American Economic Journal: Microeconomics 2017, 9(4): 42–73
https://doi.org/10.1257/mic.20140182

If Many Seek, Ye Shall Find: Search Externalities and New Goods†

By Maciej H. Kotowski and Richard J. Zeckhauser*

Consumer search serves productive roles in an economy with multiple goods. In equilibrium, search promotes the sorting of consumers among producers, thereby enabling the market for new goods, and potentially increasing welfare and profits above the benchmark case (an economy with a single good, hence, no search). When competitors are few, additional direct competitors may benefit a firm, as more sellers may encourage more consumers to search. In return, consumer search entices producers of new goods to enter. Neither of these externalities, nor the coordination problems faced by consumers and producers, is appropriately recognized in the literature. (*JEL* D11, D43, D62, D82, D83, G22)

"Too few competitors." That is hardly a complaint of firms in the economics literature. However, contrary to both the literature and standard intuition, a firm may suffer if there are too few other firms producing its product. This surprising conclusion applies in the not unusual case where two or more different goods compete, and consumers need to find the producers of a product they desire. For a viable market to emerge and to be sustained requires an appropriate interaction of consumers and producers. Sellers enter the market confident that when their numbers are sufficient, well-matched buyers will search to find them. Buyers will search only if the goods they want are available and sufficiently easy to find.

The following analogue of a common situation illustrates our argument. Juan owns and operates a taxi in a large city. Since taxis are alike in terms of quality and price, consumers rationally hire the first available cab. Taking this random matching as given, for argument's sake assume that each driver makes about 30 trips per day. Suppose Juan invests in a more comfortable car that will be appreciated by a third

* Kotowski: John F. Kennedy School of Government, Harvard University, 79 JFK Street, Cambridge, MA 02138 (email: maciej_kotowski@hks.harvard.edu); Zeckhauser: John F. Kennedy School of Government, Harvard University, 79 JFK Street, Cambridge, MA 02138 (email: richard_zeckhauser@harvard.edu). We are grateful to the editor and the referees for suggestions that greatly improved this study. We are grateful to Daniel Hojman for stimulating discussions on this and related topics. Audiences at the 2014 CEA meeting (Vancouver), the 2014 European Summer Meeting of the Econometric Society (Toulouse), and the 2015 ASSA Annual Meeting (Boston) offered appreciated constructive feedback. Part of this paper was completed when Maciej H. Kotowski was visiting the Stanford University Economics Department. He thanks Al Roth and the department for their generous hospitality. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

† Go to https://doi.org/10.1257/mic.20140182 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

of the city's taxi users. These customers are willing to pay a premium for the added comfort of a luxury vehicle, and Juan charges them accordingly. For simplicity, say Juan needs to make only 18 trips per day at this higher price to recoup his investment.[1] At first glance, Juan's entrepreneurial investment seems like a wise business move. He intuited a large untapped market, and he went after it.

Despite the large latent demand for Juan's unique service, we argue that he is unlikely to be successful. Why? Although in aggregate many consumers want added comfort, Juan is unlikely to encounter enough of them in a typical day. Indeed, unless comfort-conscious consumers alter their behavior and purposefully wait or search for his car, Juan will complete only ten trips per day and will fall short of his goal.[2] When there are very few luxury taxis in the city, even comfort-conscious consumers rationally take the first available cab. If the chance of encountering a fancy taxi is very small, why incur the cost of waiting?

Perhaps paradoxically, Juan needs more direct competitors to stand a chance of being successful with his new vehicle. If there were more luxury cabs, some consumers would wait since they could expect that a deluxe taxi would be around shortly. If enough consumers think it worthwhile to wait, Juan would be able to succeed. However, the market's outcome is still far from obvious, as competition is a double-edged sword. While more luxury vehicles may tempt consumers to seek out such taxis—effectively creating the market for this service where none existed before—the increased competition from other drivers cuts into profits. Equilibrium demands that these effects balance gracefully.

One might initially think that Juan's problem is related to the excessive costs of effective advertising,[3] or with the difficulty consumers may have in contacting him specifically for a ride. But, consider the extreme situation where Juan drives for Uber or a similar, internet-enabled, ride-sharing platform. Would a typical consumer even consider the on-demand luxury car service a realistic option if there was but one driver serving the entire city? Likely not. Rather, Juan's success requires the influx of other drivers offering a similar service and who can also be found with similar ease. That is, he needs more competition.

The benefits of more competitors apply even when each supplier is providing a slightly different good. Thus, a file-sharing service that seeks to compete with Dropbox may benefit if a few other similar services enter the market. This makes it more likely both that consumers will be aware of alternatives generally, and they will believe it worthwhile to seek a service more appropriate to their needs in terms of capabilities and price. Online streaming music and dating/match-making applications exhibit the same characteristics. In these cases, advertising is insufficient to adequately inform consumers about each product's characteristics. Consumers must search and experiment to find the product that works best for them.

[1] Although less common in North America, in many cities there is sizable variation in both the quality of vehicles in the taxi fleet and in the fares they charge. For example, in Singapore (mid-2013) the flag rate for a basic Toyota taxi is around S$3.00. This charge increases to around S$4.50 for a Mercedes taxicab.

