each variety k of firm f leads to high and low realized levels of utilities in any period $t \geq 1$, $X_{fkt} \in \{X_{fHk}, X_{fLk}\}$ with $X_{fHk}, X_{fLk} > 0$, referred to as *success* and *failure*, and is characterized by probabilities of success given the quality of the match. Let α_{fk} denote the probability the consumer receives a high utility given that the match with the brand of firm f is good, and let β_{fk} denote the probability the consumer receives a high utility given that the match with the brand of firm f is bad. We let $\pi_t = (\pi_{At}, \pi_{Bt})$, $t \geq 1$, with $\pi_{At}, \pi_{Bt} \in [0, 1]$, denote the beginning-of-period *prior* or *belief* vector that the match of the consumer and firms A and B is good. We imagine that all goods are produced at a constant marginal cost which for simplicity we have set to zero.

Let a consumer's current period expected gross utility from purchasing variety k of firm $f = A, B$ at prior π_t be given by

$$x_f(\pi_{ft}, k) \equiv \pi_{ft}X_{fk}(\bar{\theta}_f, k) + (1 - \pi_{ft})X_{fk}(\underline{\theta}_f, k), \quad (1)$$

where $X_{fk}(\bar{\theta}_f, k) = \alpha_{fk}X_{fHk} + (1 - \alpha_{fk})X_{fLk}$ and $X_{fk}(\underline{\theta}_f, k) = \beta_{fk}X_{fHk} + (1 - \beta_{fk})X_{fLk}$ denote, respectively, a consumer's current period expected gross utility conditional on the match with firm f being good or bad. A consumer's current period expected utility from purchasing variety k of firm f at t at price p is

$$u_f(\pi_{ft}, k, p) = x_f(\pi_{ft}, k) - p. \quad (2)$$

Throughout we refer to $x_f(\pi_{ft}, k)$ as *gross utility* and $u_f(\pi_{ft}, k, p)$ as *utility*. Note that if utility is linear in the quality of output, then we can also interpret X_{fHk} and X_{fLk} as the realized qualities of output. We assume that in its dealings with a given consumer, the payoff to a firm is zero if it does not trade with the consumer and the payoff to the consumer is zero if the consumer does not trade with a firm.

In the game we set up, we suppose that in each period, a firm offers a certain variety and a price for that variety. This setup is equivalent to one in which, instead, we allow each firm to offer a vector of prices for all of its varieties. Intuitively, each firm can always induce the consumer to select a desired variety by

simply charging high enough prices for the remaining varieties.

The timing of events in a period is as follows. At the beginning of any period $t \geq 1$, firms A and B simultaneously make take-it-or-leave-it offers, where an offer (k_{ft}, p_{ft}) of firm $f = A, B$ consists of a variety k_{ft} and a price p_{ft} . The consumer can accept the offer of firm A , an event denoted by $d_{At} = 1$, or accept the offer of firm B , an event denoted by $d_{Bt} = 1$, or reject both offers, an event denoted by $d_{At} = d_{Bt} = 0$. Hence, $(d_{At}, d_{Bt}) \in D = \{(1, 0), (0, 1), (0, 0)\}$. If the consumer accepts firm f 's offer, then the consumer trades with that firm, the offered price is paid, and the utility from the traded variety is realized. If the consumer rejects both firms' offers in a period, then each firm and the consumer receive a payoff of zero in the current period and meet again in the following period. The events in a given period t in chronological order are $e_t = (k_{At}, p_{At}, k_{Bt}, p_{Bt}; d_{At}, d_{Bt}; z_t)$, where $z_t = H$ denotes a success of the variety the consumer purchases at t and $z_t = L$ denotes a failure. The history of events up until the beginning of period t is $h_t = (e_1, e_2, \dots, e_{t-1})$.

Let $\pi_t = (\pi_{At}, \pi_{Bt})$ denote the prior at the beginning of period t associated with the history h_t and the initial prior π_1 . Here, the parameters of Bayes' rule depend on the selling firm as well as on the traded product variety. So, if π_{At} is the prior that the quality of the match with firm A is good at the beginning of period t and the consumer purchases variety k from firm A , then when the outcome is $z_t = H$, the updated prior π_{t+1} at the beginning of period $t + 1$ is $\pi_{t+1} = \Pi_{AHk}(\pi_t) = (\Pi_{AHk}(\pi_{At}), \pi_{Bt})$, where

$$\Pi_{AHk}(\pi_{At}) = \frac{\pi_{At}\alpha_{Ak}}{\pi_{At}\alpha_{Ak} + (1 - \pi_{At})\beta_{Ak}}. \quad (3)$$

When the outcome is $z_t = L$, however, the updated prior $\pi_{t+1} = \Pi_{ALk}(\pi_t) = (\Pi_{ALk}(\pi_{At}), \pi_{Bt})$ at the beginning of period t is

$$\Pi_{ALk}(\pi_{At}) = \frac{\pi_{At}(1 - \alpha_{Ak})}{\pi_{At}(1 - \alpha_{Ak}) + (1 - \pi_{At})(1 - \beta_{Ak})}. \quad (4)$$

The functions $\Pi_{BHk}(\pi_t)$ and $\Pi_{BLk}(\pi_t)$ are defined analogously.

We restrict attention to Markov-perfect equilibria in which the state variable at the beginning of period t is π_t . This is the state faced by firms. The state faced by the consumer when choosing which offer to accept is π_t together with the current set of offers $(k_{At}, p_{At}, k_{Bt}, p_{Bt})$. A *Markov-perfect equilibrium* consists of a strategy for the consumer $d_A(\pi_t, k_{At}, p_{At}, k_{Bt}, p_{Bt})$ and $d_B(\pi_t, k_{At}, p_{At}, k_{Bt}, p_{Bt})$, strategies $k_f(\pi_t)$ and $p_f(\pi_t)$ for firm $f = A, B$, and updating rules $\Pi_{fHk}(\pi_t)$ and $\Pi_{fLk}(\pi_t)$ for all f and $k \in K_f$ such that: *i*) given the firms' strategies, the consumer's strategy satisfies the Bellman equation

$U(\pi_t, k_A, p_A, k_B, p_B)$ given by

$$\max_{(d_A, d_B) \in D} \left\{ \sum_f d_f [(1 - \delta)u_f(\pi_{ft}, k_f(\pi_t), p_f(\pi_t)) + \delta EU(\pi'_t | \pi_t, k_f(\pi_t))] + (1 - d_A)(1 - d_B)\delta U(\pi_t) \right\}, \quad (5)$$

where, suppressing the dependence of continuation values on the firms' future offers for simplicity, $EU(\pi'_t | \pi_t, k_f(\pi_t)) = r_k(\pi_t)U(\Pi_{fHk}(\pi_t)) + [1 - r_k(\pi_t)]U(\Pi_{fLk}(\pi_t))$, with $k = k_f(\pi_t)$, is the consumer's continuation value, given the current prior π_t and the chosen variety $k_f(\pi_t)$ of firm f , and $r_k(\pi_t) = \alpha_{fk}\pi_{ft} + \beta_{fk}(1 - \pi_{ft})$ is the probability of a high utility realization when the consumer chooses variety $k = k_f(\pi_t)$ of firm f ; *ii*) given the strategies of both firm B and the consumer, firm A 's strategy satisfies the Bellman equation

$$V^A(\pi_t) = \max_{(k_A, p_A) \in F_A} \{d_A [(1 - \delta)p_A + \delta EV^A(\pi'_t | \pi_t, k_A)] + d_B \delta EV^A(\pi'_t | \pi_t, k_B(\pi_t)) + (1 - d_A)(1 - d_B)\delta EV^A(\pi_t)\}, \quad (6)$$

where $EV^A(\pi'_t | \pi_t, k_A) = r_k(\pi_t)V^A(\Pi_{AHk}(\pi_t)) + [1 - r_k(\pi_t)]V^A(\Pi_{ALk}(\pi_t))$ with $k = k_A$, for each firm $f = A, B$ we have $d_f = d_f(\pi_t, k_A, p_A, k_B(\pi_t), p_B(\pi_t))$, and F_A is the Cartesian product of K_A and $[\underline{p}, \bar{p}]$, with \underline{p} and \bar{p} appropriately chosen constants; *iii*) firm B 's strategy solves the Bellman equation analogous to (6); and *iv*) the updating rules satisfy Bayes' rule (3) and (4).⁶

To interpret (5), note that at prior π_t , given the offers $(k_A(\pi_t), p_A(\pi_t))$ and $(k_B(\pi_t), p_B(\pi_t))$, the consumer can choose to accept firm A 's offer by setting $(d_A, d_B) = (1, 0)$, or accept firm B 's offer by setting $(d_A, d_B) = (0, 1)$, or accept neither offer by setting $(d_A, d_B) = (0, 0)$.

To interpret (6), note that taking as given the consumer's acceptance strategy, $(d_A(\pi_t, \cdot), d_B(\pi_t, \cdot))$, and firm B 's offer, $(k_B(\pi_t), p_B(\pi_t))$, at prior π_t firm A makes an offer, (k_A, p_A) , realizing, of course, that if its offer is not sufficiently attractive relative to both firm B 's offer and the option of no trade, the consumer will not accept it. If the consumer accepts firm A 's offer, then the firm receives p_A from the consumer, the consumer's utility is $u_f(\pi_{At}, k_A(\pi_t), p_A(\pi_t))$, and the firm anticipates that next period the prior will be updated according to (3) and (4) with $k = k_A$. Note also that firm A thinks through the option of purposely making the consumer an offer that will be turned down in order to induce the consumer to choose the competing firm's variety, $k_B(\pi_t)$, and, thus, have the prior updated according to the information content of variety $k_B(\pi_t)$ rather than variety k_A . For simplicity, from now on we suppress the subscript t

⁶In the finite horizon case, all of the strategies, payoffs, and values are indexed by t . We have suppressed this subscript solely for simplicity.

whenever it is unambiguous.

Given our definition of equilibrium, there is a multiplicity of uninteresting Markov–perfect equilibria. This multiplicity is analogous to the one that arises in a static Bertrand game when firms A and B have different costs, say, $c_A < c_B$. The equilibrium typically considered is $p_A = p_B = c_B$. But having both firms charge some lower price p with $c_A \leq p < c_B$ is also an equilibrium as long as the consumer chooses the low–cost producer, firm A , with probability one. In this case, firm B can optimally price below marginal cost anticipating that it will never be called on to sell at that profit–losing price. We follow BV and focus attention on equilibria in which firm B chooses a price at which it would be willing to sell, should firm B be chosen by the consumer against all (equilibrium) expectations. To this end, we impose a restriction on equilibria, in the spirit of trembling–hand perfection, that generalizes the cautious equilibrium restriction that BV use.⁷

Formally, we require that when a given firm, say, firm A , sells to the consumer at prior π , the nonselling firm, firm B , makes an offer that maximizes the utility of the consumer from purchasing from B , subject to the constraint that firm B makes profits at least as large as it does on the equilibrium path. More formally, an (extended) *cautious* equilibrium is defined by the requirement that at π , firm B 's offer $(k_B(\pi), p_B(\pi))$ solves the problem

$$\max_{(k_B, p_B) \in F_B} \{(1 - \delta)u_f(\pi_B, k_B, p_B) + \delta EU(\pi'|\pi, k_B)\} \quad (7)$$

$$\text{s.t. } (1 - \delta)p_B + \delta EV^B(\pi'|\pi, k_B) \geq \delta EV^B(\pi'|\pi, k_A(\pi)). \quad (8)$$

Of course, at the solution to this problem, the constraint (8) will bind with equality so that the losing firm's equilibrium profits will equal those it would make if the consumer accepted its offer. The point of making the variety of firm B solve a maximization problem off the equilibrium path is to pin down its variety choice to be a reasonable one.

Note for later that if we let \bar{U} denote the maximized value of utility in (7), we can use duality to rewrite the problem as

$$\max_{(k_B, p_B) \in F_B} \{(1 - \delta)p_B + \delta EV^B(\pi'|\pi, k_B)\} \quad (9)$$

$$\text{s.t. } (1 - \delta)u_f(\pi_B, k_B, p_B) + \delta EU(\pi'|\pi, k_B) \geq \bar{U}. \quad (10)$$

We refer to problem (9)–(10) as the *dual form* of the cautious restriction.

⁷Note that BV consider a model in which each firm sells a single good and its only decision is the price to charge. Their cautious equilibrium restriction requires that the nonselling firm choose a price so that it is indifferent between selling the good and not selling it. Here, firms also choose which product variety to offer. Our extension requires that the choice of product variety satisfy an optimality condition.

In the following, all omitted proofs and derivation are relegated to the Appendix.

2.2 Analysis of Equilibrium

Here we characterize the equilibrium beginning with some simple preliminary features. We then turn to the efficiency and pricing properties of equilibrium, and finally we present some examples.

2.2.1 Preliminaries

In the analysis that follows, we repeatedly use a *martingale property*, that is,

$$E[x_f(\pi'_f, k_f)|\pi, k] = x_f(\pi_f, k_f), \quad (11)$$

for any $k_f \in K_f$, $f = A, B$, and $k \in K_A \cup K_B$. Note that (11) holds because x_f is linear in π_f and the mean of the posterior is the prior. This property simply states that the expected value of gross utility from any variety k_f of firm f is unchanged after the prior is updated based on the consumer's experience with either that particular variety $k = k_f$ or with any other variety $k \neq k_f$ of firm f or $f' \neq f$. (Observe that this property holds trivially when $k \in K_{f'}$, $f' \neq f$, and the consumer's taste is independent across firms, since in this case $\pi'_f = \pi_f$ after the consumer's experience with a variety of firm f' .)

We now discuss four properties of equilibrium. First, since the consumer can decline both firms' offers and each firm can guarantee that it does not trade, the payoffs of the consumer and the firms are nonnegative. Second, as we show in the Appendix, the consumer trades with a firm at any prior. So all equilibrium states belong to one of two disjoint subsets, which we denote by E_A and E_B , that partition the set of priors: when $\pi \in E_A$, the consumer accepts the offer of firm A , and when $\pi \in E_B$, the consumer accepts the offer of firm B .

Third, note that at $\pi \in E_A$, the problem of firm A can be rewritten as

$$\max_{(k_A, p_A) \in F_A} \{(1 - \delta)p_A + \delta EV^A(\pi'|\pi, k_A)\} \quad (12)$$

$$\text{s.t. } (1 - \delta)u_f(\pi_A, k_A, p_A) + \delta EU(\pi'|\pi, k_A) \geq (1 - \delta)u_f(\pi_B, k_B(\pi), p_B(\pi)) + \delta EU(\pi'|\pi, k_B(\pi)), \quad (13)$$

which requires that the expected present discounted value of the consumer's utility, conditional on the consumer accepting firm A 's offer at π , is at least as high as the expected present discounted value of the consumer's utility, conditional on the consumer accepting firm B 's offer at π . Conditions (12) and

(13) imply that firm A designs its offer so as to obtain the highest expected present discounted value of profits, subject to the constraint that the consumer weakly prefers its offer to firm B 's offer. Clearly, in equilibrium

$$(1 - \delta)u_f(\pi_A, k_A, p_A) + \delta EU(\pi'|\pi, k_A) = (1 - \delta)u_f(\pi_B, k_B(\pi), p_B(\pi)) + \delta EU(\pi'|\pi, k_B(\pi)), \quad (14)$$

so that constraint (13) holds with equality. Otherwise, firm A could raise its price, still attract the consumer, and increase its profits. Thus, the consumer must be indifferent between the two firms' offers.

Fourth, by the dual form of the cautious equilibrium restriction in (9) and (10), at a prior $\pi \in E_B$ at which the consumer accepts firm B 's offer, the problem of firm A is identical to (12) and (13) with \bar{U} , appropriately defined, replacing the right side of (13).

2.2.2 Efficiency Properties

We now turn to analyzing the efficiency properties of equilibrium. We first define our notion of efficiency. We then prove two restricted efficiency properties of equilibrium. The first property is the *match efficiency* of the variety choices of firms, namely, that the variety that each firm offers maximizes the sum of the (expected present discounted) values of the consumer's utility and that firm's profits. The second property is the *conditional efficiency* of the selling firm, namely, that given the varieties offered by each firm, the one that leads to a higher value of gross utility for the consumer is chosen. We show how these results extend existing efficiency results in the literature and help understand the circumstances, identified later, when equilibrium is inefficient.

As for efficiency, consider a planner that in each period can choose both which firm produces and which variety that firm produces. Such a planner solves the following problem:

$$W(\pi) = \max_f \left\{ \max_{k_f \in K_f} \{ (1 - \delta)x_f(\pi_f, k_f) + \delta EW(\pi'|\pi, k_f(\pi)) \} \right\}. \quad (15)$$

We say that an allocation $(f(\pi), k_f(\pi))$, where $f = f(\pi)$ denotes the identity of the producing firm, is *efficient* if it attains the value $W(\pi)$ in (15).⁸ Next, we define the *match value* $W^f(\pi) = V^f(\pi) + U(\pi)$

⁸We can think of this economy as corresponding to the general equilibrium of an economy with a continuum of ex-ante identical consumers, each of whom has an idiosyncratic match value with each firm. In this economy, the consumers own the firms and there is a numeraire good produced by a technology such that an investment of one unit of profits yields one unit of the numeraire good to the consumer. The consumer values the numeraire good in an additively separable way from the experience goods. Then, the consumer's total utility is the sum of $U(\pi)$, which is the value of the consumer's utility from the consumption of the experience goods, and of the sum of the values of firms' profits, which are used to purchase the numeraire good. The value of total utility equals $W(\pi)$. Thus, we can think of the planner as equivalently maximizing the total utility of the representative consumer.

of firm f and the consumer to be the sum of the values of the profits of firm f and the consumer's utility.