[2] Only one-third of the 30 consumers whom Juan encounters randomly are willing to pay the higher posted price.

[3] When a seller does advertise, consumers receiving and responding to the advertisement can go directly to him. Others' search is not directed and random as we assume.

This article studies large, almost competitive markets where consumers incur a cost to find the producers of desired products. It addresses the conditions under which competition, coupled with consumer search behavior, can create a market for new goods. When it does, both consumer welfare and firms' profits can rise. We argue that successful markets require that both sides take costly actions. Firms must invest to enter a market while consumers must undertake a proactive and costly search to seek out desired goods. The improved sorting of consumers due to their search behavior yields an efficiency gain that producers and consumers share. We have already sketched the essence of our argument in the market for transportation and online applications. Similar conclusions apply, for example, to restaurants, professional services, and a range of products capitalizing on twenty-first century technologies. An example at the end of our essay applies our framework to insurance markets subject to adverse selection.

We assume that consumers must search to learn producers' offers. As noted by Diamond (1971), the existence of even a small search cost can seriously impede a market's operation. In an outcome later dubbed the "Diamond paradox," the market's only equilibrium involves all firms charging the monopoly price. Applied to our motivating discussion, in this outcome all taxis charge the same high price for standard service. If consumers believe that all taxis are the same, it is easy to appreciate the situation's self-fulfilling logic.

This undesirable outcome is neither theoretically inevitable, nor often observed in the real world. It can be avoided if there is consumer-level heterogeneity and the potential for multiple goods. We take the simplest case where producers can supply one of two goods. For example, a taxi may be either a standard or a luxury vehicle, as in our motivating vignette. Initially, each consumer knows the product offered by a local, nearby supplier and he knows the distribution of available products in the economy as a whole. If he finds the nearby offer acceptable, he takes it. If it is unacceptable, he searches for a better deal.

We investigate two related questions. First, we consider the welfare and profit consequences of search and competition. When consumers search for a desired product, they expand the demand faced by that product's producers. Increased demand enhances the incentives of firms to provide the product and, as the market expands, this virtuous circle reinforces the effectiveness of consumers' search efforts. By this process, search behavior generates a market-creation externality operating through the self-sorting of consumers across producers of different goods. This externality, we believe, has been neglected in previous examinations of markets where consumers search for desired goods. (We discuss the related literature below.) To highlight one consequence, even though consumer search is costly, and their search tamps down producers' market power, the resulting equilibrium can Pareto-dominate the benchmark case. Compared to Diamond's single-good, monopoly-price equilibrium, multiple goods flourish in the market, consumers enjoy greater expected surplus, and firms earn higher profits.

How might such an equilibrium come about? That is our second focus and it sounds two recurrent themes—product differentiation and price dispersion. We construct equilibria where multiple goods trade at multiple prices and we show how the combined incentives offered by product differentiation and price dispersion support the appealing welfare consequences noted above. Two considerations render this conclusion nontrivial. First, consumers must have the incentive to search. Hence, some firms must promise consumers a good deal. Second, and cutting against the first point, firms face a free-rider problem when consumers are already searching. No firm wishes to be the one that promises the best deal as it can often get away with leaving just a little less surplus to searching consumers than other firms. However, if all firms succumb to this temptation, the constellation of incentives supporting the market unravels—consumers' reason to search evaporates. In this regard, our story goes beyond a simple coordination game. One side (firms) faces a steady incentive for *mis*coordination among its members. Reconciling these opposing forces is the focus of our technical analysis.

Consumer advocates often encourage people to "shop around" for a better deal. Many fail to follow that advice.[4] In some markets search offers little immediate gain. Hence, as in the case of taxi services in Juan's city when there are few luxury cabs, even luxury-seeking consumers may rationally eschew search. On the other hand, in some markets extensive search appears to be consumers' modus operandi. The market for air transport or packaged vacations is one example, at least for many travelers.[5] Our study emphasizes the transition from one search regime to another. In our framework, it is a lost cause for one or a few consumers alone to search for a better deal. Rather, welfare gains come only when many consumers search simultaneously. Although a search-intensive equilibrium may appear wasteful at first glance—many consumers are pursuing economic rents all the while incurring direct search costs—we argue that by changing producers' incentives Pareto-improving gains can be realized. Producers' incentives change on two margins. First, they face pressure on the intensive margin as better-informed consumers curtail their market power. Second, producers gain, now profitable, new opportunities on the extensive margin. When consumers search, they self-sort among producers thereby creating an opening for new products and services. When the second effect predominates, efficiency is enhanced, and welfare and profits can both rise.