Our first result is that the variety offered by a firm maximizes the match value between that firm and the consumer, and that this match value solves the following programming problem:

$$W^A(\pi) = \max\left\{\max_{k_A \in K_A} \{(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)\}, \quad (16)$$

$$(1 - \delta)[x_B(\pi_B, k_B(\pi)) - p_B(\pi)] + \delta EW^A(\pi'|\pi, k_B(\pi))\}, \quad (17)$$

where $(k_B(\pi), p_B(\pi))$ is the equilibrium strategy of firm B . Note that $W^A(\pi)$ is the maximum of two values: the first value (of the inner maximization problem) is the sum of the values of firm A 's profits and the consumer's utility when firm A sells at prior π and the prior is updated using the variety choice of firm A ; the second value is this same value when firm B sells at prior π and the prior is updated using the variety choice of firm B .

Proposition 1. (*Match Efficiency under Duopoly*) *For an arbitrary firm, say, firm A , the variety $k_A(\pi)$ offered by that firm, both when it sells and when it does not sell, maximizes the match value $W^A(\pi)$. Moreover, the match value $W^A(\pi)$ solves (16).*

This match efficiency result combines in a simple way the properties of the best responses of the consumer and a firm. Notice that the strategy of firm B enters the match efficiency problem of A and vice versa. Yet, when there are only two firms, the match efficiency problem of each firm can be reduced to a simpler problem, which we term the *autarky* problem, in which the strategy of the other firm does not enter. The autarky problem of, say, firm A is a restricted planning problem in which the firm maximizes the value of gross utility of the consumer but is restricted to choosing one of its own varieties, that is,

$$\hat{W}^A(\pi_A) = \max_{k_A \in K_A} \{(1 - \delta)x_A(\pi_A, k_A) + \delta E\hat{W}^A(\pi'_A|\pi_A, k_A)\}, \quad (18)$$

$$E\hat{W}^A(\pi'_A|\pi_A, k_A) = r_{k_A}(\pi_A)\hat{W}^A(\Pi_{AHk}(\pi_A)|\pi_A, k_A) + [1 - r_{k_A}(\pi_A)]\hat{W}^A(\Pi_{ALk}(\pi_A)|\pi_A, k_A).$$

We refer to $\hat{W}^A(\pi_A)$ and $\hat{k}_A(\pi_A)$ as the autarky value of output and the autarky variety choice of firm A . We use similar notation for the autarky problem of firm B .⁹

Corollary 1. (*Autarky Result*) *The match value $W^f(\pi)$ equals the autarky value of firm f , $\hat{W}^f(\pi_f)$, and the product variety offered by firm f coincides with the autarky choice $\hat{k}_f(\pi_f)$, $f \in \{A, B\}$.*

⁹In a continuous time labor market framework, Felli and Harris (1996, 2006) show that the value function of the pairwise-team between a firm and a worker coincides with the autarky value of that firm.

Clearly, when a firm is the selling firm it maximizes the value of utility of the consumer over its products. The key step is showing this also holds when the firm is not the selling firm. The step exploits the double indifference property of the equilibrium: Bertrand competition implies that the consumer is indifferent between purchasing from the selling firm and the nonselling firm whereas the cautious equilibrium implies that the nonselling firm is indifferent equilibrium between selling and not. Putting these pieces together gives the result.

We now turn to establishing a second efficiency property of equilibrium. To this end, consider the *conditional* planning problem in which the planner can choose which firm produces but is restricted to choosing one of the two varieties, $\hat{k}_A(\pi_A)$ or $\hat{k}_B(\pi_B)$, that the two firms offer in equilibrium. We let the value of this problem be $W^*(\pi)$, where

$$W^*(\pi) = \max_f \left\{ (1 - \delta)x_f(\pi_f, \hat{k}_f(\pi_f)) + \delta EW^*(\pi'|\pi, \hat{k}_f(\pi_f)) \right\}. \quad (19)$$

Proposition 2. (*Conditional Efficiency*) *The value of gross utility in equilibrium equals $W^*(\pi)$.*

Here, we develop two immediate implications of our efficiency results. The first is that if each firm has only one variety, then the equilibrium is efficient. The second is that if all varieties of each firm are equally informative, then the equilibrium is efficient.

Corollary 2. (*BV's Efficiency*) *If each firm has only one variety, then the equilibrium allocation is efficient.*

The proof is immediate, since in this case conditional efficiency implies efficiency. Note that Corollary 1 is simply a restatement of the main result in BV, namely, their Theorem 1. This corollary covers both the finite and the infinite horizon cases.

Next, consider the case in which all varieties of a given firm have the same informativeness, in that there exist (α_A, β_A) and (α_B, β_B) such that

$$\alpha_{fk} = \alpha_f \text{ and } \beta_{fk} = \beta_f \text{ for } f = A, B \text{ and } k = 1, \dots, K_f. \quad (20)$$

The next corollary extends these efficiency results to an environment in which each firm produces multiple varieties but all of them are equally informative.

Corollary 3. (*Efficiency with Equally Informative Varieties*) *If all varieties of a given firm are equally informative as in (20), then the equilibrium allocation is efficient.*

2.2.3 Dynamic Pricing and Compensating Price Differentials

We use the match value $W^f(\pi) = V^f(\pi) + U(\pi)$ of firm $f = A, B$ to characterize the equilibrium pricing rule in the following proposition.

Proposition 3. (*Dynamic Pricing*) *The price charged by the selling firm, say, firm A, is*

$$p_A(\pi) = x_A(\pi_A, k_A(\pi)) - x_B(\pi_B, k_B(\pi)) + \frac{\delta}{1-\delta} [EW^B(\pi'|\pi, k_A(\pi)) - EW^B(\pi'|\pi, k_B(\pi))]. \quad (21)$$

To help clarify this pricing formula, note that (21) can be rewritten as

$$x_A(\pi_A, k_A(\pi)) - p_A(\pi) + \frac{\delta}{1-\delta} EW^B(\pi'|\pi, k_A(\pi)) = x_B(\pi_B, k_B(\pi)) + \frac{\delta}{1-\delta} EW^B(\pi'|\pi, k_B(\pi)),$$

which simply states that the match value of the non-selling firm is the same both on and off the equilibrium path. In particular, the left side is the sum of payoffs to firm B and the consumer when the consumer purchases from firm A whereas the right side is the sum of payoffs to firm B and the consumer when the consumer purchases from firm B . This result, in turn, follows from the feature of equilibrium that the price of the selling firm makes the consumer indifferent between purchasing from the selling firm and purchasing from its competitor, and the cautious equilibrium restriction that implies that the nonselling firm is indifferent between not selling and selling at the offered price and variety.

By Proposition 3, the price of the selling firm is the sum of two components. The first is a *static component* equal to the difference in the gross utility of the consumer from purchasing from the selling firm and from the nonselling firm. This component corresponds to the familiar static Bertrand pricing rule. The second component is the difference between the nonselling firm's match value of the information acquired when variety $k_A(\pi)$ is purchased and of the information acquired when variety $k_B(\pi)$ is purchased. This term can be interpreted as a *compensating price differential* that compensates the consumer for the foregone opportunity to learn about her taste for the brand of firm B when purchasing from firm A .

We can link the *sign* and *size* of the compensating price differential to the informativeness of the firms' varieties. To see how, note from Corollary 1 that the match value of B is the autarky value $\hat{W}^B(\pi_B)$ and that with independent priors, the variety $k_A(\pi)$, equal to the autarky choice $\hat{k}_A(\pi_A)$, is uninformative about the varieties of firm B . Hence, $E\hat{W}^B(\pi'_B|\pi_B, k_A(\pi)) = \hat{W}^B(\pi_B)$ and the compensating price differential reduces to

$$\hat{W}^B(\pi_B) - E\hat{W}^B(\pi_B|\pi_B, \hat{k}_B(\pi_B)). \quad (22)$$

Next, note that the autarky value function $\hat{W}^f(\pi_f)$ in (18) is convex in π_f . (This result is standard and a proof can be easily adapted from Banks and Sundaram (1992).) Since belief updating induces a mean-preserving spread to the prior and the mean-preserving spread of a convex function increases its value, the second term in (22) is larger than the first term. Hence, this difference is always negative. It is convenient to define variety k to be *more informative* than variety k' , with $k, k' \in K^A \cup K^B$ if at any prior π , the posterior distribution under variety k is a mean-preserving spread of the posterior distribution under variety k' . Hence, we have established the following result.

Corollary 4. (*Compensating Price Differential*) *In duopoly, the compensating price differential of the selling firm is negative and more negative the more informative is the competitor's offered variety.*

This corollary implies that in order to compensate the consumer for the lost information about the brand of the nonselling firm, the selling firm sets a lower price in this dynamic game than it would have in the corresponding static game. Here since priors are independent across firms, buying from one firm, say firm A , means giving up information on tastes for firm B 's products. The selling firm must compensate the consumer for this foregone information. This result, at some level, shows that the model predicts a long-lived version of penetration pricing: as long as consumers tastes are relatively uncertain, firms must price below their static prices. (Note, as we show later, that when priors are correlated across firms the compensating price differential can be either positive or negative.)

Under this pricing rule, the model can also generate a pattern of seemingly random, temporary price discounts for a given product. To see why, imagine that as the prior evolves, the consumer switches back and forth between purchasing from firm A and firm B . In particular, suppose that the consumer buys from firm A at prior π , then buys from firm B at prior π' , and then eventually returns to buy from firm A at prior π'' . Suppose also that firm A offers the same variety throughout. For example, imagine a consumer purchases a Gucci belt, then a Versace sweater, then another Gucci belt and interpret Gucci as firm A and Versace as firm B . In such a scenario, it is easy to generate a pattern of temporary price discounts. Basically, if at prior π'' firm B offers a more informative variety than at prior π , then firm A will find it optimal to offer a lower price for the same product at π'' than at π . In our example, this means that the second Gucci belt will be priced lower than the first one.

Based on Proposition 3, we can revisit the result of BV that the pricing rule of the selling firm is static. It turns out that this result depends on each firm producing only one variety.

Corollary 5. (*BV's Static Pricing*) *If each firm has only one variety, then the pricing rule is static in that*

the price charged by the selling firm, say, firm A , is

$$p_A(\pi) = x_A(\pi_A, k_A(\pi)) - x_B(\pi_B, k_B(\pi)). \quad (23)$$

To prove this result, recall from Corollary 1 that the match value $W^B(\pi)$ reduces to the autarky value $\hat{W}^B(\pi_B)$. In turn, since firms have only one variety, the autarky value of firm B reduces to $x_B(\pi_B)$. Hence, by the martingale property,

$$EW^B(\pi'|\pi, k_A(\pi)) - EW^B(\pi'|\pi, k_B(\pi)) = x_B(\pi_B) - x_B(\pi_B) = 0.$$

An analogous argument holds when firm B is the selling firm. This proves Corollary 5. An obvious generalization of this corollary is that even when the winning firm has multiple varieties, the price of the winning firm is static whenever the losing firm has only one variety.

Finally, consider the case in which all varieties of a brand are equally informative. By Corollary 3, the equilibrium allocation is efficient. The pricing rule of the selling firm, however, is *not* static. To see why, note that since the match value of each firm is convex in π , the expected value $EW^B(\pi'|\pi, k_B(\pi))$ typically differs from $EW^B(\pi'|\pi, k_A(\pi))$. So, in general the compensating differential in (21) is nonzero. Note that this result, together with Corollary 3, also implies that static pricing is not necessary for efficiency.

2.2.4 Examples of Inefficiency

So far we have examined situations in which equilibrium is efficient. Here, we show that when firms produce multiple varieties of different degrees of informativeness about consumers' tastes, equilibrium may be inefficient. We begin by providing some intuition on the nature of this inefficiency. Then, we consider two economies in which equilibrium is inefficient. In the first economy, the market *underprovides* information relative to the efficient allocation whereas in the second economy, the market *overprovides* it.

An Intuition for Inefficiency. The conditional efficiency result of Proposition 2 implies that if firms offer the right varieties, namely, those the planner would choose, then the consumer chooses the same variety as the planner does. Thus, inefficiencies can only be due to firms offering varieties different from the efficient ones.

Why would the firms and the planner choose different varieties? The key reason is that when choosing among its varieties, the selling firm takes into account only how its choice affects its own profits and,

by match efficiency, the consumer's utility. In particular, the selling firm has no incentive to weigh the impact of its choice on the ability of the nonselling firm to attract the consumer in the future and, thus, on the nonselling firm's future profits. In contrast, the planner maximizes the value of the consumer's gross utility, that is, the sum of the values of the two firms' profits and the consumer's utility. Hence, unlike the selling firm, the planner internalizes the impact of the choice of variety by the selling firm on the profit of the nonselling firm. These differing objectives between the selling firm and the planner lead to differing choices of varieties.

To show how inefficiencies emerge, we focus on two simple economies, which show that both underprovision and overprovision of information can occur. These examples admit two interpretations. First, they can be interpreted as instances of two-period economies in which either the future is discounted at rate $\widehat{\delta} = \delta/(1 - \delta)$ or, equivalently, period 1 payoffs are discounted at rate $1 - \delta$ and period 2 payoffs are discounted at rate δ , where period 2 possibly stands for an arbitrarily long future. Second, they can be interpreted as instances of infinite horizon economies in which utility realizations are *dependently* distributed over time. Specifically, in this second interpretation, we let the consumer's gross utility be the same as in our original formulation, but we assume that firms and consumer no longer receive informative signals about the state of the world from period $t = 2$ on.

To understand this interpretation, note that in our environment so far, agents face a compound lottery that depends on the state of the world, θ : when the consumer purchases variety k of firm f , with probability π_t the consumer experiences X_{fHk} with probability α_{fk} and X_{fLk} with probability $1 - \alpha_{fk}$, whereas with probability $1 - \pi_t$ the consumer experiences X_{fHk} with probability β_{fk} and X_{fLk} with probability $1 - \beta_{fk}$. In the two economies considered here, we reduce this compound lottery to a single lottery that is independent of the state of the world from period $t = 2$ onward: with probability $\alpha_{fk}\pi_{ft} + \beta_{fk}(1 - \pi_{ft})$, where $\pi_{ft} = \pi_{f2}$, the consumer experiences X_{fHk} and with the complementary probability the consumer experiences X_{fLk} . Hence, the prior is not updated after period 2. The precise way we formulate this dependently distributed stochastic process for utility realizations is discussed in the Appendix. We find this interpretation useful because it also helps clarify that the results in BV2 on efficiency for a deterministic economy with payoffs that eventually become constant (see their Section 6) do not immediately extend to a multivariety stochastic economy with arbitrary dependently distributed payoffs. In both economies we assume that $X_{fHk} > X_{fLk}$ for all k and f so that the consumer's gross utility increases with the prior.

Economy 1: Information Underprovision. Assume that firm A has two varieties: a perfectly informative variety, variety 2, with $\alpha_{A2} = 1$ and $\beta_{A2} = 0$, and a moderately informative variety, variety 1, with

$1 > \alpha_{A1} > \beta_{A1} > 0$. Firm B has only one variety that, for simplicity, is uninformative, with $\alpha_{B1} = \beta_{B1}$. The initial prior vector is denoted by (π_A, π_B) . Since the probability of success of firm B 's variety does not vary with θ_B , we can simplify notation and denote the consumer's gross utility from the purchase of firm B 's single variety by x_B . We assume that utilities are ordered as in Figure 1,

$$x_A(1, 1) > x_A(1, 2) > x_B > x_A(0, 2) > x_A(0, 1) \text{ and } x_A(\Pi_{AL1}(\pi_A), 1) > x_A(\Pi_{AL1}(\pi_A), 2) > x_B, \quad (24)$$

so that if the consumer purchases variety 1 of firm A in period 1 and experiences a failure, then at the resulting prior, $\Pi_{AL1}(\pi_A)$, variety 1 of firm A is the most attractive in the market.

Consider the planner's problem. If the planner chooses the more informative variety of firm A , variety 2, in the first period, then after a success the planner chooses variety 1 and after a failure it chooses the variety of firm B by (24). If, instead, the planner chooses the less informative variety of firm A , variety 1, in the first period, then in the second period the planner chooses variety 1 both after a success and after a failure by (24). Thus, by the martingale property in (11), the planner's value when it chooses variety 1 reduces to $x_A(\pi_A, 1)$. Assume that the planner in the first period prefers the more informative variety of firm A in that

$$(1 - \delta)x_A(\pi_A, 2) + \delta[\pi_A x_A(1, 1) + (1 - \pi_A)x_B] > x_A(\pi_A, 1), \quad (25)$$

with $x_A(\pi_A, 1) > x_B$.

Consider next the equilibrium. If firm A chooses variety 1 in the first period, then Bertrand competition in the second period implies that firm A wins and sells variety 1 to the consumer in that period as well regardless of the utility realization in the first period. Hence, the match value of firm A when it chooses variety 1 in the first period coincides with the value of the planner, $x_A(\pi_A, 1)$. If, instead, firm A chooses the more informative variety in period 1, variety 2, then after a success firm A sells variety 1 to the consumer (by (24)), but after a failure firm B sells to the consumer at price $p_B = x_B - x_A(0, 2)$, which implies that the match value of firm A after a failure is $x_A(0, 2)$. The match value of firm A when it chooses the less informative variety in the first period exceeds that when it chooses the more informative variety if

$$x_A(\pi_A, 1) > (1 - \delta)x_A(\pi_A, 2) + \delta[\pi_A x_A(1, 1) + (1 - \pi_A)x_A(0, 2)]. \quad (26)$$

Equilibrium is inefficient when (25) and (26) simultaneously hold.¹⁰ The key idea behind inefficiency

¹⁰Parameter values that satisfy these two inequalities as well as (24) are: $x_A(1, 1) = 8.2$, $x_A(1, 2) = 7.5$, $x_B = 6.25$, $x_A(0, 2) = 3.25$, $x_A(0, 1) = 2.5$, $\alpha_{A1} = 0.55$, $\beta_{A1} = 0.45$, $\alpha_{A2} = 1$, $\beta_{A2} = 0$, $\delta = 0.6$, and $\pi_A = 0.75$. Note that α_{B1} , β_{B1} , and π_B play no role in these calculations.

is that the planner values varieties according to the sum of the values they generate for the consumer and both firms, whereas firm A values varieties only according to their match value, namely, their value to firm A and the consumer. This difference in valuations leads the planner and the firms to make different decisions. To see how this difference in valuations shows up, note that for either variety of firm A , the first period payoffs of the planner and of the match of firm A and the consumer agree. The second period payoffs of the planner and of the match of firm A and the consumer following a success also agree. At both of these nodes, firm A sells to the consumer. In the second period after a failure of variety 2, however, firm A loses to firm B : the planner's payoff at this node is x_B , which is the sum of the consumer's utility, $x_B - p_B$, firm B 's profits, p_B , and firm A 's profits, zero, whereas the payoff of the match of firm A and the consumer is only the sum of firm A 's profits, zero, and the consumer's utility, $x_B - p_B = x_A(0, 2)$, which is smaller than x_B .

In sum, information is underprovided in equilibrium because firm A strategically chooses the less informative variety anticipating that by doing so, it will not lose the consumer to firm B after a failure and so will end up making more profits.

Economy 2: Information Overprovision. Assume, as before, that firm A has two varieties: a perfectly informative variety, variety 2, with $\alpha_{A2} = 1$ and $\beta_{A2} = 0$, and a moderately informative variety, variety 1, with $1 > \alpha_{A1} > \beta_{A1} > 0$. Firm B has only one variety that is now moderately informative. The initial prior vector is denoted by (π_A, π_B) . We assume that utilities are ordered as in Figure 2,

$$x_A(1, 1) > x_B(1, 1) > x_A(1, 2) > x_A(0, 2) > x_B(0, 1) > x_A(0, 1), \quad (27)$$

and that a failure of variety 1 of firm A implies that firm B has the most attractive variety,

$$x_B(\pi_B, 1) > x_A(\Pi_{AL1}(\pi_A), 1) > x_A(\Pi_{AL1}(\pi_A), 2). \quad (28)$$

Consider the planner's problem. Under (27) and (28), if the planner chooses either variety 1 or 2 of firm A in the first period, then after a success the planner chooses variety 1 in the second period, whereas after a failure it chooses the variety of firm B . Assume that the planner prefers the less informative variety in that

$$(1 - \delta)x_A(\pi_A, 1) + \delta\{r_1(\pi_A)x_A(\Pi_{AH1}(\pi_A), 1) + [1 - r_1(\pi_A)]x_B(\pi_B, 1)\} > (1 - \delta)x_A(\pi_A, 2) + \delta[\pi_A x_A(1, 1) + (1 - \pi_A)x_B(\pi_B, 1)]. \quad (29)$$

Consider now the equilibrium. Suppose that firm A chooses the perfectly informative variety, variety 2, in the first period. Under (27), after a success firm A sells variety 1 to the consumer in the second period, whereas after a failure firm A loses the consumer to firm B , which sells to the consumer at price $x_B(\pi_B, 1) - x_A(0, 2)$. Hence, the payoff to firm A and the consumer in the second period is $x_A(1, 1)$ after a success but, unlike the planner's payoff of $x_B(\pi_B, 1)$, it is only $x_A(0, 2)$ after a failure and equals the consumer's utility, $x_B(\pi_B, 1) - p_B$, from purchasing from firm B .

If, instead, firm A chooses variety 1 in the first period, then after a success firm A retains the consumer and sells variety 1, but after a failure firm A loses the consumer to firm B and the payoff to firm A and the consumer is just the consumer's utility, $x_B(\pi_B, 1) - p_B = x_A(\Pi_{AL1}(\pi_A), 1)$. Hence, the match value of firm A from choosing variety 1 in the first period is $x_A(\pi_A, 1)$ by the martingale property (11). We assume that the match value of firm A in the first period when it chooses the more informative variety, variety 2, exceeds that when it chooses the less informative variety, variety 1, so that

$$(1 - \delta)x_A(\pi_A, 2) + \delta [\pi_A x_A(1, 1) + (1 - \pi_A)x_A(0, 2)] > x_A(\pi_A, 1). \quad (30)$$

When (29) and (30) are simultaneously satisfied, equilibrium is inefficient.¹¹ As in our first example, by comparing (29) with (30), we see that at the nodes at which firm A sells, the payoff to the planner and to firm A and the consumer agree. But at the nodes at which firm A loses the consumer to firm B , namely, in the second period after a failure of either variety, the planner's payoff is $x_B(\pi_B, 1)$, which is the sum of the consumer's utility, $x_B(\pi_B, 1) - p_B$, firm B 's profits, p_B , and firm A 's profits, zero. Instead, at these nodes the payoff to firm A and the consumer is just the consumer's utility, respectively, $x_B(\pi_B, 1) - p_B = x_A(0, 2) < x_B(\pi_B, 1)$ after a failure of variety 2 and $x_B(\pi_B, 1) - p_B = x_A(\Pi_{AL1}(\pi_A), 1) < x_B(\pi_B, 1)$ after a failure of variety 1.

Thus, information is overprovided because firm A strategically chooses the most informative variety by placing less value than the planner does on the nodes at which firm B is the selling firm.¹²

¹¹Parameter values that satisfy (29) and (30), as well as assumptions (27) and (28), are: $x_A(1, 1) = 14.2$, $x_B(1, 1) = 14.1$, $x_A(1, 2) = 9.1$, $x_A(0, 2) = 9.05$, $x_B(0, 1) = 9$, $x_A(0, 1) = 1.5$, $\alpha_{A1} = 0.085$, $\beta_{A1} = 0.082$, $\alpha_{A2} = 1$, $\beta_{A2} = 0$, $\delta = 0.78$, $\pi_A = 0.85$, and $\pi_B = 0.83$.

¹²Note that Felli and Harris (2006) extend Felli and Harris (1996) by adding a training program to the technology of each firm. This program produces constant output regardless of a worker's ability but participating in it nonetheless generates information about ability—for instance, because of the monitoring by supervisors. They show that training is overprovided in a stochastic learning-by-doing version of this model in which training affects the mean but not the variance of beliefs about a worker's unknown productivity.

3 Application to Smartphone and Tablet Market

We turn here to our empirical application of competition between Apple and Samsung in the smartphone and tablet market.

3.1 Data Description

We apply our theory to the market for smartphones and tablets produced by Apple and Samsung. In recent years, these two firms have been the largest in the smartphone market, with approximately a 20% market share each. We collect data on buyer purchases of new-in-box smartphones and tablets from posted-price transactions in eBay.com from 2014–2017. We focus on items in these categories that have been matched by the seller to an item in one of several commercially available catalogs for U.S. cell phones and tablets. These catalogs also contain average user reviews (on a scale of 0 to 5) for each product. Examples of items in our dataset are “Apple iPhone 5S 16GB”, “Apple iPad 2 16GB, Wi-fi,” “Samsung Galaxy Prevail,” or “Samsung Galaxy Note SGH-I717 16GB.” As these are all new-in-box items readily available at retail outlets, we abstract from any pricing decision of individual sellers, who are unlikely to have any market power in this setting—the median seller only sells one item in our data—and, instead, treat buyers as facing the brand-new device price plus a random eBay seller discount, as we describe below. Importantly, we only use eBay consumers’ decisions to estimate demand parameters—for instance, the X_{fHk} ’s and X_{fLk} ’s from above for each product—and do not estimate any parameter related to the supply side of the market.¹³

We have transaction-level data for new-in-box smartphones and tablets on eBay. Our sample consists of 28,640 cases in which a given buyer purchased a smartphone or tablet and then, at some point in the future, purchased another smartphone or tablet. Each observation contains information on the product purchased, the price paid, and the buyer id. An advantage of this data is that it allows us to track multiple purchases for an individual buyer over time. For example, we can measure the likelihood of purchasing an Apple product conditional on purchasing this brand in the past. Clearly, the transactions in our dataset are not exhaustive, but still represent a fairly large sample that displays variation across consumers in the variety of products purchased, that is, phones rather than tablets, and their brands. We also observe a smaller dataset consisting of 1,279 cases—79 for Apple’s products and 1,200 for Samsung’s products—in which a given buyer purchased *and rated* the same product twice in a row.¹⁴

¹³For example, we do not estimate marginal costs from a pricing first-order condition, but instead take marginal cost estimates from industry tear-down cost reports.

¹⁴These ratings are a relatively new feature of eBay, and are actual *product* ratings left by buyers, similar to product ratings

It is important to note that the purchase and rating behavior of these eBay buyers may not reflect that of the broader smartphone and tablet market. Even among the eBay users in our data, we are only able to observe their phone and tablet purchases on eBay, not through other outlets. We impose two sample restrictions that help us focus on buyers whose behavior may be closer to those of off-eBay buyers. First, to eliminate buyers who are only looking for older-model phones, we restrict our attention to buyers purchasing a type of phone or tablet, for instance, “Samsung Galaxy Note SGH-I717 16GB”, within six months from the time that particular device first appeared in eBay transaction data during our sample period. Second, to eliminate buyers purchasing large quantities of devices over short time periods—potentially to re-sell them, rather than as final consumers—we focus on cases in which at least 90 days passed between any two consecutive purchases of a buyer. It is *after* imposing these restrictions that we arrive at a sample of 28,640. Even with these restrictions, our results should be interpreted as an exercise in examining how a shift to personalized prices would affect consumer surplus, profits, and overall welfare for this specific population of buyers, holding fixed their behavior outside of this market.

3.2 Identification and Estimation Approach

We believe consumers in our data are not currently facing the first-degree price discrimination we have discussed so far. As discussed above, we interpret our baseline model as describing a *counterfactual* world. Indeed, on the platform from which our data comes from, manufacturers like Apple and Samsung cannot first-degree price discriminate—based on a consumer’s previous positive or negative experiences with the product or based on any other characteristic. Therefore, we treat the transaction and ratings data that we observe as though generated from a dynamic experience goods model akin to the model described above but in which consumers face *uniform* prices rather than personalized prices. We use the empirical model described below to estimate consumer preference and information parameters, and then compare the uniform-pricing world to a counterfactual world with personalized pricing so as to assess the impact of price discrimination on consumer surplus, producer surplus, and total welfare.

The key parameters of the model to estimate are $(X_{fHk}, X_{fLk}, \alpha_{fk}, \beta_{fk})$ for each firm f and variety

on Amazon.com, distinct from the long-standing system of *seller* ratings on eBay. Specifically, ratings of a product are intended to correspond to a buyer’s evaluation of her experience with the product, regardless of which individual eBay seller sold her the product, whereas seller ratings—see, for instance, Nosko and Tadelis 2015—are intended to correspond to a buyer’s evaluation of a particular seller on eBay. There may be a large number of sellers offering any given product, and vice versa. eBay product ratings are only available for certain well-defined, cataloged products as described in the previous paragraph. For example, an eBay listing for an old high school yearbook would have no product ratings, but the seller of this product may have seller ratings. We examined the text corresponding to some of the product ratings in our data and found that some buyers were incorrectly basing their reviews on their opinion of the seller rather than the product, but a majority appeared to correctly understand that the product rating is intended to be based on the buyer’s opinion of the product.

k . We consider two firms, Apple (firm 1) and Samsung (firm 2) each offering two products, a phone or a tablet. Thus, for simplicity, in the empirical application, we consider all Apple phones to be a single product, “Apple phone”, all Apple tablets to be a single product, and similar. Consumers can also choose an outside option rather than one of these four products. To map a consumer’s problem more closely into a standard dynamic discrete choice setting, we change notation here and use the index $j \in J = \{1, 2, 3, 4, 5\}$, rather than the index (f, k) , to denote the outside option ($j = 1$) and the four inside-option products.¹⁵ With this notation, the parameters to estimate are $(X_{Hj}, X_{Lj}, \alpha_j, \beta_j)$.

We now describe the estimation of the $\{X_j\}$ parameters for the four products. We assume consumers have independent tastes for Apple’s products and Samsung’s products. Consumers solve a recursive choice problem in deciding among the five options as a function of their current beliefs, $\pi = (\pi_A, \pi_S)$, about being a good match with the products of either firm. In the model, in period 1, a consumer picks one of the five product options, and then experiences a high (H) or low (L) utility outcome. In the following period, the consumer again chooses between these five options, but in this case the consumer has an updated belief about being a good match for the firm whose product she just experienced.

The consumer’s problem of choosing which product to buy is given by the following Bellman equation:

$$U(\pi, \varepsilon_i) = \max_j \{v_j(\pi) + \varepsilon_{ji}\},$$

where $\varepsilon_i = \{\varepsilon_{ji}\}$ represents an i.i.d. choice-specific idiosyncratic disturbance that we introduce into the empirical model to augment the model considered in the previous section. This shock captures random product-specific price discounts that consumer i may receive on eBay, as well as any idiosyncratic preference shocks consumers may experience for products. We will assume ε_{ji} follows an Extreme Value Type I (EV-I) distribution with scale parameter τ . The choice-specific value function, $v_j(\pi)$, is given by

$$\begin{aligned} v_j(\pi) = & \{(1 - \delta)X_{Hj} + \delta E[U(\Pi_{Hj}(\pi), \varepsilon')]\} \gamma_j(\pi) \\ & + \{(1 - \delta)X_{Lj} + \delta E[U(\Pi_{Lj}(\pi), \varepsilon')]\} [1 - \gamma_j(\pi)] - p_j, \end{aligned}$$

where $\gamma_j(\pi) = \alpha_j \pi_{f(j)} + \beta_j (1 - \pi_{f(j)})$, where $\pi_{f(j)}$ is π_A if j is a product of firm A and is π_B otherwise. Note that $p_{f(j)}$ in the choice-specific value function is the brand-new price for product j . We approximate

¹⁵Note that in our data we also observe phone and tablet purchases that do not correspond to Samsung or Apple products. We consider these purchases to be *outside option* purchases. This is in contrast to some demand settings in the literature, where an outside option purchase is considered to be an instance in which a consumer is not observed purchasing any of the products in the researcher’s data. The fact that we observe these non-Apple and non-Samsung purchases allows us to formalize the non-uniform time period between a consumer’s purchases, as we describe at the end of this subsection.

this price from eBay data using purchases within the first six months that a given type of phone or tablet, for instance, “Apple iPhone 5S 16GB”, appears in the data, and computing the 90th percentile across these purchases for a given product j . We specifically choose the 90th percentile to best approximate the price that a consumer would likely pay for the product if purchased outside of eBay. The price consumers actually face *on any given eBay listing* is not this price, but rather a discounted price that can widely vary across transactions. We treat this discount as entering a consumer’s utility through the term ε_{ji} .¹⁶

Each consumer has an initial prior π_0 , which is a random variable—a vector with a prior for Apple and prior for Samsung—drawn from a density $g(\pi_0)$. Since a consumer’s priors are independent across the two firms, we only need to model a marginal density for the prior for each firm. We use a Beta(η_{f1}, η_{f2}) distribution for each firm f , where η_{f1} and η_{f2} are parameters to be estimated—two parameters for Apple’s products and two for Samsung’s products.

Our estimation approach relies on using pairs of consecutive purchases by a given consumer. For an individual consumer i who purchases product k followed by product j , her contribution to the likelihood function is given by

$$\ell_i = \int_0^1 \left\{ \left(\frac{e^{v_j(\Pi_{Hk}(\pi_0))/\tau}}{\sum_l e^{v_l(\Pi_{Hk}(\pi_0))/\tau}} \right) \gamma_k(\pi_0) + \left(\frac{e^{v_j(\Pi_{Lk}(\pi_0))/\tau}}{\sum_l e^{v_l(\Pi_{Lk}(\pi_0))/\tau}} \right) [1 - \gamma_k(\pi_0)] \right\} \frac{e^{v_k(\pi_0)/\tau}}{\sum_l e^{v_l(\pi_0)/\tau}} g(\pi_0; \eta) d\pi_0. \quad (31)$$

The model allows us to estimate the information parameters α_j and β_j , which appear in the expressions for $\gamma_j(\cdot)$ above and in those for the updated beliefs $\Pi_{Hj}(\cdot)$ and $\Pi_{Lj}(\cdot)$, in an initial step outside of the likelihood function. For this estimation, we rely on our consumer ratings data. This data is more sparse than our transaction data. Because of this feature of the ratings data, we find it helpful to impose the constraint that $\beta_j = 1 - \alpha_j$, which is a common assumption in the theoretical literature. We also assume that $\alpha > 1/2$.

For each product j , we consider a consumer’s outcome to be a success/good outcome if the consumer records a 5-star rating for the corresponding purchase, which occurs for about 79% of ratings, and a failure/bad outcome otherwise. It can be shown that for a consumer who purchases the same product twice in a row with an initial prior belief of $\pi_{f(j)}$ before the first purchase,

$$\Pr(z_t = H, z_{t+1} = L) = \Pr(z_t = L, z_{t+1} = H) = \pi_{f(j)}\alpha_j(1 - \alpha_j) + (1 - \alpha_j)\alpha_j(1 - \pi_{f(j)}) = \alpha_j(1 - \alpha_j),$$

where $\Pr(z_t = H, z_{t+1} = L)$ is the probability the consumer experiences a good outcome after the first

¹⁶An alternative approach would be to use eBay prices directly, but such an approach would require taking a stand on the eBay prices of the goods a consumer did *not* purchase, which is not obvious from the data.

purchase and a bad outcome after the second purchase, and $\Pr(z_t = L, z_{t+1} = H)$ is the analogous probability in the reverse case. These probabilities can be observed in the data. In practice, we find that one or the other of these two probabilities is equal to zero for some products. We therefore average these two probabilities and estimate α_j as the value that satisfies

$$\hat{\Pr}(z_t = H, z_{t+1} = L) + \hat{\Pr}(z_t = L, z_{t+1} = H) = 2\alpha_j(1 - \alpha_j),$$

where the symbol $\hat{(\cdot)}$ represents the empirical frequency of the event of one good outcome and one bad outcome among all cases in which a consumer purchased and rated product j twice in a row.

With estimates of α_j in hand, we then estimate the other parameters of the model by maximum likelihood, with the likelihood function formed by aggregating (31) over all consecutive purchases observed in the data. Estimation of the parameters $\{X_{Hj}, X_{Lj}\}$ requires normalizing those of one product to zero: we set $X_{L1} = X_{H1} = 0$ for the non-Apple/non-Samsung product. The full list of parameters we estimate in this likelihood step is thus $\{X_{Lj}\}_{j \in \{2,3,4,5\}}$, and $\{X_{Hj}\}_{j \in \{2,3,4,5\}}$, the vector of distribution parameters η of the density $g(\pi_0)$, and τ , which is the variance of the EV-I term. Specifically, we estimate these parameters using maximum likelihood with value function constraints, adopting the MPEC approach of Su and Judd (2012)—this approach is an efficient alternative to the standard approach of a nested fixed-point algorithm with a maximum likelihood objective function. For the calculation of the value function in the evaluation of continuation values, we use Chebyshev interpolation, and for the integration required to evaluate the likelihood function, we rely on Gauss-Chebyshev quadrature (following Judd (1998) in both cases).