Three related claims constitute the primary economic contributions of our paper. First, the benefits of costly search activity are intimately tied with the aggregate prevalence of search behavior. Searching alone is often ineffective as a desired product or good is likely not on offer. If many search, desired products will spring up. Resolving this coordination problem and moving an economy from a low-search to a high-search equilibrium often yields benefits for consumers and producers alike. To sustain such coordinated action, firms must ensure that consumers' costly activities are worthwhile. Thus, and second, firms often benefit from direct competitors. When this is the case, each consumer is sure that his investment is not in vain. Third, intensive search behavior among consumers facilitates the introduction of

[4] For examples, see Pratt, Wise, and Zeckhauser (1979). Even for goods that are standardized, relatively expensive, repeat purchases, and for which search is in principle easy, prices can vary dramatically. For example, Medigap insurance—-a product to cover co-payments and deductibles in Medicare—-is standardized by the government. In 2003 in Colorado, the monthly premiums for a 75-year-old for Plan C coverage ranged from \$58 to \$271 per month. See Cutler and Zeckhauser (2004).

[5] Actively searching in these markets may be worthwhile as price dispersion is common (Borenstein and Rose 1994; Gerardi and Shapiro 2009). See also Sengupta and Wiggins (2014).

new goods to the possible benefit of all. Despite the costs and inconvenience of search, consumers gain from the new product's presence. Suppliers of the novel good gain from their new commercial success. And, most surprisingly, suppliers of *old* products can gain as well. They can now take advantage of the enhanced sorting of consumers among products offered in the economy. Our fourth contribution flows from a narrower technical perspective. We identify necessary and sufficient conditions for the existence of multiple equilibria in our model and, in particular, the existence of an equilibrium featuring price dispersion among all goods in the economy and search by all types of consumers. As we explain below, it is precisely in such cases that beneficial outcomes described above are most likely to occur.

After providing context for our study in Section I, we introduce our model in Section II. Sections III and IV present the equilibrium analysis. We sketch extensions to our model in Section V. To highlight but one variation, we extend our model to allow for adverse selection, which we then use to study insurance markets. We show how search activity can sustain welfare- and profit-enhancing insurance products through better consumer sorting. An online Appendix collects proofs and other technical remarks.

I. Related Literature

Since Stigler (1961) and Rothschild (1973), a large literature has examined consumer search and market equilibrium. The findings can be quite surprising. As Diamond (1971) showed, even a small search cost can squash any possible benefits from competition and lead to an all-charge-the-monopoly-price outcome. We take the monopoly-price outcome as the benchmark case, and investigate how *price dispersion* and *product differentiation* can together support Pareto-superior outcomes when consumers search. The equilibria we construct depend on the combined presence of *search externalities and complementarities* between consumers and firms. Our study relates to prior contributions on these three dimensions, as explained below.

Price Dispersion.—Price dispersion is a commonly observed phenomenon and many theories have been developed to explain its presence.[6] Reinganum (1979) focuses on firm-level heterogeneity. Stahl (1989, 1996) considers differences among consumers in their search costs. Burdett and Judd (1983) propose a nonsequential/noisy search process. More generally, Burdett and Judd (1983) observe that the key driver behind price dispersion in nearly all models (and likely in practice as well) are differences in consumers' ex post information concerning the prices charged by firms. This will be true in our setting as well. In our model, as in Diamond (1987), differences in consumers' willingness to pay for a product help to drive dispersed prices. Some consumers will be more keen to search for cheap or new products than others. Albrecht and Axell (1984) exploit the same type of heterogeneity to support wage dispersion in their model of search unemployment. In contrast to these models,

[6] We only cite a representative sample from this large literature. Stiglitz (1989) and Baye, Morgan, and Scholten (2006) provide surveys. See also McCall and McCall (2008, chapter 10).

our setting allows firms to further specialize in producing different products. This cross-cutting dimension is critical for many of the beneficial welfare conclusions that we identify.

Though anecdotal evidence for price dispersion abounds, systematic documentation across a variety of goods is less common. Recently, Kaplan and Menzio (2014) have documented the patterns of price dispersion across a wide range of consumer products. Their data allow for a decomposition of variance in prices and they find that search frictions account for 35 to 55 percent of the price differences observed across identical goods. They attribute the bulk of the remainder to inter-temporal price discrimination. Our analysis focuses on search frictions.

Product Differentiation.—Wolinsky (1986) examines how product differentiation motivates consumers to search. His framework represents a commonly encountered approach integrating search and product differentiation.[7] In his model, a consumer must search to learn both the price and his own idiosyncratic valuation for each firm's particular product. In our setting, by contrast, each consumer has a persistent type and firms producing objectively similar products are viewed similarly. Hence, we depart from the standard model of product differentiation in a monopolistically competitive market. Though we allow for multiple goods, our environment also differs from a "multi-product" search setting in which a consumer buys multiple products, perhaps from different suppliers (Burdett and Malueg 1981; McAfee 1995; Gatti 2000; Rhodes 2014; Zhou 2012). In our model, a consumer demands but one item; we relax this assumption in Section V.