We now describe our specification of the discount factor. In the data, the time between purchases varies across buyers, so there is no clear “period length” common to all buyers. Therefore, we model each purchasing consumer as facing a Poisson arrival rate λ of the need to purchase a new device. Under this Poisson rate assumption, the expected time between purchases is then an exponentially distributed random variable, which we denote by \mathcal{T} , with mean $1/\lambda$. In the data, we observe a truncated moment of this distribution, $E[\mathcal{T}|\underline{T} \leq \mathcal{T} \leq \bar{T}]$. The corresponding theoretical equivalent of this truncated moment is

$$\frac{e^{-\underline{T}\lambda} - e^{-\bar{T}\lambda}}{e^{-\underline{T}\lambda} - e^{-\bar{T}\lambda}} - \frac{1}{\lambda},$$

where \underline{T} and \bar{T} are the minimum and maximum observed lengths of time, corresponding to 90 and 870 days in the data, respectively. We equate these moments to solve for λ . We assume consumers face an annual rate of time discounting of $\rho = 0.05$. We combine this standard discounting with the random

purchase opportunity Poisson rate to obtain an adjusted discount factor δ that corresponds to the discount factor consumers would face if given *one* purchase opportunity per year. This adjustment yields a discount factor of $\delta \equiv \tilde{\delta}/[1 - (1 - \lambda)\tilde{\delta}]$ with $\tilde{\delta} = 1/(1 + \rho)$.

When simulating counterfactuals to compare the uniform-pricing case to the personalized-pricing case, we use a marginal cost of \$200 for each product, based on marginal cost estimates for both Apple’s products and Samsung’s products from industry tear-down reports.

3.3 Results

The estimated parameters are shown in Table 1. The units of these estimates are dollars per consumer per year, and are relative to the outside option, which yields a normalized payoff of zero. Therefore, all estimates in dollars should be interpreted as an amount above or below the outside option utility. We find that the probability of a good experience if a consumer is a good match for either Apple or Samsung, described by the parameters α_j ’s, is above 0.96 for all products, which implies both that utility experiences are fairly informative about the quality of the match between a consumer and a firm’s brand, and that on average consumers who are a good match for a given brand experience much higher utility than consumers who are a bad match. As α_j is highest for Samsung tablets and lowest for Apple tablets, these products are, respectively, the most and the least informative ones—the ordering of products by their (Blackwell) informativeness is equivalent to their ordering by their values of α_j when $\alpha_j = 1 - \beta_j$.¹⁷ Note that relative to the marginal cost of \$200 of each product, Samsung tablets are also the product with the lowest markup and characterized by the lowest utility realizations after a good or a bad experience.

The estimated prior distributions are shown in Figures 6 and 7. We estimate similar distributions of initial priors for Apple’s products and Samsung’s products, although Apple priors place more weight near 0 and less weight at values close to 1 than Samsung priors. Hence, on average consumers hold low priors that they are a good fit for Apple’s products but high priors that they are a good fit for Samsung’s products. These estimates reflect the large market share of Samsung phones and the low market shares of Apple’s products within the eBay data, despite these latter products yielding higher utility experiences as implied by the estimates of $\{X_{Hj}\}$ and $\{X_{Lj}\}$. A prior close to 0 for Apple’s products can be interpreted as a belief of “I am just not an Apple person” whereas a belief close to 1 signifies the opposite. Casual observation outside of our empirical setting suggests that there exists a significant fraction of loyal Apple enthusiasts.

¹⁷By the Blackwell criterion of informativeness, product j is more informative than product k if the posterior beliefs reached after experiencing product j at the end of a period second-order stochastically dominate those reached after experiencing product k .

We take our estimates of these priors as suggesting that, among active eBay participants, these Apple enthusiasts may be less common.

Using these estimated parameters, we simulate two different worlds. First, we simulate the main model in our paper, where firms can first-degree price discriminate based on consumers' purchase history and experiences with the purchased products. Specifically, we first determine the product that each firm offers a consumer of a given prior and the price each firm charges; we next calculate the product the consumer will buy; and we then compute consumer surplus, producer surplus (variable profits), and total welfare as the sum of consumer and producer surplus. Note that in this counterfactual world, consumers do not face any ε_{ij} shock and do not purchase the outside option in equilibrium, as consistent with our theoretical model. Therefore, in order to generate outcomes in the uniform-price world to which our model can be compared, we simulate consumer choices in the uniform-price world as well, absent price or preference shocks and the ability for consumers to purchase the outside option. For this latter exercise, we hold prices fixed at the prices used in estimation, that is, the 90th percentile of prices for products within the first six months of their release. We then compute consumer surplus, producer surplus, and total welfare for this uniform-price world as well.

Figure 8 displays the estimated consumer surplus for the uniform-price and price-discrimination worlds at each value of the prior for a consumer match with Apple and Samsung. In these figures, as in Table 1, units are relative to the outside option (and thus can be negative). We find that consumer surplus is higher under price discrimination for values of the priors close to $(0, 0)$ and $(1, 1)$ and along the radius of prior values connecting these two extreme points, but lower at priors close to $(0, 1)$ or $(1, 0)$, that is, for consumers who believe they are a poor fit for one firm and a good fit for the other. This is consistent with the intuition from the static model of Thisse and Vives (1988), which demonstrates that in a horizontal price-differentiation world, consumers who are located geographically close to the middle of a Hotelling line with a firm at each end will actually benefit from first-degree price discrimination, because such discrimination will lead to intense competition to serve these consumers. Consumers near the extremes of the Hotelling line, on the other hand, are held captive and may be hurt by first-degree price discrimination.

Figure 9 displays the price paid by consumers at each value of the prior. We find that prices are (mostly) *lower* under price discrimination for values of the prior for Apple's products below 0.5, but *higher* for values of the prior for Apple's products above 0.5 than under uniform-pricing. This result reflects the fact that utility is higher for consumers who are a good match for Apple regardless of their

prior for Samsung. Note that the lowest prices charged in the price-discrimination world are not directly along the 45-degree line in the Apple and Samsung prior space because the firms are not symmetric in that their estimated X_{Lj} 's, X_{Hj} 's, and α_j 's differ. Figure 10 further shows the sum of the estimated profits of Apple and Samsung under the two different scenarios. We find that under price discrimination, profits are lower for most priors except for priors close to $(0, 1)$ or $(1, 0)$, as consistent with the findings in Figures 8 and 9.

Finally, Figure 11 demonstrates that the comparison of price discrimination and uniform pricing in terms of total welfare is qualitatively similar to the comparison of the two pricing schemes in terms of consumer surplus. That is, total welfare is higher for consumers for whom competition is more intense under price discrimination, namely, consumers with priors close to $(0, 0)$ and $(1, 1)$ and along the radius of prior values connecting these two extreme points. As consumers who are a good match for just one brand experience lower utility—and those who are a good match for Apple face much higher prices under price discrimination—welfare at these other priors is lower under price discrimination. Figure 11 shows, however, that, unlike consumer surplus, total welfare increases even at many priors far from $(0, 0)$ and $(1, 1)$ under personalized pricing.

Table 2 reports the welfare measures discussed averaged over the values of the priors for Apple and Samsung' products, when moving from a world of uniform pricing to one of personalized pricing. Namely, the statistics in this table are obtained by integrating the differences in consumer surplus, prices paid, firm profits, total profits, and total welfare across the uniform and price discrimination worlds from Figures 8–11 over all possible priors with respect to the density of π_0 . We find that, on average, consumer surplus would increase by about \$108 per year per consumer if firms could offer personalized prices. Apple's profits per consumer would decrease by about \$83 per year, Samsung's profits would decrease by about \$25, and total welfare would be nearly unaffected, with an increase of only \$0.12 per consumer per year.

These results thus underscore the importance of accounting for the effect on competition of the greater opportunities for price discrimination that “Big Data” may make increasingly feasible when evaluating the impact of “Big Data”. Provided access to information about consumers' behavior and experiences is not too limited or asymmetric across firms, our empirical analysis suggests that there is scope for a significant fraction of consumers to benefit from personalized pricing.

4 Conclusion

Shiller et al. (2014) predicts that over the next decade or so personalized pricing will become increasingly

common, especially for goods sold online. Many of these goods, ranging from clothing and electronics to furniture and food, are differentiated varieties of experience goods. The existing literature is largely silent on both the pricing and efficiency properties of equilibrium in these markets in such circumstances. In this paper, we have proposed a simple model to shed light on these issues. In terms of pricing, we have argued that in stark contrast to the related literature on duopoly markets in which firms price compete for just one consumer, prices no longer have the simple static form familiar from Bertrand competition. Rather, prices contain a variety-specific dynamic component that reflects the relative informativeness of competing varieties. In terms of efficiency, we have clarified the limits to which the efficiency results in the duopoly case with one-variety firms can be extended, and we have provided simple, intuitive examples of the type of inefficiencies that may arise in more general environments with multiple varieties and firms. Finally, we have made precise the sense in which intensifying competition between firms can restore efficiency and give rise to static, competitive pricing.

We have allowed firms to personalize their prices to consumers based on consumers' purchases and experiences. In the context of labor markets, the idea that in many professions, such as professional sports and academia, wages are personalized in that they are highly dependent on a worker's current and past performance is well-accepted. In the context of product markets, the advances in information technology that allow sellers to offer personalized prices, especially in on-line markets, are recent and fast spreading (see Fudenberg and Villas-Boas (2007, 2012)). In this sense, we think of our analysis as relevant for a growing segment of the product market.

Crucially, the inefficiencies we highlight emerge not because of the direct information spillovers identified by the literature on strategic experimentation in nonmarket settings (as in Bolton and Harris (1999)). Rather, inefficiencies arise here because in Markov models of Bertrand competition, there is no means for groups of losing firms to discipline the behavior of the winning firm, either with sticks (coordinated punishments) or carrots (transfers to other firms or the consumer) in order to induce an alternative choice of variety by the winning firm or to induce an alternative firm to win.

Lastly, using eBay data, we have provided evidence on the gains associated with the sophisticated forms of price discrimination considered here for the market for smartphones and tablets. We find that a large fraction of consumers benefit from the introduction of price discrimination—in particular those with less asymmetric beliefs about their tastes for different brands for whom competition intensifies. Consumers more certain about their tastes for one firm, however, can be worse off under price discrimination.

References

- Angeletos, G.-M. and Pavan, A. (2007). Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142.
- Athey, S. and Mobius, M. (2012). The impact of news aggregators on internet news consumption: The case of localization. In *Workshop on the Economics of Web Search and Social Network. Sixth ACM International Conference on Web Search and Data Mining 2013*, page 2.
- Atkeson, A., Hellwig, C., and Ordoñez, G. (2014). Optimal regulation in the presence of reputation concerns. *The Quarterly Journal of Economics*, 130(1):415–464.
- Banks, J. S. and Sundaram, R. K. (1992). Denumerable-armed bandits. *Econometrica: Journal of the Econometric Society*, pages 1071–1096.
- Bergemann, D. and Välimäki, J. (1996). Learning and strategic pricing. *Econometrica: Journal of the Econometric Society*, pages 1125–1149.
- Bergemann, D. and Välimäki, J. (2000). Experimentation in markets. *The Review of Economic Studies*, 67(2):213–234.
- Bergemann, D. and Välimäki, J. (2006a). Dynamic price competition. *Journal of Economic Theory*, 127(1):232–263.
- Bergemann, D. and Välimäki, J. (2006b). Dynamic pricing of new experience goods. *Journal of Political Economy*, 114(4):713–743.
- Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.
- Bonatti, A. (2011). Menu pricing and learning. *American Economic Journal: Microeconomics*, 3(3):124–63.
- Calzolari, G. and Pavan, A. (2006). On the optimality of privacy in sequential contracting. *Journal of Economic theory*, 130(1):168–204.
- Eeckhout, J. and Weng, X. (2015). Common value experimentation. *Journal of Economic Theory*, 160:317–339.
- Erdem, T. (1998). An empirical analysis of umbrella branding. *Journal of Marketing Research*, pages 339–351.
- Erdem, T. and Chang, S. R. (2012). A cross-category and cross-country analysis of umbrella branding for national and store brands. *Journal of the Academy of Marketing Science*, 40(1):86–101.
- Fudenberg, D. and Villas-Boas, J. M. (2007). *Behavior-Based Price Discrimination and Customer Recognition*. Elsevier Science, Oxford.
- Fudenberg, D. and Villas-Boas, J. M. (2012). *Price Discrimination in the Digital Economy*. Oxford University Press.
- Gallivan, Q. (2012). Why big data is all retailers want for christmas.

- Heilman, C. M., Bowman, D., and Wright, G. P. (2000). The evolution of brand preferences and choice behaviors of consumers new to a market. *Journal of Marketing Research*, 37(2):139–155.
- Hellwig, C. and Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *The Review of Economic Studies*, 76(1):223–251.
- Judd, K. L. (1998). *Numerical methods in economics*. MIT press.
- Keller, G. and Rady, S. (2010). Strategic experimentation with poisson bandits. *Theoretical Economics*, 5(2):275–311.
- Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.
- Klein, N. and Rady, S. (2011). Negatively correlated bandits. *The Review of Economic Studies*, 78(2):693–732.
- Mitchell, M. F. (2000). The scope and organization of production: firm dynamics over the learning curve. *The Rand journal of economics*, pages 180–205.
- Nosko, C. and Tadelis, S. (2015). The limits of reputation in platform markets: An empirical analysis and field experiment. Technical report, National Bureau of Economic Research.
- Oh, W. and Lucas Jr, H. C. (2006). Information technology and pricing decisions: Price adjustments in online computer markets. *MIS quarterly*, pages 755–775.
- Shapiro, C. and Varian, H. R. (1999). *Information rules: a strategic guide to the network economy*. Harvard Business School Press.
- Shiller, B. and Waldfogel, J. (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. et al. (2014). First-degree price discrimination using big data. *mimeo*.
- Su, C.-L. and Judd, K. L. (2012). Constrained optimization approaches to estimation of structural models. *Econometrica*, 80(5):2213–2230.
- Thisse, J.-F. and Vives, X. (1988). On the strategic choice of spatial price policy. *The American Economic Review*, pages 122–137.
- Villas-Boas, S. B. and Villas-Boas, J. M. (2008). Learning, forgetting, and sales. *Management Science*, 54(11):1951–1960.

A Proofs

Proof that Consumer Trades at any Prior: To prove this result, suppose by way of contradiction that there exists an equilibrium in which the consumer rejects both offers at some prior. If this occurs, then the consumer receives zero utility in the current period and the prior is unchanged. By our Markov assumption, this result implies that the consumer must decline both firms' offers in all future periods as well and, hence, end up with an expected present discounted value of zero. For this choice to possibly be optimal, the consumer must be offered a value of zero or less from a purchase from either firm. But if one firm offers the consumer such a trade, it is optimal for the other firm to offer a price $p_f \in (0, x_f(\pi, \tilde{k}_f))$ for variety $\tilde{k}_f \in \arg \min_k x_f(\pi, k)$ and attract the consumer, since even at the lowest prior and for the least profitable variety, gross utility is positive if $X_{fHk}, X_{fLk} > 0$ for all f and k . Hence, $x_f(\pi, \tilde{k}_f) > p_f$. Thus, such an equilibrium cannot exist and the consumer must purchase a variety at each prior. \square

Proof of Proposition 1: As mentioned, the set of priors can be partitioned into priors in E_A , at which firm A sells, and priors in E_B , at which firm B sells. At $\pi \in E_A$, firm A weakly prefers selling to the consumer to not selling and having firm B sell variety $k_B(\pi)$. Hence, firm A 's optimality implies

$$\max_{(k_A, p_A) \in F_A} [(1 - \delta)p_A + \delta EV^A(\pi' | \pi, k_A)] \geq \delta EV^A(\pi' | \pi, k_B(\pi)). \quad (32)$$

Moreover, at equilibrium, the optimality of firm A 's price implies that the constraint (13) holds as an equality so that

$$(1 - \delta)u_f(\pi_A, k_A, p_A) + \delta EU(\pi' | \pi, k_A) = (1 - \delta)u_f(\pi_B, k_B(\pi), p_B(\pi)) + \delta EU(\pi' | \pi, k_B(\pi)). \quad (33)$$

Using the definition of u_f in (2), substituting p_A obtained from (33) into (32), and using the definition of $W^A(\pi)$ implies that we can rewrite (32) as

$$\max_{k_A \in K_A} [(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi' | \pi, k_A)] \geq (1 - \delta)[x_B(\pi_B, k_B(\pi)) - p_B(\pi)] + \delta EW^A(\pi' | \pi, k_B(\pi)).$$

At $\pi \in E_B$, firm B is selling and, by the definition of equilibrium, firm A is indifferent between selling and not selling so that

$$\max_{(k_A, p_A) \in F_A} [(1 - \delta)p_A + \delta EV^A(\pi' | \pi, k_A)] = \delta EV^A(\pi' | \pi, k_B(\pi)). \quad (34)$$

At such a prior, consumer optimality implies that the consumer prefers purchasing from the selling firm to purchasing from the firm's competitor so that

$$(1 - \delta)u_f(\pi_A, k_A, p_A) + \delta EU(\pi' | \pi, k_A) \leq (1 - \delta)u_f(\pi_B, k_B(\pi), p_B(\pi)) + \delta EU(\pi' | \pi, k_B(\pi)). \quad (35)$$

Using the definition of u_f , substituting p_A obtained from (35) into (34), and using the definition of $W^A(\pi)$, we obtain

$$\max_{k_A \in K_A} [(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi' | \pi, k_A)] \leq (1 - \delta)[x_B(\pi_B, k_B(\pi)) - p_B(\pi)] + \delta EW^A(\pi' | \pi, k_B(\pi)).$$

Combining the two cases establishes the proposition. \square

Proof of Corollary 1: The proof amounts to using the cautious restriction to show that the match value

expression can be rewritten as

$$W^A(\pi) = \begin{cases} \max_{k_A \in K_A} \{(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)\} & \text{for } \pi \in E^A \\ \max_{k_A \in K_A} \{(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)\} & \text{for } \pi \in E^B \end{cases}, \quad (36)$$

so $W^A(\pi) = \hat{W}^A(\pi_A)$. Clearly, for any $\pi \in E^A$, the value has the form in the top branch. Now, to establish that the value has the form in the bottom branch, we proceed as follows. Formally, consider the problem faced by, say, firm A , summarized in (12) and (13). Clearly, firm A solves this problem at any prior π at which it sells to the consumer. The dual form of the cautious restriction implies that firm A also solves such a problem at any prior π at which it does not sell to the consumer. As we have discussed, in any solution the constraint (13) holds as an equality, otherwise firm A could increase its profits by marginally increasing the price charged to the consumer and still attracting the consumer. Hence, we can use (14) to obtain an expression for p_A by using the definition of u_f in (2). Substitute the resulting expression for p_A back into (12) to rewrite firm A 's problem as

$$V^A(\pi) = \max_{k_A \in K_A} [(1 - \delta)x_A(\pi_A, k_A) + \delta EU(\pi'|\pi, k_A) + \delta EV^A(\pi'|\pi, k_A) - C^A(\pi)], \quad (37)$$

where $C^A(\pi) = (1 - \delta)u_f(\pi_B, k_B(\pi), p_B(\pi)) + \delta EU(\pi'|\pi, k_B(\pi))$ is simply an additive constant that is not affected by the actions (k_A, p_A) . Denote by $U(\pi|A)$ the consumer's equilibrium value from purchasing from firm A and by $U(\pi|B)$ the consumer's equilibrium value from purchasing from firm B . Recalling that $U(\pi|A) = U(\pi|B)$, consider the following simple steps: add the equilibrium value $U(\pi|A)$ to the left side of (37) and the equilibrium value $U(\pi|B)$ to the right side of (37), shift the term $C^A(\pi)$ outside of the maximization operator, simplify the right side of the resulting expression using the fact that $U(\pi|B) = C^A(\pi)$, and, finally, substitute $W^A(\pi) = V^A(\pi) + U(\pi)$ for the sum of the values of firm A 's profits and the consumer's utility. We then arrive at

$$W^A(\pi) = \max_{k_A \in K_A} \{(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)\}. \quad (38)$$

Since $x_A(\pi_A, k_A)$ does not depend on π_B and only the first component of $\pi = (\pi_A, \pi_B)$ is updated when a variety of firm A is consumed, it follows that the programs in (18) and (38) have the same value. \square

Proof of Proposition 2: By way of contradiction, suppose that at the initial prior π , the sum of the values of firms' profits and the consumer's utility under the solution to the conditional planning problem $W^*(\pi)$ is strictly greater than the sum of the values of firms' profits and the consumer's utility in the equilibrium. By the one-shot deviation principle for dynamic programming, for this to be the case there must exist a one-shot deviation from the equilibrium plan of firm choice that leads to a higher sum of values.

In particular, suppose without loss that at prior π , the planner prefers that firm B sells to the consumer but that in the equilibrium firm A is the selling firm. We introduce some notation to denote the values of firms' profits and the consumer's utility. To this end, given the choice of variety functions $\hat{k}_A(\pi_A)$ or $\hat{k}_B(\pi_B)$, let

$$V^A(\pi|A) = (1 - \delta)p_A(\pi) + \delta EV^A(\pi'|\pi, \hat{k}_A(\pi_A)) \quad (39)$$

denote the value of firm A 's profits when it sells $\hat{k}_A(\pi_A)$ at prior π , and let

$$V^A(\pi|B) = \delta EV^A(\pi'|\pi, \hat{k}_B(\pi_B)) \quad (40)$$

denote the value of firm A 's profits when firm B sells $\hat{k}_B(\pi_B)$ at prior π . We use analogous notation for the value of firm B 's profits in these two cases. We let $U(\pi|A)$ and $U(\pi|B)$ denote the corresponding values for the consumer. Our contradiction hypothesis is that at π , the planner prefers that B is chosen by

the consumer over A , that is,

$$V^A(\pi|B) + V^B(\pi|B) + U(\pi|B) > V^A(\pi|A) + V^B(\pi|A) + U(\pi|A), \quad (41)$$

but that firm A sells in equilibrium, which implies that $V^A(\pi|A) \geq V^A(\pi|B)$.

Recall that (14) must hold in equilibrium, so that $U(\pi|A) = U(\pi|B)$. Since in equilibrium firm B is not selling, the value of its profits from losing must equal the value of the profits it would have obtained by winning. Thus, $V^B(\pi|A) = V^B(\pi|B)$. Using these two equalities, $U(\pi|A) = U(\pi|B)$ and $V^B(\pi|A) = V^B(\pi|B)$, to simplify (41), we obtain that

$$V^A(\pi|B) > V^A(\pi|A), \quad (42)$$

which implies that firm A can raise its profits by increasing its price marginally and losing the consumer to firm B . But this contradicts profit maximization by firm A and thus establishes the desired result. \square

Proof of Corollary 3: The proof is immediate. Under (20), the continuation match value of any firm depends on the identity of the selling firm but no longer depends on the particular variety that such a firm offers. That is, the value $E\hat{W}^f(\pi'_f|\pi_f, k_f)$ does not depend on k_f . In a slight abuse of notation, we rewrite this value as simply $E\hat{W}^f(\pi'_f|\pi_f, f)$, meaning that the future state does not depend on the particular variety that firm f offers. We similarly denote $EW^*(\pi'|\pi, k_f)$ by $EW^*(\pi'|\pi, f)$. In such a case, the choice of variety of firm f is static, namely, firm f selects the variety that yields the highest gross utility in the current period. Thus, under (20) we can rewrite (18) as

$$\hat{W}^f(\pi_f) = \max_{k_f \in K_f} \left\{ (1 - \delta)x_f(\pi_f, k_f) + \delta E\hat{W}^f(\pi'_f|\pi_f, f) \right\}.$$

So, the equilibrium variety offered by firm f , $\hat{k}_f(\pi_f)$, solves the static problem $\max_{k_f \in K_f} x_f(\pi_f, k_f)$. Since the future state depends on the identity of the selling firm but not on its choice of variety, conditional efficiency implies

$$W^*(\pi) = \max_f \left\{ (1 - \delta) \max_{k_f \in K_f} x_f(\pi_f, k_f) + \delta EW^*(\pi'|\pi, f) \right\}. \quad (43)$$

Once again using the fact that the future state does not depend on the variety choice of any firm, $W^*(\pi)$ in (43) can be equivalently rewritten as

$$W^*(\pi) = \max_f \left\{ \max_{k_f \in K_f} [(1 - \delta)x_f(\pi_f, k_f) + \delta EW^*(\pi'|\pi, f)] \right\}.$$

Thus, $W^* = W$ and equilibrium allocations are efficient. \square

The Dependently Distributed Stochastic Process of the Examples: The original stochastic process for utility realizations for any variety k of a firm f , given the state of the world $\theta \in \{\bar{\theta}, \underline{\theta}\}$, is $\Pr(X_{fkt} = X_{fHk}|\theta)$ and $\Pr(X_{fkt} = X_{fLk}|\theta)$. This process for each variety and firm is i.i.d., depends on the state of the world, but does not depend on the history of realized utility. Of course, the realizations of the process for some particular X_{fkt} will be relevant at t only if the consumer chooses that particular variety k of firm f at t . We now construct a stochastic process that coincides with the original process in period 1 but after period 1 is dependently distributed in that the process from period 2 on, for any given firm or variety, depends only on the realization of utility in period 1. In this sense, according to the new stochastic process, utility realizations become uninformative about the state of the world after period 1. Also, the

beginning of period 2 priors govern the realizations of utilities from period 2 onward.

As mentioned, we think of this economy as somewhat loosely corresponding to a stochastic analog of the economy that BV2 consider in Section 6 of their paper. These authors consider an infinite horizon oligopoly model with deterministic time-varying payoffs that are thought of as capturing either habit formation or learning-by-doing. The economy in Section 6 of their paper assumes that payoffs for a product become constant after a certain number of uses. Note that BV2 find a unique cautious Markov perfect equilibrium that is efficient, whereas we find a cautious Markov-perfect equilibrium that is inefficient. In this sense, our results clarify that the deterministic results by BV do not extend to a multivariety stochastic economy with arbitrary dependently distributed payoffs.

Formally, let $a_t = (f_t, k_t, z_t)$ denote the *realized experience* of the consumer in period t that records the selling firm f_t at t , the variety sold k_t , and the realized outcome $z_t \in \{H, L\}$ with that variety. Equivalently, a_t records that in period t , the consumer experienced some particular X_{fHk} or X_{fLk} for variety $k = k_t$ of firm $f = f_t$. Let $a^t = (a_1, a_2, \dots, a_{t-1})$ denote the history of such experiences. As noted, in period 1 the process is informative and coincides with the original one. Starting from a prior vector $\pi_1 = (\pi_{A1}, \pi_{B1})$ suppose, for concreteness, that in period 1, the consumer experiences variety $k = k_1$ of firm $f_1 = A$ and has outcome $z_1 = H$, so that $a^2 = (A, k, H)$. Then, in period 2 the updated prior is $\pi_2 = (\pi_{A2}, \pi_{B2}) = (\Pi_{AHk}(\pi_{A1}), \pi_{B1})$. The probability distribution over possible utility levels for all the varieties that could be chosen in period 2 is

$$\Pr(X_{Ak2} = X_{AHk}|a^2) = \alpha_{Ak}\pi_{A2} + \beta_{Ak}(1 - \pi_{A2}) \text{ and } \Pr(X_{Bk2} = X_{BHk}|a^2) = \alpha_{Bk}\pi_{B2} + \beta_{Bk}(1 - \pi_{B2}).$$

Notice that this probability distribution does not depend on the state of the world θ but rather depends only on the realization of period 1 utility. Under this new process, the consumer's gross utility in period 2 is the same as in our original formulation, except now the signals in period 2 are uninformative.

In periods $t \geq 2$, the same distributions apply regardless of the consumer's realized experiences after period 1; that is, in periods $t \geq 2$, $\Pr(X_{Akt} = X_{AHk}|a^t) = \Pr(X_{Ak2} = X_{AHk}|a^2)$ and $\Pr(X_{Bkt} = X_{BHk}|a^t) = \Pr(X_{Bk2} = X_{BHk}|a^2)$.

In short, we can interpret our examples as corresponding to infinite horizon economies with dependently distributed stochastic processes. \square

Proof of Proposition 3: As argued, it is optimal for the winning firm, here firm A , to charge a price so that (14) holds, which using the definition of u_f in (2) can be written as

$$x_A(\pi_A, k_A(\pi)) - p_A(\pi) + \frac{\delta EU(\pi'|\pi, k_A(\pi))}{1 - \delta} = x_B(\pi_B, k_B(\pi)) - p_B(\pi) + \frac{\delta EU(\pi'|\pi, k_B(\pi))}{1 - \delta}. \quad (44)$$

Next, since the losing firm must be indifferent between not selling to the consumer and selling, we have that

$$\delta EV^B(\pi'|\pi, k_A(\pi)) = (1 - \delta)p_B(\pi) + \delta EV^B(\pi'|\pi, k_B(\pi)). \quad (45)$$

Solving (45) for $p_B(\pi)$ and substituting the resulting expression into (44) gives (21). \square

Proof of Proposition 4: Divide the set of priors into E_A , at which firm A sells, and its complement, E'_A , at which it does not. At $\pi \in E_A$, firm A weakly prefers selling to the consumer to not selling and having the second-best firm $f = f_2(\pi)$ sell variety $k_f(\pi)$. Hence, firm A 's optimality implies

$$\max_{(k_A, p_A) \in F_A} [(1 - \delta)p_A + \delta EV^A(\pi'|\pi, k_A)] \geq \delta EV^A(\pi'|\pi, k_f(\pi)). \quad (46)$$

Moreover, at equilibrium, the optimality of firm A 's offer implies that the many-firm version of constraint

(13) holds as an equality so that

$$(1 - \delta)u_f(\pi_A, k_A, p_A) + \delta EU(\pi'|\pi, k_A) = (1 - \delta)u_f(\pi_f, k_f(\pi), p_f(\pi)) + \delta EU(\pi'|\pi, k_f(\pi)), \quad (47)$$

where $f = f_2(\pi)$. Solving the equality in (47) for p_A , substituting the resulting expression into (46), and using the definition of $W^A(\pi)$ implies that

$$\max_{k_A \in K_A} [(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)] \geq (1 - \delta)[x_f(\pi_f, k_f(\pi)) - p_f(\pi)] + \delta EW^A(\pi'|\pi, k_f(\pi)).$$

At $\pi \in E'_A$, firm A is not selling and, by the definition of equilibrium, firm A is indifferent between selling and not selling so that

$$\max_{(k_A, p_A) \in F_A} [(1 - \delta)p_A + \delta EV^A(\pi'|\pi, k_A)] = \delta EV^A(\pi'|\pi, k_f(\pi)), \quad (48)$$

where $f = f_1(\pi)$. At such a prior, by the definition of equilibrium, the consumer at least weakly prefers the selling firm so that

$$(1 - \delta)u_f(\pi_A, k_A, p_A) + \delta EU(\pi'|\pi, k_A) \leq (1 - \delta)u_f(\pi_f, k_f(\pi), p_f(\pi)) + \delta EU(\pi'|\pi, k_f(\pi)), \quad (49)$$

with $f = f_1(\pi)$. We can equivalently express (49) as

$$(1 - \delta)x_A(\pi_A, k_A) + \delta EU(\pi'|\pi, k_A) - (1 - \delta)[x_f(\pi_f, k_f(\pi)) - p_f(\pi)] - \delta EU(\pi'|\pi, k_f(\pi)) \leq (1 - \delta)p_A(\pi),$$

which, together with (48), implies that

$$\begin{aligned} \max_{k_A \in K_A} [(1 - \delta)x_A(\pi_A, k_A) + \delta EU(\pi'|\pi, k_A) - (1 - \delta)[x_f(\pi_f, k_f(\pi)) - p_f(\pi)] - \delta EU(\pi'|\pi, k_f(\pi)) \\ + \delta EV^A(\pi'|\pi, k_A)] \leq \delta EV^A(\pi'|\pi, k_f(\pi)). \end{aligned} \quad (50)$$

Using the definition of $W^A(\pi)$ and noting that $x_f(\pi_f, k_f(\pi))$, $p_f(\pi)$, and $\delta EU(\pi'|\pi, k_f(\pi))$ do not depend on (k_A, p_A) , we can simplify (50) to obtain

$$\max_{k_A \in K_A} [(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)] \leq (1 - \delta)[x_f(\pi_f, k_f(\pi)) - p_f(\pi)] + \delta EW^A(\pi'|\pi, k_f(\pi)).$$

Combining the two cases establishes the proposition. \square

Proof of Corollary 6: Recall the pricing rule in (21). Here the analog of Corollary 1 applies so that $W^f(\pi)$ coincides with the autarky value $\hat{W}^f(\pi)$, which is convex in π . Since a mean-preserving spread of a convex function increases its value, it is immediate that if the variety of the selling firm, here $k_A(\pi)$, is more informative than that of the nonselling firm, here $k_B(\pi)$, then the compensating price differential is positive. Likewise, if $k_B(\pi)$ is more informative than $k_A(\pi)$, then the compensating price differential is negative. \square

Proof of Proposition 6: We solve for the equilibrium by backward induction.

Second Period. Note, as before, that in period 2 the firm with the highest value, $x_f(\pi_{f2})$, sells to the consumer, where π_{f2} is the beginning-of-period prior. Hence, given the winning firm in the first period and this prior, condition (62) determines the second-period outcomes.

First, suppose that A won in the first period. If a success occurred, then $\pi_{A2} = 1$, A wins in the second period, B is the second-best firm, the price charged by A is $p_{A2} = x_A(1) - x_B(\pi_B)$, and the consumer's utility is $x_B(\pi_B)$. If a failure occurred, then $\pi_{A2} = 0$, B wins in the second period, C is the second-best

firm, the price charged by B is $p_{B2} = x_B(\pi_B) - x_C$, and the consumer's utility is x_C .

Next, suppose that B won in the first period. If a success occurred, then $\pi_{B2} = 1$, B wins in the second period, C is the second-best firm, the price charged by B is $p_{B2} = x_B(1) - x_C$, and the utility of the consumer is x_C . If a failure occurred, then $\pi_{B2} = 0$, C wins in the second period, A is the second-best firm, the price charged by C is $p_{C2} = x_C - x_A(\pi_A)$, and the utility of the consumer is $x_A(\pi_A)$.

Finally, suppose that C won in the first period. Since its variety is uninformative, the prior does not change. Hence, B wins in the second period regardless of the utility realized in the first period. Clearly, C is the second-best firm, the price charged by B is $p_{B2} = x_B(\pi_B) - x_C$, and the utility of the consumer is x_C .

First Period. Consider the behavior of firm C in the first period. In equilibrium, firm C never sells and hence earns a profit of zero. Suppose that this firm deviates, charges a price p_{C1} , and attracts the consumer. The present value of profits from this deviation is $(1 - \delta)p_{C1}$, since under this deviation in the second period firm B sells to the consumer after both a high and a low utility is realized in the first period. Hence, to make firm C indifferent between selling and not selling, it must be that $p_{C1} = 0$. Thus, the consumer's utility from purchasing from firm C in the first period is $U(\pi|C) = (1 - \delta)x_C + \delta x_C = x_C$.