Search Externalities and Complementarities.—As noted above, our key economic conclusions stem from the presence of search externalities and complementarities. Several earlier studies hint at similar implications to those that we identify. In line with our analysis, Pratt, Wise, and Zeckhauser (1979) empirically document considerable price dispersion across a range of standardized products.[8] They also suggest a precursor to the market-creation/expansion implications of search behavior that we examine. Their example of expensive and inexpensive motels along a road could find a welcome home in this paper. Nevertheless, the environment analyzed here differs from their setting and accommodates conclusions outside of their model. To highlight but one difference, we show how an equilibrium with search can increase firms' expected profits even above those with a monopoly price for a single good. This outcome cannot occur in Pratt, Wise, and Zeckhauser (1979).

Formally, many of our conclusions rest on a strong feed-back effect between search activity and production profitability. Diamond (1982) has argued for the presence of such an effect based on his analysis of a barter economy. His model assumes that goods are exchanged on a one-for-one basis and therefore does not include price dispersion. In an important paper, Burdett and Judd (1983) rely on

[7] Wolinsky (1986) builds on Perloff and Salop (1985). See also Wolinsky (1984). Anderson and Renault (1999, 2000); Armstrong, Vickers, and Zhou (2009); and Larson (2013) are more recent studies in this vein.

[8] This observation continues to be salient despite the advent of new technologies that have, arguably, lowered search costs (Baye and Morgan 2001; Baye, Morgan, and Scholten 2006).

the same complementarity between consumers' search activity and firms' pricing strategies to support multiple equilibria in a search model. They also show that multiple price-dispersion equilibria can coexist with different levels of consumer search intensity. In one equilibrium, a small fraction of consumers learn the prices posted by more than one firm. In another, a large fraction of consumers sample more than one firm.[9] In their model, firms' profits decline relative to the monopoly-price case whenever more consumers search to learn the posted price of more than one firm. This effect is to be expected and is present in some cases of our model as well. However, as we explain in Section IVB, firms' profits can also rise if they are also able to specialize in producing different goods. With product differentiation, search activity also facilitates the sorting of consumers among producers. Thus, firms can compensate for their reduced market power with better targeted pricing.[10]

Cachon, Terwiesch, and Xu (2008) identify a "market-expansion effect" tied to search activity. They model an oligopolistic market where firms can expand product lines. Unlike our model, Cachon, Terwiesch, and Xu's (2008) firms offer a continuum of firm-differentiated products. Each consumer has an idiosyncratic taste for each good and must search to learn the price and appeal of each good. The fine-grained idiosyncrasies in consumer tastes coupled with their cost specification produces a unique equilibrium without price dispersion. In contrast, our model generates high-, moderate-, and no-search equilibria, often with dispersed prices.

Market-expansion effects are also noted in the literature on agglomeration economies. Its central idea is that concentrations of sellers attract increased numbers of consumers by reducing consumers' search costs (Dudey 1990). New York's Diamond District or the London Silver Vaults are prime examples. The micro-level strategic interaction and competitive incentives in our model are unrelated to those stemming from agglomeration effects. Our consumers are not drawn to a particular locale.

While our model considers a nonmonetary economy, some features of our environment have been considered in recent search-theoretic models involving money, following Kiyotaki and Wright (1989). Both Camera, Reed, and Waller (2003) and Fong and Szentes (2005), for example, allow agents to endogenously specialize their production across an array of differentiated products. Respectively, they show how money can increase specialization in production or increase the production of costly, but higher quality goods. A key component in our framework is the strategic complementarity among producers of a similar product, which requires a costly investment to produce. As noted by Lester, Postlewaite, and Wright (2012), a similar strategic complementarity exists for costly information acquisition when assets trade as currency: as more agents recognize an asset as a safe medium of exchange, it becomes more worthwhile for others to do the same. Finally, our main conclusions concerning welfare in equilibria with intensive search and high price dispersion play

[9] Burdett and Judd (1983) study a model with non-sequential search and a model with noisy search. We are referring to their model with nonsequential search where the fraction of consumers learning one price is endogenously determined in equilibrium.

[10] Acemoglu and Shimer (2000) identify a similar phenomenon when studying labor markets. They argue that better search-related sorting promotes "technology dispersion." Their dispersion and the differentiation that we study are roughly analogous.

off the counterbalancing of both intensive- and extensive-margin effects. In a monetary model, such effects have been studied by Rocheteau and Wright (2005), though they rely on a considerably different market mechanism than we do.

Externalities of various sorts have always been a theme in the search literature. Effects such as the thick-market externality, the congestion externality, and the sorting externality have been identified and studied (Shimer and Smith 2001; Burdett and Coles 1997, 1999). We relate our analysis to these findings in the conclusion, after our model's details and implications have been spelled out.

II. The Model

Our model features a large number of firms and consumers. Each firm chooses a good to produce and its price. Consumers differ in their willingness to pay for goods. Each consumer searches sequentially among the firms to satisfy his consumption needs. First, we describe the two sides of the market. Subsequently, we introduce the market mechanism and search process. We comment on our model's interpretation before concluding this section.

Firms.—Firms move first in our economy. Each firm $j \in \mathbb{Z}$ commits to a single contract—a good-price pair, $\sigma_j = \langle x_j, p_j \rangle$ —that each passing consumer may either accept or reject; $x_j \in \{0, 1\}$ is the good produced by firm j and p_j is its per unit price. For simplicity, assume that good 0 can be produced at zero marginal cost and with no fixed cost. Good 1 can be produced with a constant marginal cost $c \geq 0$; moreover, if a firm produces good 1 it must additionally incur a fixed cost of $\phi > 0$. This fixed cost, for example, may represent an investment in a product-specific technology. Thus, a taxi driver may invest in a fancier car or a restaurant may opt for a more stunning decor.[11]

A firm's profit depends on the number of consumers who accept its contract offer. This number will be determined in equilibrium, but for now we observe that if n consumers accept $\sigma_j = \langle x_j, p_j \rangle$, then firm j's profit is $\pi(n \mid \sigma_j) = n p_j - x_j \cdot (nc + \phi)$.[12]

Consumers.—Each consumer $i \in \mathbb{Z}$ is willing to pay $z > 0$ for good 0. Consumers differ in their willingness to pay for good 1. This variation is described by the random variable V_i, which we call the consumer's type. A helpful interpretation is to consider good 0 to be an established, standardized product and good 1 to be a novel product with more heterogenous demand. For simplicity, suppose V_i takes on one of two values, $\underline{v}$ or $\overline{v}$, where $0 \leq c < \underline{v} < \overline{v}$. Consumers' types are independently and identically distributed and $\gamma := \Pr[V_i = \overline{v}]$. This distribution is common knowledge.

Goods 0 and 1 are substitutes. A consumer wishes to buy at most one unit of either good, but not both. He needs but one cab ride, or one dinner. (Two cab rides

[11] Fishman and Levy (2012) also allow firms to make a costly investment in product quality, thus generating product differentiation. In their model, the outcome of this investment is stochastic.

[12] Implicitly, we also allow a firm to forego production entirely. We do not model this decision directly as it is a weakly dominated action. We thank an anonymous referee for suggesting this notational simplification.

would take him away from his destination. Two dinners would give him indigestion.) Thus, if a type-v_i consumer accepts $\sigma_j = \langle x_j, p_j \rangle$ —he buys one unit of good x_j at price p_j from firm j—his immediate consumption payoff is $u(\sigma_j \,|\, v_i) = x_j v_i + (1 - x_j)z - p_j$. A consumer's purchase is voluntary. He is free to exit the market for an immediate payoff of zero.

The Market Mechanism.—Although a consumer knows the aggregate distribution of contracts in the economy, he does not know the specific contract offers of all firms. We assume that a consumer initially knows only the offer of one local firm; he must search to learn the offers of others. Our market unfolds as follows:

(1) Each firm j simultaneously commits to a single feasible contract, $\sigma_j \in \Sigma$. The set of feasible contracts is $\Sigma = \{\langle x, p \rangle : x \in \{0, 1\}, p \in \mathbb{R}_+\}$.

(2) Without cost, each consumer i learns the contract offer of firm $j = i$. Thus, at the start of the market consumer i has access to contract σ_i.

(3) For $t = 0, 1, \ldots,$ consumer i may (a) decide to accept contract σ_{i+t}, (b) search for a better option, or (c) exit the market.

  (a) If the consumer accepts σ_{i+t}, he receives an immediate payoff of $u(\sigma_{i+t} \,|\, v_i)$ and he exits the market. The firm offering the chosen contract supplies the good and receives its quoted price as payment.

  (b) If the consumer decides to search, he incurs a search cost of $s > 0$ and he learns the contract offer of firm $i + t + 1$. Step 3 subsequently repeats with σ_{i+t+1} being the available contract offer.

  (c) If the consumer exits the market, he receives a payoff of zero and does not return.

REMARK 1: *If consumer i accepts σ_j after sampling t (non-local) firms, his total payoff is $u(\sigma_j \,|\, v_i) - ts$. If he exits the market without making a purchase, this total payoff is $-ts$.*

REMARK 2: *Our model assumes that consumers cannot return to a previously sampled firm. That is, search is "without recall." This assumption accords well with our motivating taxicab example. Once a taxi drives by, it is usually gone for good. Anticipating our equilibrium analysis, we emphasize that our argument also holds when search allows for recall* (*DeGroot* 1970).[13] *For example, when consumers search for a new cellphone plan, they can reconsider previously encountered options*

[13] If search allows for recall, at each date t consumer i will be able to pick his favorite option from the contracts he has encountered on his search $\{\sigma_i, \sigma_{i+1}, \ldots, \sigma_{i+t}\}$ or to continue searching. Except for the extension in Section VC to increasing search costs, the analysis is unchanged if search with recall is allowed. A unified exposition motivated our modeling choice.

without much hassle. In such applications, search costs often stem from the challenge of evaluating a product's qualities or fine print specifics rather than physically "finding" the good.[14]

REMARK 3: *Each consumer knows the offer of a "local" firm. This is equivalent to giving the consumer one free sample before possibly embarking on a more extensive search. This is a common assumption (Salop and Stiglitz* 1982*; Stahl* 1989*; Acemoglu and Shimer* 2000, *among others*). *Unlike studies following Varian* (1980), *we do not posit the existence of a group of consumers who are perfectly informed about all firms' offers.*

Strategies and Equilibrium.—In many models generating price dispersion, firms adopt mixed strategies. A similar conclusion applies here, and we denote a mixed strategy for firm j as $\psi_j \in \Delta(\Sigma)$.[15]

When firm j posts a contract, the number of consumers who eventually accept it is a random variable, N_j. The distribution of N_j will depend on σ_j, on other firms' strategies, ψ_{-j}, and on consumers' search strategies. Suppressing these latter elements in our notation, firm j's expected profit is $\Pi_j(\sigma_j) := E\left[\pi(N_j \mid \sigma_j) \mid \sigma_j, \psi_{-j}\right]$.