Next, consider the behavior of firm B in period 1. In equilibrium, firm B sells only after a failure by firm A in period 1. Since in this case firm C is the second-best firm in the second period, firm B 's value of profits in equilibrium is $\delta(1 - \pi_A)[x_B(\pi_B) - x_C]$. Now this value must equal the value that firm B would obtain if it deviated in period 1 and attracted the consumer at price p_{B1} , which from our second-period analysis would imply a value of profits of $(1 - \delta)p_{B1} + \delta\pi_B[x_B(1) - x_C]$ and a value of the consumer's utility of

$$U(\pi|B) = (1 - \delta)[x_B(\pi_B) - p_{B1}] + \delta[\pi_B x_C + (1 - \pi_B)x_A(\pi_A)]. \quad (51)$$

The price that makes firm B indifferent between deviating or not satisfies

$$(1 - \delta)p_{B1} = \delta(1 - \pi_A)[x_B(\pi_B) - x_C] - \delta\pi_B[x_B(1) - x_C]. \quad (52)$$

Substituting p_{B1} from (52) into (51), we obtain that the value of the consumer's utility from choosing firm B is

$$U(\pi|B) = (1 - \delta)x_B(\pi_B) + \delta[\pi_B x_B(1) + (1 - \pi_B)x_A(\pi_A)] - \delta(1 - \pi_A)[x_B(\pi_B) - x_C].$$

We will assume that $U(\pi|B) \geq U(\pi|C)$, so that firm B is the second-best firm.

Finally, consider the behavior of firm A in period 1. The value of firm A 's profits from selling must be larger than the value of firm A 's profits when firm A raises its price so much that it loses the consumer to the second-best firm, here firm B . Hence, it must be that

$$(1 - \delta)p_{A1} + \delta\pi_A[x_A(1) - x_B(\pi_B)] \geq 0, \quad (53)$$

where we have used our second-period analysis to determine that if firm B , the second-best firm, attracts the consumer in the first period, then firm A 's profits in the second period are zero. In the first period, firm A will charge a price p_{A1} that makes the consumer indifferent between accepting its offer and accepting firm B 's offer. The value of the consumer's utility from purchasing from firm A in the first period is

$$U(\pi|A) = (1 - \delta)[x_A(\pi_A) - p_{A1}] + \delta[\pi_A x_B(\pi_B) + (1 - \pi_A)x_C],$$

whereas the consumer's utility from purchasing from firm B is $U(\pi|B)$. To make the consumer indifferent

between firms A and B in the first period, p_{A1} must be such that $U(\pi|A) = U(\pi|B)$, so

$$p_{A1} = x_A(\pi_A) - x_B(\pi_B) + \frac{\delta}{1 - \delta} [x_B(\pi_B) - \pi_B x_B(1) - (1 - \pi_B)x_A(\pi_A)]. \quad (54)$$

The term in brackets in (54) is the compensating price differential as (57) requires, $EW^B(\pi'|\pi, k_A(\pi)) - EW^B(\pi'|\pi, k_B(\pi))$, which here is negative since $x_A(\pi_A) > x_B(0)$.

Lastly, by substituting the expression in (54) into (53), we can combine the condition that firm A obtains a nonnegative value of profits in the first period and the condition $U(\pi|B) \geq U(\pi|C) = x_C$ into a single one, namely,

$$(1 - \delta)x_A(\pi_A) + \delta[\pi_A x_A(1) + (1 - \pi_A)x_C] \geq U^B \geq x_C. \quad (55)$$

Conditions (62) and (55) are sufficient for our equilibrium to exist. This equilibrium is inefficient when (63) is greater than (64). These conditions can all be satisfied, so that the equilibrium exists and is inefficient for both small and large δ .¹⁸

We can relate our example to our earlier intuition for inefficiency by noting that in the example $V^C(\pi|A) = 0$ and $V^A(\pi|B) = 0$, so that (61) is equivalent to $V^C(\pi|B) > V^A(\pi|A)$. It is easy to verify that this condition holds here. \square

Proof of Proposition 8: Suppose the horizon T is finite. Also suppose that in the last period, the equilibrium dictates that some variety k with prior π_{kT} is sold by two firms, say, A and B . Since both of these firms have identical technologies for producing this variety, Bertrand competition implies that these firms price at marginal cost, here zero. Now consider period $T - 1$. Since in period T profits will be zero both on and off the equilibrium path for all varieties, the game from period $T - 1$ on becomes static and Bertrand competition between the two (or more) firms selling the equilibrium variety drives prices to their marginal costs. By backward induction, this logic applies to every period. Hence, firms always charge prices equal to their marginal costs. To establish efficiency, note that since firms price at marginal cost, all firms' profits are zero. Thus, the match value of each firm is the value of the consumer's gross utility. Therefore, each firm chooses a variety that maximizes the objective function of the planner. Since, in turn, the planner's objective coincides with that of the consumer, the optimality of the consumer's choice implies that the equilibrium allocation coincides with the planning solution.

Suppose now that the horizon is infinite. The limit of the finite horizon strategies of the consumers and the firms is an equilibrium of the infinite horizon economy. In that equilibrium, as just proved, firms charge prices equal to their marginal costs and the equilibrium is efficient. \square

B Correlated Tastes

So far we have assumed that taste is independent across brands. We now discuss an alternative extreme in which the consumer's taste is *perfectly correlated* across all brands, so that the consumer has either a good or bad match with all firms. One example is the market for fashion clothing. Suppose that for some brands with a more distinctive style, say Gucci, the consumer's gross utility is very steep in the prior: consumers who are a good match for fashion clothing enjoy Gucci products very much whereas consumers who are a bad match enjoy them very little. For other brands with a less distinctive style, say Ralph Lauren, the

¹⁸Parameter values producing these outcomes are: $x_A(1) = 11, x_B(1) = 9, x_C = 5, x_A(0) = 2, x_B(0) = 1, \alpha_{A1} = \alpha_{C1} = 1, \beta_{A1} = \beta_{C1} = 0, \delta = 0.95, \pi_A = 0.3$, and $\pi_C = 0.6$. Alternatively, $x_A(1) = 13.35, x_B(1) = 8.34, x_C = 5, x_A(0) = 0.57, x_B(0) = 0.01, \alpha_{A1} = \alpha_{C1} = 1, \beta_{A1} = \beta_{C1} = 0, \delta = 0.45, \pi_A = 0.3$, and $\pi_C = 0.6$. It is also easy to find parameter values such that firm B is the selling firm and firm C is the second-best firm but the planner prefers firm A to all other firms.

consumer's gross utility is relatively flat in the prior: consumers who are a good match for fashion clothing enjoy Ralph Lauren products only slightly more than consumers who are a bad match.

We formalize these ideas with a simple modification of our setup. Here, the state of the world is a scalar $\theta \in \{\bar{\theta}, \underline{\theta}\}$, where $\bar{\theta}$ means that the consumer is a good match for all firms in the market and $\underline{\theta}$ means that the consumer is a bad match for all firms, albeit to different degrees as captured by the possible utility realizations X_{fHk} and X_{fLk} , and probabilities of success for each firm f and variety k , α_{fk} and β_{fk} . The prior at t is now a scalar, π_t , which denotes the probability that the match between the consumer and all brands is good. The rest of the environment and the definition of equilibrium are the same as in the earlier setup. In this environment except for Corollary 4, all the earlier propositions and corollaries immediately apply. Next, we formalize the modification of Corollary 4 that applies to this case.

Corollary 6. (*Compensating Price Differential with Correlated Tastes*) *In a duopoly with correlated consumer tastes, the compensating price differential is negative when the selling firm offers a less informative variety and is positive when the selling firm offers a more informative variety than the nonselling firm.*

Thus, when its variety is less informative than that of its competitor, the firm prices its variety *below* the statically optimal price and when its variety is more informative than that of its competitor, the firm prices its variety *above* the statically optimal price. Hence, the selling firm is bound by competition to offer a discount for a variety that is at a competitive disadvantage in terms of the information it conveys, but it can charge a premium for a variety that allows the acquisition of more information than the best alternative in the market. Moreover, as priors change and the second-best variety changes, a firm can switch from pricing below the static price to above the static price and vice-versa.

Imagine now an unwary observer watching the pricing behavior of a firm when it introduces a product. In the first case above, that observer would see prices systematically below the static price, and might be tempted to interpret the pricing behavior as an type of penetration pricing. Likewise, in the second case above, that observer would see prices systematically above the static price and might be tempted to interpret the pricing behavior as a type of price skimming. Unbeknownst to this observer, however, is that both types of pricing behavior are simply the result of the firm charging a compensating differential for its product.

Moreover, if as priors change, the pattern of products offered by firms changes then the model predicts what looks to outside observers a pattern of temporary and seemingly random price increases and decreases. Hence, our equilibrium can generate not just patterns of random, temporary price discounts for a variety sold to a consumer continuously purchasing from a *same* firm, which are reminiscent of those documented in the empirical literature, but also richer patterns of random price increases and decreases, depending on the difference in gross utility and informativeness across the two firms' varieties.¹⁹

As before, in equilibrium information can be either underprovided or overprovided. Consider first the case of information underprovision. Since in our earlier example, displayed in Figure 1, we supposed that firm B had only one uninformative variety, this same example applies (trivially) to the case of correlated priors. It is immediate to extend this example to the case in which the consumer's gross utility from purchasing firm B 's variety vary with the prior and obtain the same results.

More interesting is the case of information overprovision. The example of overprovision is similar to our earlier example in the independent case. Suppose, as depicted in Figure 3, that firm A has a perfectly informative variety, variety 2, and a less informative variety, variety 1, whereas B has one variety that is moderately informative. Briefly, in the planning solution the less informative variety of A is chosen in the first period. After a success this same variety is chosen, whereas after a failure firm B 's variety

¹⁹Although these pricing strategies are typically referred to as applying to new products, in our model they can be thought equivalently as pricing strategies for new products offered to existing consumers or as pricing strategies for existing products offered to consumers who are new to a market or are otherwise uncertain about their tastes for existing products. For empirical evidence on learning by consumers new to a market about existing products, see, for instance, Heilman et al. (2000).

is chosen. In equilibrium, firm A chooses the more informative variety and attracts the consumer in the second period, both after a success and after a failure.²⁰

The intuition for the inefficiency is simple. Here, firm A chooses the variety, variety 2, that allows it to retain the consumer in the second period regardless of the outcome in the first period. The planner, instead, chooses a variety for which firm A will lose the consumer after a failure in the first period. Firm A 's and the planner's different valuations at the nodes at which firm A loses to firm B account for the discrepancy between the choices in equilibrium and those by the planner. In sum, firm A strategically chooses the more informative variety because by doing so it ensures that it will retain the consumer in period 2 after both success and failure and hence earn higher profits than when it chooses the less informative variety.

C An Oligopoly Model

So far we have focused on a market with two firms, each of which has a unique production technology. Now, we consider a market with more than two firms and allow for the possibility that a technology is adopted by more than one firm. The markets considered here can be interpreted as resulting from different configurations of entry costs. For example, we think of the economy in which each technology is adopted by at least two firms as one in which the cost of adoption of all technologies is zero. We first show that the counterpart of our earlier match efficiency result holds but that conditional efficiency may fail. We then show that our dynamic pricing result immediately extends to this economy.

To show how conditional efficiency may fail, we first examine a market with three firms, each of which produces a single variety. We show that, unlike in the duopoly case, the equilibrium may be inefficient and typically features nonstatic pricing.

We then turn to examining how intensifying competition can restore efficiency. We start with an economy with many heterogeneous technologies but in which each technology is adopted by at least two firms. We show that greater competition between firms with the same technology drives prices to their competitive level and leads to efficiency. Yet, competitive pricing for purchased varieties is not sufficient for efficiency. We illustrate this point in an economy in which all but one firm is competitive and that last is a monopolist. Specifically, we show that even if the only non-competitive firm never sells in equilibrium, so that all observed prices are competitive, the equilibrium may be inefficient.

Finally, we consider an environment in which firms can sell multiple differentiated varieties but priors are variety specific rather than firm specific. We interpret this economy to be one in which retail stores sell a mixture of branded varieties. In this case, even if no two technologies are the same, pricing is competitive and the equilibrium is efficient as long as each variety can be offered by at least two firms.

Throughout this section we focus on the case of independent priors. It should be clear that our results extend in the natural way to the case of correlated priors, as they did in the duopoly case.

C.1 Match Efficiency and Dynamic Pricing

Let there be an arbitrary number of firms indexed by $f \in \{A, B, C, \dots\}$. Each firm has a technology defined as in the duopoly case. We let $\pi_t = (\pi_{At}, \pi_{Bt}, \dots)$ denote the beginning-of-period t priors. The definition of equilibrium here is the obvious extension of our earlier definition.

We begin by showing that the natural generalization of the match efficiency result of Proposition 1 holds. The first part of this result is that the variety choice of any firm f maximizes the match value of firm f , namely, $W^f(\pi) = V^f(\pi) + U(\pi)$. The second part of this result is that the match value solves

²⁰Parameter values that produce these outcomes are: $x_A(1, 1) = 15$, $x_B(1, 1) = 14.1$, $x_A(1, 2) = 11$, $x_A(0, 2) = 10.56$, $x_B(0, 1) = 10.4$, $x_A(0, 1) = 1.5$, $\alpha_{A1} = 0.380$, $\beta_{A1} = 0.379$, $\alpha_{A2} = 1$, $\beta_{A2} = 0$, $\delta = 0.75$, and $\pi_A = \pi_B = 0.9$.

a simple programming problem. To set up this problem, we introduce notation to denote the selling firm and the second–best firm at any prior π . To this end, note that an equilibrium partitions the set of priors into sets $\{E_A, E_B, E_C, \dots\}$, where firm A sells at priors $\pi \in E_A$, firm B at priors $\pi \in E_B$, and so on. We can use equilibrium strategies to define a function $f_1(\pi)$ that identifies at π the selling firm, which we refer to as the *first–best* or *best* firm at π . Similarly, we can also define a function $f_2(\pi)$ that identifies at π the firm whose offer at π yields to the consumer the same value of utility as the selling firm’s offer. We let $f_2(\pi)$ denote the *second–best* firm at prior π . (Clearly such a firm exists: if there is more than one such firm, then choose, say, the one with the highest technology index f .) We establish that the match value of any firm, say, firm A , solves

$$W^A(\pi) = \max\left\{\max_{k_A \in K_A} \{(1 - \delta)x_A(\pi_A, k_A) + \delta EW^A(\pi'|\pi, k_A)\}, \quad (56)$$

$$(1 - \delta)[x_f(\pi_f, k_f(\pi)) - p_f(\pi)] + \delta EW^A(\pi'|\pi, k_f(\pi))\},$$

where at $\pi \in E_A$, $f = f_2(\pi)$, and at $\pi \notin E_A$, $f = f_1(\pi)$. To interpret (56), when firm A sells at π , in the current period the sum of firm A ’s profits and the consumer’s utility is just the consumer’s gross utility, $x_A(\pi_A, k_A)$, whereas the future value is the value to this pair when the prior is updated using firm A ’s offered variety. Likewise, when firm A does not sell at π , in the current period firm A ’s profits are zero, so in the current period this sum is just the consumer’s period utility, $x_f(\pi_f, k_f(\pi)) - p_f(\pi)$, from a purchase from the best firm $f = f_1(\pi)$, whereas the future value is the value to firm A and the consumer when the prior is updated using firm f ’s offered variety.

Proposition 4. (*Match Efficiency under Oligopoly*) *For an arbitrary firm, say, firm A , the offered variety $k_A(\pi)$ by that firm, both when it sells and when it does not sell, maximizes the match value $W^A(\pi)$. Moreover, the match value $W^A(\pi)$ solves (56).*

The proof is an immediate extension of that of Proposition 1. The next proposition is also an immediate extension to the many–firm case of our earlier result.

Proposition 5. (*Dynamic Pricing*) *Let $f = f_2(\pi)$ denote the second–best firm. The price of the selling firm, say, firm A , is*

$$p_A(\pi) = x_A(\pi_A, k_A(\pi)) - x_f(\pi_f, k_f(\pi)) + \frac{\delta}{1 - \delta} [EW^f(\pi'|\pi, k_A(\pi)) - EW^f(\pi'|\pi, k_f(\pi))]. \quad (57)$$

Note that under oligopoly, the match value of a firm typically does not coincide with its autarky value as it did in the duopoly case, so we cannot reduce the compensating price differential in (57), the term in brackets, to the form in (22). Briefly, the reason for this lack of coincidence is that the first term in the maximization operator in (56) need not coincide with the second term, so this match surplus problem does not reduce to the autarky problem. In particular, if $\pi \in E^A$, then typically firm A strictly prefers selling to not selling (and the consumer is tied) so the first term is strictly larger than the second. Instead, if $\pi \notin E^A$ and A is not the second–best firm, then firm A is indifferent between selling and not selling but the consumer typically strictly prefers buying from the best firm than from firm A . Hence, the second term is strictly larger than the first term. In neither case does this value reduce to autarky. (Of course, if two firms, say, firms A and B , are the only ones that at any priors are the best and the second–best firms, then the match value reduces to the autarky value, since in this case all other firms are effectively irrelevant.)

C.2 Failure of Efficiency and Static Pricing

We show that efficiency and static pricing both may fail even when all firms have one variety. This result shows the limits to which the two results on efficiency and static pricing in BV's case of two firms with one variety can be extended. In terms of the failure of efficiency, we shed light on a type of inefficiency that is different from that in our earlier economies. Even though, trivially, firms choose the right varieties, the consumer buys from the *wrong firm* in equilibrium. Hence, the inefficiency stems from the failure of conditional efficiency. In this sense, our analysis shows that increasing competition simply by adding an additional firm can actually lead to inefficiency. In terms of pricing, the failure of static pricing stems from the fact that the match value of a firm no longer reduces to the autarky value. In particular, the continuation match value of the second-best firm can vary with the identity of the selling firm even if all firms have only one variety, which implies that the compensating price differential is nonzero.