A consumer's strategy is a search-and-purchase plan. In each period, he must decide whether to accept an available offer or to search further so as to maximize his expected utility net of search costs. His beliefs regarding the distribution of contracts in the economy are a critical input informing this plan. We represent these beliefs by $\tilde{\psi}_i \in \Delta(\Sigma)$. When a consumer searches, we assume that he views each encounter with a new firm as an independent draw from $\tilde{\psi}_i$.

With the above ideas in place, we can define an equilibrium in our market. We restrict our analysis to symmetric equilibria and we will often drop the "symmetric" qualifier. The definition conforms to the usual requirements. Firms' strategies constitute a Nash equilibrium; each consumer's search behavior maximizes his expected welfare given his beliefs; and each consumer's beliefs are correct.

DEFINITION 1: *A* (*symmetric*) *equilibrium is a tuple of production strategies, beliefs, and search strategies such that:*

(*i*) *Each firm adopts a common strategy* $\psi_j = \psi^*$, *which maximizes its expected profit given the strategies of all other firms and consumers' search strategies and beliefs.*

(*ii*) *Each consumer believes that the distribution of contracts in the economy is* $\tilde{\psi}_i = \psi^*$.

(*iii*) *The search strategy of each consumer maximizes his expected utility given* $\tilde{\psi}_i$.

[14] Ellison and Wolitzky (2012) discuss how firms may manipulate consumers' costs of learning product attributes and prices.

[15] Equivalently, we can reframe our model so that firms follow pure strategies and equilibrium will be described by the fraction of firms offering specific contracts. For example, Pratt, Wise, and Zeckhauser (1979) take this route.

(*iv*) *All consumers* (*of a particular type*) *adopt the same search strategy*.

Context and Interpretations.—Our model's simplicity accommodates several interpretations and links easily with a range of models of market competition. First, we assume that consumers acquire market information sequentially. Thus, we depart from a thoroughly "classic" model where perfect information reigns supreme. If consumers enjoyed perfect information (and could immediately contract with their preferred provider), then our model reduces to a case of Bertrand competition—each consumer will rush to the firm promising the greatest surplus.

Second, our model features an orderly search process. Given our assumptions about consumer beliefs, our model behaves as if a consumer samples from a large set of ex ante identical firms, as is standard in the search literature. Our search process considerably simplifies the analysis of consumer sorting. It lets us easily track the expected inflow of local and non-local consumers to each firm.

Third, we are silent regarding the model's time dimension. While the "t" superscript suggests the passing of physical time and we often speak of "periods of search," we stake no claim on the model's temporal calibration. We view "t" as an index describing the extent of search undertaken by a consumer. Hence, our economy can unfold asynchronously with different consumers searching at different paces or different times. The key restriction is that firms cannot screen consumers based on how much previous search they have undertaken or on their arrival "time." All firms post a single price, which is a common practice.

Finally, we assume an equal number of consumers and firms. Our analysis is qualitatively unchanged if each firm is initially matched with m consumers each of whom may accept its offer or engage in search. The $m = 1$ case simplifies exposition.

III. Consumer Search Behavior

How should a typical consumer search? If he believes that all firms are offering contracts according to $\tilde{\psi} \in \Delta(\Sigma)$, his optimal search strategy involves a cutoff rule (see Lemma A.1 in the online Appendix). For each type of consumer v_i, there exists a critical value $u^*(v_i)$, such that if $u(\sigma_j | v_i) \geq \max\{u^*(v_i), 0\}$, then he stops searching and accepts σ_j. If $u^*(v_i) > \max\{u(\sigma_j | v_i), 0\}$, the consumer takes a pass on σ_j and samples his next available option. Otherwise, he exits the market. Following DeGroot (1970), the cutoff value characterizing the search strategy is the unique solution to the equation $u^*(v_i) = E_{\tilde{\psi}}[\max\{u^*(v_i), u(\sigma | v_i), 0\}] - s$, where σ is distributed according to $\tilde{\psi}$.

This strategy has two noteworthy characteristics. First, it is stationary and independent of the extent of prior search activity. This characteristic is a direct consequence of the constant search costs and does not carry over to a more general setting (see Section VC). Second, a consumer's beliefs $\tilde{\psi}$ are constant and do not adjust in response to the observed sample of contracts. Though in equilibrium a consumer's beliefs are correct, he does not draw any inferences from "off-equilibrium path" occurrences. Being surprised by an unexpectedly good deal will not prime the consumer to anticipate even better offers elsewhere. Both characteristics considerably simplify the analysis and are standard features of nearly all search models.