C.2.1 The Logic of Inefficiency

Here we show where the logic of our earlier conditional efficiency result breaks down with more than two firms. This breakdown can easily be seen with three firms. In this case, the contradiction step in the proof of Proposition 2 no longer holds. Intuitively, in the equilibrium losing firms care about the identity of the winning firm because the information revealed by the consumer's experience with the winning firm's variety affects the ability of such losing firms to attract the consumer in the future. The planner weighs the preferences of all such losing firms when choosing an optimal plan, whereas, in the equilibrium, the preferences of the third-best or any lower firm for which firm sells to the consumer has no impact on the equilibrium. Indeed, the equilibrium does not restrict how much any losing firm prefers any other losing firm over the winning firm to sell to the consumer: equilibrium just requires all losing firms to be indifferent between winning the consumer and losing the consumer to the selling firm. This difference in how the preferences of losing firms are taken into account in equilibrium and by the planner is the source of the inefficiency.

To see why, let us retrace the steps of the contradiction argument in the proof of Proposition 2 and see where the logic fails. To start, suppose that in the equilibrium in the first period at the prior π , firm A is the best firm, firm B is the second best and firm C is the third best. But, suppose by way of contradiction that at that prior, the planner prefers firm B to sell in the first period rather than firm A . Let $V^f(\pi|f')$ and $U(\pi|f')$ denote, respectively, the value of firm f 's profits, $f = A, B, C$, and the consumer's value when the consumer purchases from firm $f' = A, B, C$. For firm A to prefer to sell in the equilibrium rather than to raise its price and lose to firm B , it must be that

$$V^A(\pi|A) \geq V^A(\pi|B), \quad (58)$$

whereas for the planner to prefer firm B to firm A , it must be that

$$V^A(\pi|B) + V^B(\pi|B) + V^C(\pi|B) + U(\pi|B) > V^A(\pi|A) + V^B(\pi|A) + V^C(\pi|A) + U(\pi|A). \quad (59)$$

Here, the values on the left side of (59) are obtained by evaluating profits and utility off the equilibrium path when firm B sells at prior π .

Now, since B is the second-best firm, the consumer is indifferent between purchasing from firm B and purchasing from firm A so that $U(\pi|B) = U(\pi|A)$. By the cautious restriction, firm B is indifferent between selling to the consumer and losing to firm A , so that $V^B(\pi|B) = V^B(\pi|A)$. Using these two equalities to simplify (59), we obtain that

$$V^A(\pi|B) + V^C(\pi|B) > V^A(\pi|A) + V^C(\pi|A). \quad (60)$$

As long as the third–best firm, here firm C , has a sufficiently strong preference for the second–best firm, firm B , selling to the consumer rather than firm A selling to the consumer, that is, as long as

$$V^C(\pi|B) - V^C(\pi|A) > V^A(\pi|A) - V^A(\pi|B), \quad (61)$$

no contradiction arises between (58) and (60). Hence, here an inefficient equilibrium is consistent with profit maximization by all firms. (Recall that when there are only two firms, the only losing firm is necessarily the second best firm and that firm is indifferent between selling or not so that (60) reduces to $V^A(\pi|B) > V^A(\pi|A)$ which clearly contradicts (58). No such contradiction arises here.)

One way to think about this situation is that the third–best firm, here C , would be willing to reimburse the winning firm, firm A , for the loss of profits that A would incur if it purposely lost to firm B . If firm C did so, then firm C 's net gain in moving from an equilibrium in which A sells to an outcome (not an equilibrium) in which B sells would be

$$V^C(\pi|B) - V^C(\pi|A) - [V^A(\pi|A) - V^A(\pi|B)] > 0,$$

where the term in brackets equals firm C 's payments to firm A . That the net gain for firm C is positive follows from (61). And, given these payments firms A and B and the consumer would all be indifferent and firm C would be strictly better off.

C.2.2 Inefficiency under Oligopoly

Consider an economy with three firms, A , B , and C , with one variety each.²¹ The economy, as before, can be interpreted as an instance either of a two–period economy in which the future is discounted at rate $\hat{\delta} = \delta/(1 - \delta)$ or of an infinite horizon economy in which utility realizations from the varieties of firms A and B are dependently distributed over time. The infinite horizon interpretation is the natural analog of our interpretation in the duopoly examples, discussed in the Appendix. Specifically, we assume that if the variety of either firm $f = A, B$ has not been purchased by the consumer in period $t = 1$, then *regardless of the state of the world*, in each period $t \geq 2$ utility realizations are distributed as follows: with probability γ_f , the consumer receives utility X_{fHk} with probability α_{fk} and utility X_{fLk} with probability $1 - \alpha_{fk}$; with probability $1 - \gamma_f$, the consumer receives utility X_{fHk} with probability β_{fk} and utility X_{fLk} with probability $1 - \beta_{fk}$. We let $\gamma_f = \pi_{f2}$ so that, given this modified information structure, $\pi_{ft} = \pi_{f2}$ at any $t > 2$. Then, the following derivations apply to both the finite horizon and the infinite horizon cases.

As noted above, our infinite horizon interpretation with a dependently distributed stochastic process can be thought of somewhat loosely as a stochastic analog of the deterministic economy that BV2 consider in Section 6 of their paper. They focus on an economy with time–varying payoffs and in Section 6 assume that payoffs for a product become constant after a certain number of uses. For that economy, BV2 find a unique cautious Markov–perfect equilibrium that is efficient, whereas we find a cautious Markov–perfect equilibrium that is inefficient. Thus, our result clarifies that BV2's results on efficiency in a deterministic economy do not immediately generalize to a stochastic economy with arbitrary dependent stochastic processes.

We now turn to the details. Since each firm has one variety, we simplify notation by letting $x_f(\pi_f)$ denote the consumer's gross utility for the single variety of firm f . We let firm A and firm B have perfectly informative varieties and firm C have an uninformative variety. We assume that at the initial

²¹We can think of the market considered before as also consisting of firms A , B , and C , but with C operating such a dominated technology compared with those of A and B that C never sells.

prior $\pi = (\pi_A, \pi_B, \pi_C)$, gross utilities satisfy the restrictions implicit in Figure 4, including

$$x_A(1) > x_B(1) > x_B(\pi_B) > x_C > x_A(\pi_A) > x_A(0) > x_B(0), \quad (62)$$

where the constant x_C equals $x_C(\pi_C)$. Note that Figure 4 shows how the consumer's gross utility from purchasing from a firm varies as a function of the prior about that firm. Of course, as the prior, say, π_A , varies, the consumer's gross utility from purchasing from firms B and C does not vary. (In this sense, for convenience, Figure 4 combines the three graphs of $x_f(\pi_f)$ against π_f for $f = A, B, C$ into one graph.)

We make assumptions so that in period 1 the planner chooses firm B . From (62) it follows that after a success of the variety of firm B , the planner chooses firm B again, but after a failure the planner chooses firm C for a total value of gross utility of

$$(1 - \delta)x_B(\pi_B) + \delta [\pi_B x_B(1) + (1 - \pi_B)x_C]. \quad (63)$$

In the equilibrium, instead, firm A sells in the first period. After a success, firm A also sells in the second period, but after a failure firm B sells for a total value of gross utility of

$$(1 - \delta)x_A(\pi_A) + \delta [\pi_A x_A(1) + (1 - \pi_A)x_B(\pi_B)]. \quad (64)$$

The price charged by the selling firm, say, firm A , is

$$p_A(\pi) = x_A(\pi_A, k_A(\pi)) - x_f(\pi_f, k_f(\pi)) + \frac{\delta}{1 - \delta} [EW^f(\pi'|\pi, k_A(\pi)) - EW^f(\pi'|\pi, k_f(\pi))], \quad (65)$$

where $f = f_2(\pi)$ is the second-best firm at π .

Proposition 6. (Oligopoly Inefficiency) *If (63) is greater than (64) then the equilibrium is inefficient.*

The sufficient conditions for this equilibrium to exist are (62) together with a simple inequality we derive in the proof.

C.3 Intensifying Competition

We now consider two ways of intensifying competition between firms that can restore efficiency. First, we suppose that two or more firms adopt each technology. For a motivation, imagine a market for prescription and over-the-counter drugs in which two pharmacies, say, CVS and Walgreens, sell exactly the same varieties of antibiotics, toothpaste, pain relievers, and so on. We interpret this case as one in which technologies can be freely adopted by firms. We show that if every technology is adopted by at least one firm, then efficiency is restored and pricing is competitive. We then use a variant of this example in which one firm has an exclusive technology, whereas all other firms do not, to show that the converse is not true: even if all purchased varieties are priced competitively, the equilibrium may not be efficient.

Second, we suppose that priors are variety specific and interpret this case as corresponding to a market in which retailers sell different mixes of potentially overlapping varieties. Since priors are variety specific and the same variety can be sold by multiple sellers, here brands in our earlier sense are not relevant. Here a retail firm simply sells collections of varieties, each of which has no information about another variety. For example, consider three supermarket chains, say, Safeway, Cub Foods, and Rainbow, and suppose, for concreteness, that each of them sells two of three cereals' varieties: Frosted Flakes, Cocoa Puffs, and Harvest Crunch. Specifically, suppose that Safeway sells Frosted Flakes and Cocoa Puffs, Cub Foods sells Cocoa Puffs and Harvest Crunch, and Rainbow sells Frosted Flakes and Harvest Crunch. We interpret this

case as one in which each variety of a technology can be freely adopted by the retailer. We show that, even though firms are heterogeneous, the intense competition between varieties also restores efficiency.

C.3.1 Competition by Technology

Consider a market in which firms operate different technologies $\{A, B, C, \dots\}$ that produce multiple product varieties, and assume that each technology has been adopted by at least two firms. We refer to this economy as a *replica economy*. In our earlier duopoly setup, since each firm was uniquely identified by its technology, we could have interpreted the prior to be equivalently about the match of the consumer with a particular firm or about the match of the consumer with a particular technology for producing goods. Here, multiple firms use the same technology, and we define the prior to be about the match of the consumer with each technology. This setup captures the idea that when firms sell varieties produced with a same technology, consumers learn about the value of their idiosyncratic matches with these varieties regardless of where the consumers purchase them. For example, suppose both Walgreens and CVS sell an identical variety of Colgate toothpaste. An unsuccessful experience with the toothpaste changes a consumer's valuation of the product as well as other Colgate oral care products such as toothbrushes, regardless of whether a consumer purchased the toothpaste from Walgreens or CVS.

We start with an extreme case in which all of the products of Walgreens and CVS are from Colgate. Below we relax this assumption by considering a more general case in which firms sell a mixture of varieties of different underlying brands, so that Walgreens might sell Colgate toothpaste and Palmolive dish washing soap, whereas CVS might sell Colgate toothpaste and Oral-B toothbrushes and so on. We provide conditions under which the equilibrium is efficient even if no two firms have the same technologies.

A Replica Economy. Suppose that there are (at least) two firms operating each technology. In a slight abuse of notation, we refer to the two firms that operate technology A as A_1 and A_2 , and we use similar notation for the firms operating the remaining technologies.

First, consider prices. If some firm, say, firm A_1 , sells at a given prior π , then the second-best firm is firm A_2 . Since firms A_1 and A_2 have identical technologies, they both choose the same variety at this prior. By (57), these two facts imply that $p_A(\pi) = 0$: the intense competition between firms with the same technology makes them bid their prices down to their marginal costs, which are zero. Since this outcome occurs at each prior, the profits of all firms are zero.

Next, consider the variety choices of each firm. Since the value of profits of each firm, $V^f(\pi)$, is zero, the match value of each firm, $W^f(\pi) = V^f(\pi) + U(\pi)$, reduces to the consumer's value, $U(\pi)$. Moreover, since at each prior the prices of all firms are zero, the consumer's utility, $u_f(\pi_f, k(\pi), p(\pi))$, reduces to $x_f(\pi_f, k(\pi))$ so that the match value of each firm reduces to the value of the consumer's gross utility. Thus, conditional on being chosen by the consumer, each firm offers the same variety that the planner would choose. Finally, since the consumer purchases from the firm whose offer leads to the highest value of gross utility, consumer optimality implies that the consumer selects the same firm that the planner would select. Hence, equilibrium is efficient.

Proposition 7. (*Replica Economy*) *In a replica economy, offered varieties are efficient and firms charge prices equal to their marginal costs.*

An Almost Replica Economy: One Monopolist and Many Competitive Firms. In the replica economy, all firms had replicas and acted in a competitive fashion. Here, we imagine that all but one firm has a replica. In equilibrium, the firms with replicas are competitive in that they price at marginal cost. The firm without a replica is a monopolist in that it exclusively operates a technology. Not surprisingly, the

monopolist can make inefficient variety choices in such an environment. We find it interesting that even if the monopolist never sells in equilibrium, its mere presence can cause the competitive firms to make inefficient variety choices. The lesson we draw here is that even though an outside observer armed only with price data would infer that the equilibrium is competitive and therefore efficient, the observer would be mistaken. In this sense, competitive pricing is not sufficient for efficiency.

For concreteness, suppose that technology A is adopted by only one firm but that technology B is adopted by two firms, denoted by B_1 and B_2 for simplicity. We use similar notation for firms using the remaining technologies. We refer to the firms with technologies $\{B, C, D, \dots\}$ as the *competitive firms*.

Using the same logic as in Proposition 7, it is immediate that the competitive firms price their varieties at their marginal cost (of zero) in all periods and, thus, make zero profits. Nonetheless, the competitive firms may make inefficient choices of varieties. Intuitively, the competitive firms choose those varieties that maximize the value of the consumer's utility. In contrast, the planner chooses the varieties of the competitive firms so as to maximize the sum of the values of the consumer's utility and the monopolist's profits.

To see how this difference in valuations can lead to inefficiencies—for instance, to information being underprovided compared with the efficient level—consider an economy in which a monopolist with technology A faces competitive firms with technology B . Suppose for simplicity that technology B has two varieties, with variety 1 somewhat informative and variety 2 perfectly informative, whereas technology A has only one variety, which is assumed for simplicity to be uninformative. Since the value $x_A(\pi_A)$ is independent of π_A , we denote this value simply by x_A . Suppose utilities are ordered as in Figure 5, so that

$$x_B(1, 1) > x_B(1, 2) > x_A > x_B(0, 2) > x_B(0, 1). \quad (66)$$

In this case, the second-period outcomes are as follows. If a firm of type B sells variety 1 in the first period, it sells the same variety again in the second period regardless of the utility realized in the first period. If a firm of type B sells variety 2 in the first period, then in the second period it sells variety 1 after a success and loses the consumer to firm A after a failure.

It is possible to find conditions under which in equilibrium the consumer purchases from a firm of type B variety 1 in the first period, whereas under the planner's solution the consumer would purchase variety 2 in the first period. For this to be the case for the planner, it must be that

$$(1 - \delta)x_B(\pi_B, 2) + \delta[\pi_B x_B(1, 1) + (1 - \pi_B)x_A] > x_B(\pi_B, 1). \quad (67)$$

In equilibrium, a firm of type B chooses variety 1 if the match value from variety 1 is greater than that from variety 2, that is,

$$x_B(\pi_B, 1) > (1 - \delta)x_B(\pi_B, 2) + \delta[\pi_B x_B(1, 1) + (1 - \pi_B)x_B(0, 2)]. \quad (68)$$

Equilibrium is inefficient when (67) and (68) simultaneously hold.²²

In this equilibrium, all purchased varieties are sold by the competitive firms, which set their prices equal to their marginal costs. The subtle inefficiency here is that the mere presence of the monopolist lurking off the equilibrium path can imply that the competitive firms make inefficient variety choices.

If we imagine that the monopolist had to pay a cost to adopt its technology, then since the monopolist never sells in equilibrium, an implication of this example is that the monopolist would make negative profits. Hence, the monopolist should not have adopted its technology in the first place. With a slight

²² Parameter values that produce these outcomes and satisfy these two inequalities, as well as the assumptions in (66), are: $x_B(1, 1) = 10$, $x_B(1, 2) = 8.5$, $x_A = 4$, $x_B(0, 2) = 3$, $x_B(0, 1) = 2$, $\alpha_{B1} = 0.60$, $\beta_{B1} = 0.57$, $\alpha_{B2} = 1$, $\beta_{B2} = 0$, $\delta = 0.55$, and $\pi_B = 0.6$. Note that α_{A1} , β_{A1} , and π_A play no role in these calculations.

modification to the economy, we can eliminate this prediction. Suppose that we introduce a new period, period 0, in which the monopolist must decide whether or not to pay the adoption cost, and then nature stochastically draws vectors of priors for period 1. Under some of these priors, the prior π_B about the technology of the competitive firms will be low enough that the monopolist will sell at a positive price and make positive profits in the continuation game. We suppose that the distribution of initial priors is such that the value of the monopolist's profits is positive. Then, the above example refers to the branch of this modified game with a relatively high prior π_B at which the monopolist does not sell.²³

C.3.2 Competition by Variety

Suppose now that each consumer has a different prior about each variety, but at least two firms can produce each variety. Hence, each firm has a set of varieties that overlaps only partially with that of other firms. Formally, let there be K varieties with state of the world $\theta = (\theta_1, \theta_2, \dots, \theta_K)$ and $\theta_k \in \{\bar{\theta}_k, \underline{\theta}_k\}$, $1 \leq k \leq K$. The prior at t is $\pi_t = (\pi_{1t}, \dots, \pi_{Kt})$. Each firm f has a subset $K_f \subset K$ of these varieties, each of which also belongs to the variety set of at least another firm. For instance, firms A and B can produce variety k , and the prior about this variety is updated in the same way regardless of which of the two firms sells variety k . Here we show that even if each firm sells different bundles of varieties, the intense competition between varieties restores efficiency.²⁴ To relate this setup to our motivating example with three supermarket chains, note that having firms A and B both “produce” the same variety, say Kellogg’s Frosted Flakes, means that two supermarket chains A and B buy goods from a wholesaler at some constant cost, which for notational simplicity we set to zero.

Proposition 8. *(Overlapping Varieties) Consider the variety-specific prior economy in which firms have different subsets of varieties, but each variety can be offered by two or more firms. With a finite horizon, the equilibrium is efficient and firms charge prices equal to their marginal costs. With an infinite horizon, there exists an efficient equilibrium in which firms charge prices equal to their marginal costs.*

²³ It is easy to construct examples in which the monopolist sells and the compensating price differential is positive.

²⁴ It is possible to construct examples in which the equilibrium is inefficient when only a single variety in the market is not produced by two or more firms.

Table 1: Parameter Estimates and Prices

	X_L	X_H	Uniform Price	α
Apple phone	67	1,000	640	0.9882
Apple tablet	-101	1,021	560	0.9677
Samsung phone	-444	550	400	0.9802
Samsung tablet	-751	340	340	0.9977

Table 2: From Uniform Pricing to Personalized Pricing

	Change (\$ per consumer per year)
Consumer surplus	108.30
Prices paid	-41.79
Apple profit	-83.19
Samsung profit	-24.99
Total profit	-108.18
Total welfare	0.12

Figure 1. Underprovision of Information (Firm-Specific Priors)

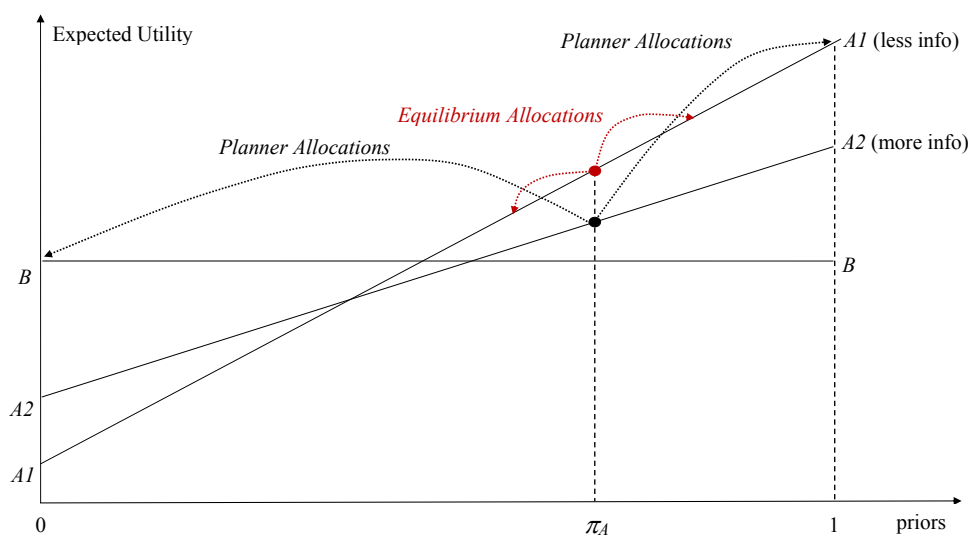


Figure 2. Overprovision of Information (Firm-Specific Priors)

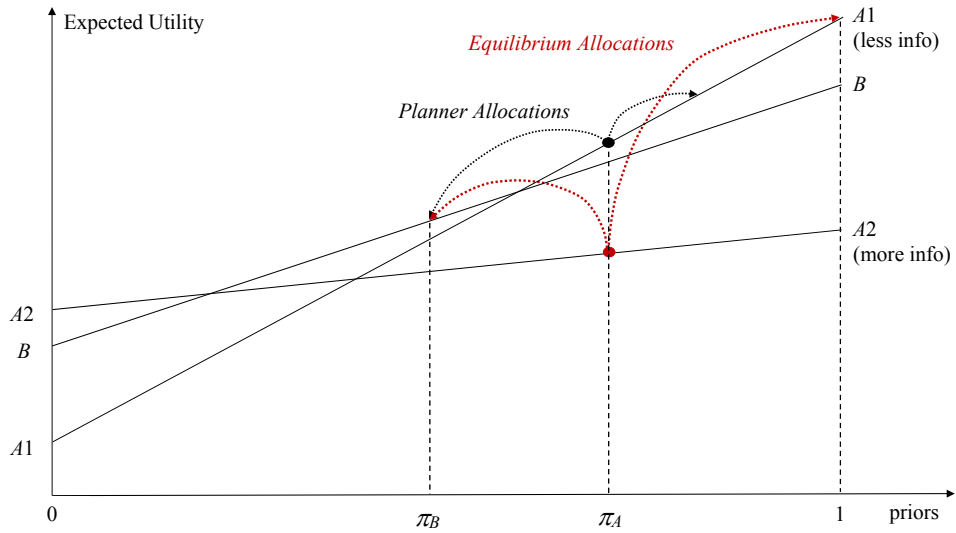


Figure 3. Overprovision of Information (Correlated Priors)

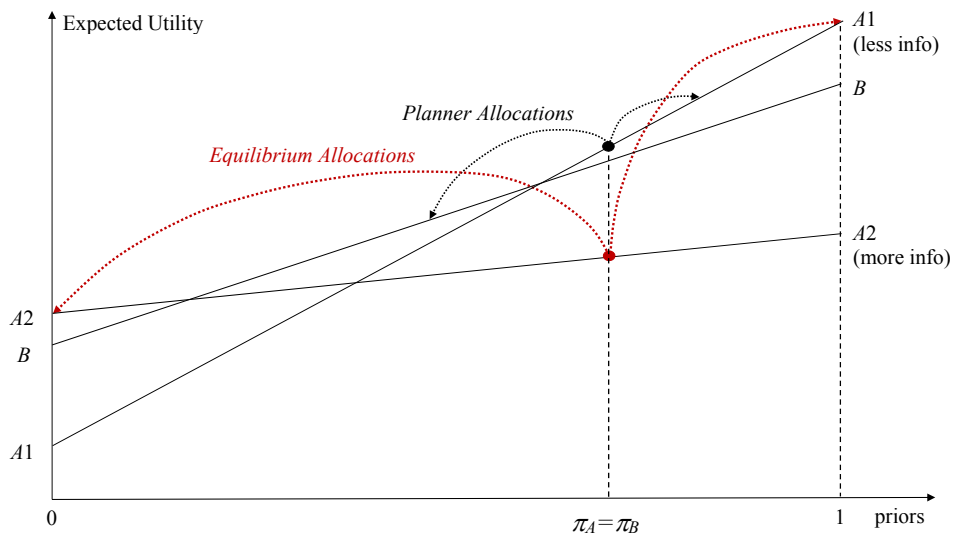


Figure 4. Inefficiency of Oligopoly with Three Firms (Firm-Specific Priors)

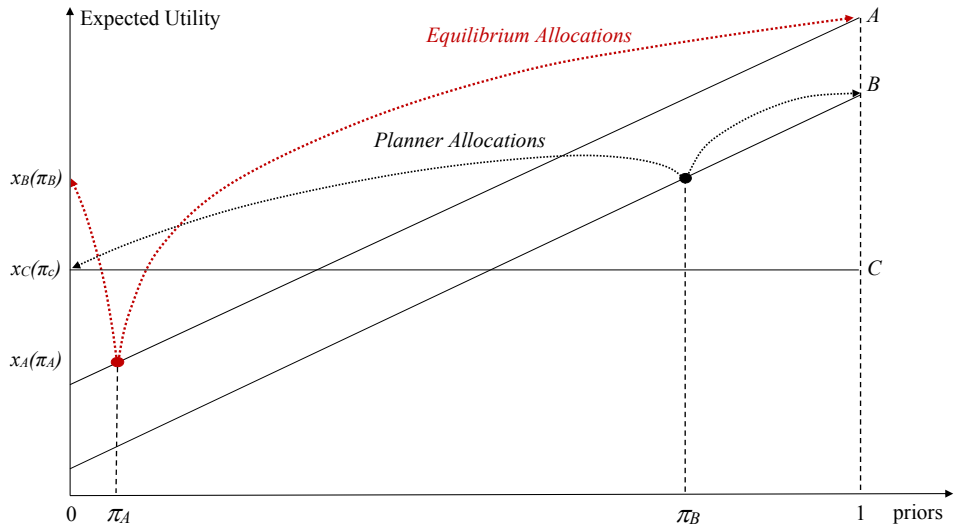


Figure 5. Underprovision of Information With One Monopolist and Many Competitive Firms (Technology-Specific Priors)

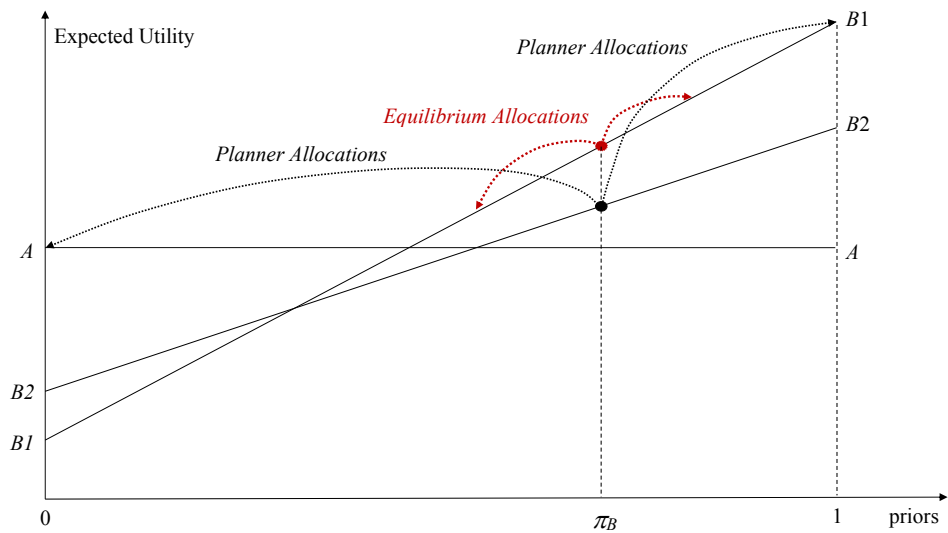


Figure 6: Prior Distribution for Apple

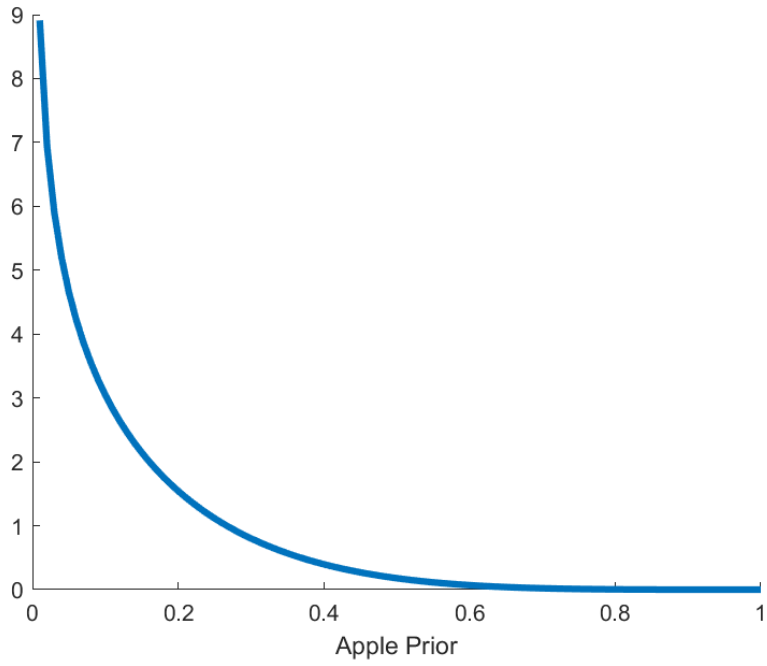


Figure 7: Prior Distribution for Samsung

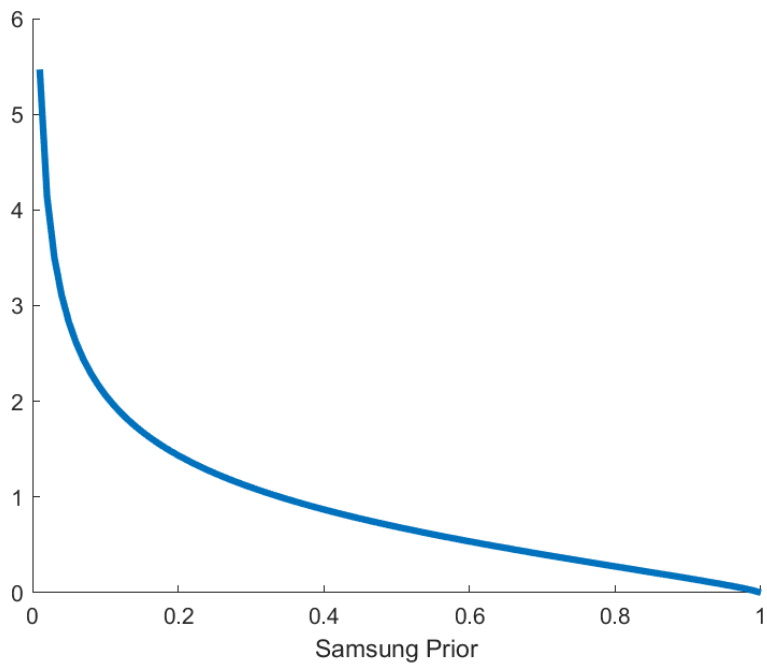


Figure 8: Counterfactual Results for Consumer Surplus

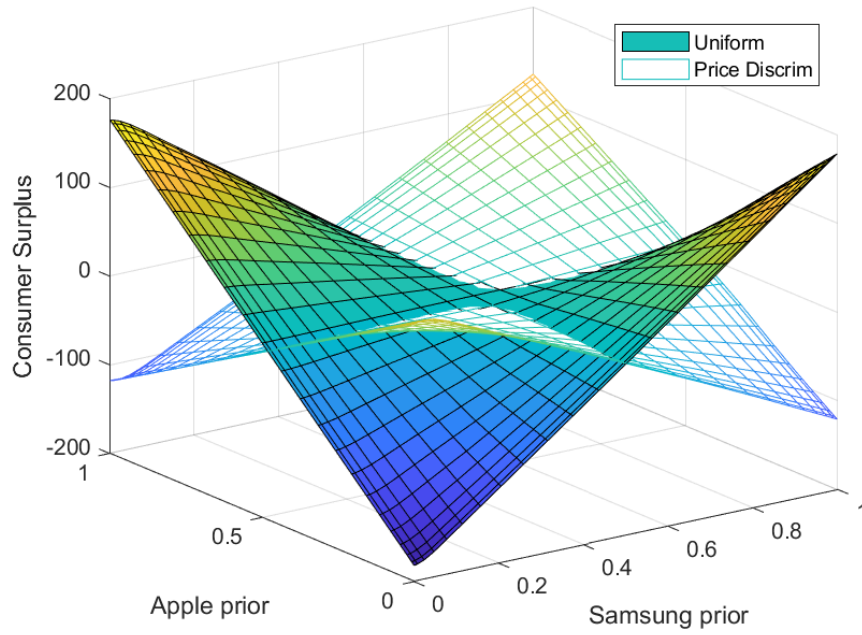


Figure 9: Counterfactual Results for Prices

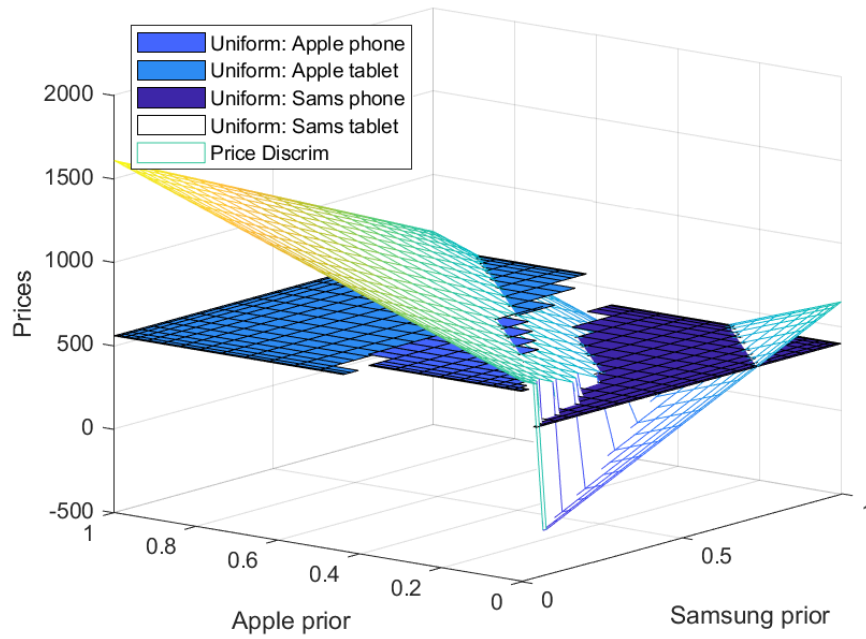


Figure 10: Counterfactual Results for Profits (Apple and Samsung)

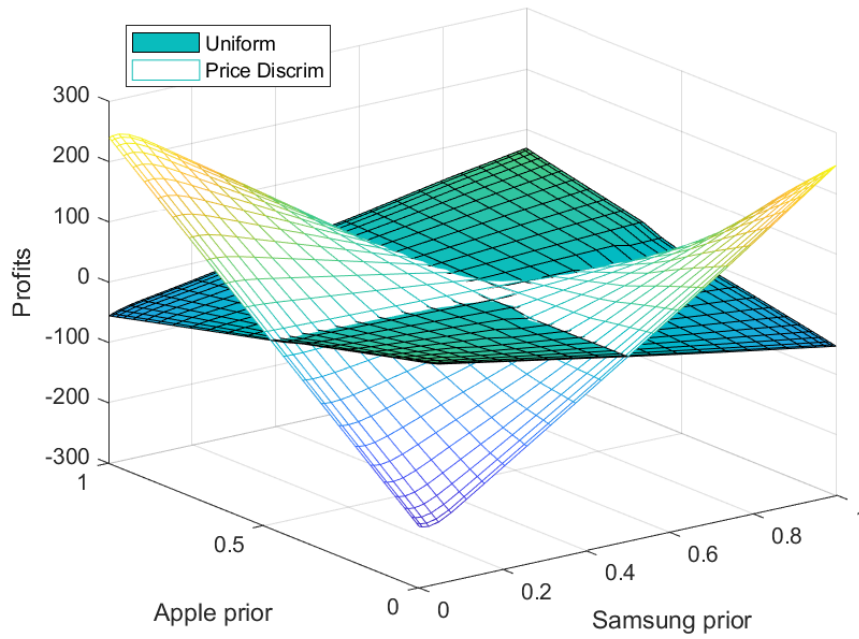


Figure 11: Counterfactual Results for Total Welfare