Since consumer behavior is straightforward, we will typically identify equilibria by referring only to firms' strategies. When we do so, our description implicitly assumes that consumers hold correct beliefs, and that their search behavior is described by the above cutoff strategy.

IV. Market Equilibrium

We begin our analysis by introducing a maintained parameter restriction.

ASSUMPTION A-1: $\underline{v} - c - \phi < (1 - \gamma) z$ *and* $\gamma(\overline{v} - c) - \phi < z$.

When Assumption A-1 holds, producing good 1 in an economy without consumer search is comparatively unprofitable. The assumption tilts our setting against the conclusions that we are working toward and focuses our attention to (in our opinion) interesting cases.

The case with neither search, nor product differentiation, nor price dispersion will serve as our benchmark. This equilibrium, which always exists, reproduces the well-known Diamond (1971) paradox. A single product is sold at its monopoly price. As all consumers know this to be true, search never pays. As Burdett and Judd (1983) show, this result is common across search models whenever consumers learn the price posted by a single firm.

THEOREM 1: *There exists a unique equilibrium with no product differentiation. In this equilibrium all firms offer good* 0 *at the monopoly price* z.

From a consumer's perspective, Theorem 1 describes this economy's "bad equilibrium." Since consumers expect all firms to offer the same contract, they have no incentive to search. Acting like a local monopolist, each firm extracts the surplus from its captive market and earns a profit of z. If search costs are sufficiently high or there are too few consumers who value good 1 highly, this will be the economy's only equilibrium.

As noted above, the equilibrium described in Theorem 1 is neither inevitable nor particularly common. Goods usually differ both physically and in price. Similar products vary by quality and the same product can vary in price from store to store. Both forms of variation suggest alternative equilibrium arrangements that may dominate the benchmark case. More precisely, we say that an equilibrium features *product differentiation* if there is a positive probability that two firms produce different goods. Similarly, *price dispersion* exists if there is a positive probability that two firms set different prices for the same good. Our analysis below considers progressively more complex equilibria that differ on these two characteristics.

A coordination failure lurks behind Theorem 1. Given Assumption A-1, the sustainable production of good 1 requires that some consumers make their purchase from a nonlocal supplier. This is only possible after a costly search. We say that an equilibrium exhibits *search* if there is a positive probability that at least one consumer searches for at least one period. For search to be worthwhile, a consumer needs to be reasonably certain that it will yield a worthwhile prize, and that it will not be exceptionally costly. The first criterion means that some consumers must anticipate receiving a positive surplus from some contract(s) in the economy. The second criterion means that an adequate number of firms must offer sufficiently attractive contracts.

Once multiple goods are traded simultaneously, a large set of equilibrium contracts may emerge. This, however, does not happen and there is only a small set of relevant cases.

THEOREM 2: *Consider an equilibrium with product differentiation. There are at most two distinct contracts offered by firms producing good* 0, *and there are at most two distinct contracts offered by firms producing good* 1.

The intuition behind Theorem 2 is straightforward and has been noted before by Albrecht and Axell (1984) and Diamond (1987), among others. Each posted equilibrium price must correspond to the reservation price of some consumer. Suppose good 1 trades in equilibrium at, say, three prices: $\underline{p} < p' < \bar{p}$. Suppose that the cheapest contract is acceptable to both type-$\underline{v}$ and type-$\bar{v}$ consumers, and the most expensive contract is acceptable only to type-$\bar{v}$ consumers. But now, some firm has a profitable deviation. If the price p' is acceptable to both types of consumers, then any firm charging $\underline{p}$ can raise its price without hurting its sales volume. Similarly, if p' is accepted only by type-$\bar{v}$ consumers, a firm charging p' can boost its profits by raising its price. While a full argument must consider additional scenarios, the reasoning is similar.

Aspects of Theorem 2 are similar to those identified by Curtis and Wright (2004), who propose a "law of two prices" within a monetary search model. The need to simultaneously satisfy a family of incentive constraints supporting search behavior limits firms' freedom to set prices in equilibrium. Despite this similarity, as we show below, our analysis leads to somewhat different conclusions concerning the incidence of price dispersion equilibria than those proposed by Curtis and Wright (2004). For instance, they show that equilibria with multiple prices occur for a wider range of parameters than equilibria with one price. Example 1, presented below, suggests the converse conclusion in our model.

Theorem 2 allows for a considerable simplification of our analysis, as there are at most four classes of equilibria we need to consider. One class exhibits no price dispersion. Two goods are available and each good sells at a uniform price. The other three classes feature varying degrees of price dispersion. We focus on the two extreme cases in this quartet: product differentiation and no price dispersion, and product differentiation with both goods selling at multiple prices. We comment briefly on the remaining cases—one good trading at multiple prices—in this section's conclusion. Neither of these cases generate equilibria that Pareto-dominate the benchmark scenario. Thus, they are comparatively less interesting than the other cases considered.

A. *Product Differentiation without Price Dispersion*

Starting with the first case, we begin with a lemma pinning down the equilibrium prices. Implicitly, it also defines the equilibrium search pattern, as explained below.

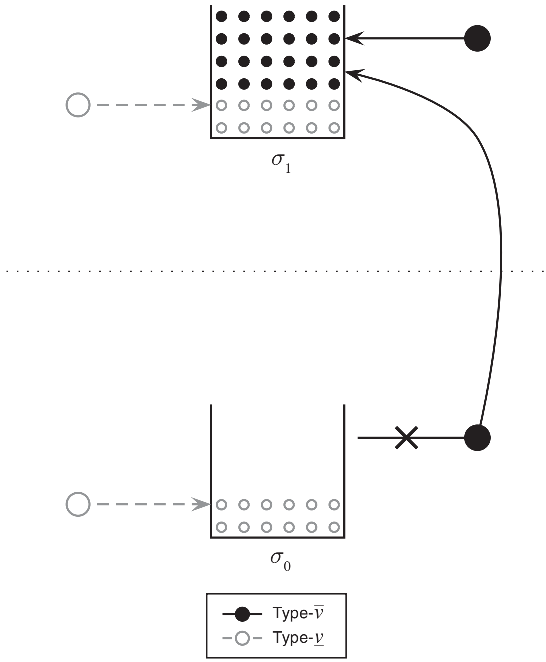

FIGURE 1

Notes: The flow of consumers in an equilibrium with product differentiation but no price dispersion. Type-$\underline{v}$ consumers accept either contract, while type-$\bar{v}$ consumers search for σ_1.

LEMMA 1: *Let $\sigma_0 = \langle 0, p_0 \rangle$ and $\sigma_1 = \langle 1, p_1 \rangle$ be the two contracts offered in an equilibrium with product differentiation and no price dispersion. Then $p_0 = z$ and $p_1 = \underline{v}$.*

Two immediate implications flow from Lemma 1. First, type-$\underline{v}$ consumers will never search in an equilibrium absent price dispersion. All offered contracts guarantee them a payoff of zero. Second, given Assumption A-1, a firm producing good 1 must in expectation sell its product to more than one consumer. Hence, type-$\bar{v}$ consumers must search and accept σ_1. This is a simple case of consumers self-sorting between products, a recurrent theme in both our model and real-world markets. The implied pattern of search and the flow of consumers is presented in Figure 1. After matching with a firm offering $\sigma_0 = \langle 0, p_0 \rangle$, a type-$\bar{v}$ consumer is unsatisfied; he searches until he finds a firm offering $\sigma_1 = \langle 1, p_1 \rangle$. This flow is depicted by the solid arrows leading to $\sigma_1 = \langle 1, p_1 \rangle$. In contrast, a type-$\underline{v}$ consumer finds both σ_0 and σ_1 acceptable. We can summarize the resulting equilibrium as follows.

THEOREM 3: *Let*

$$(1) \qquad \lambda_1^* = \frac{\gamma(\underline{v} - c)}{\phi - (1 - \gamma)(\underline{v} - c - z)}.$$

There exists an equilibrium with product differentiation and no price dispersion if and only if

$$s \le \lambda_1^* \left(\bar{v} - z - c - \frac{\phi}{1 - \gamma + \gamma/\lambda_1^*} \right) \tag{2}$$

and

$$s \le \frac{(\lambda_1^*)^2 (1 - \gamma)(\underline{v} - c)}{\gamma}. \tag{3}$$

In such an equilibrium, a typical firm offers the contract $\sigma_1 = \langle 1, \underline{v} \rangle$ *with probability* λ_1^* *and* $\sigma_0 = \langle 0, z \rangle$ *with probability* $\lambda_0^* = 1 - \lambda_1^*$. *On the equilibrium path, a type-*$\bar{v}$ *consumer searches until he finds a firm offering the contract* σ_1, *which he accepts. A type-*$\underline{v}$ *consumer does not search, and accepts the first contract available.*

As illustrated in Example 1, (2), (3), and Assumption A-1 can be simultaneously satisfied for a wide range of parameters. By inspection, we observe that (2) and (3) hold when search costs are sufficiently low ($s \to 0$). Conversely, if there are too few type-$\bar{v}$ agents ($\gamma \to 0$) such an equilibrium will not exist.[16]

The equilibrium in Theorem 3 sorts consumers between goods. However, this sorting is incomplete as good 1 is consumed by a mixed group of consumers. This outcome is crucial lest the market unravel. If a firm charges a price slightly more than $\underline{v}$ for good 1, its demand will fall. Hence, producers of good 1 face a sharp incentive to maintain prices at a (relatively) low level. Coincidentally, that same low price motivates type-$\bar{v}$ consumers to search and it is their inflow that makes producing good 1 a profitable undertaking.

Several intriguing comparative static implications follow easily from Theorem 3. First, equilibrium prices are independent of search costs and are driven by consumers' valuations. Second, the equilibrium distribution of contracts does not vary with search cost. Instead, the propensity of a firm to produce good 1 increases ($\lambda_1^* \uparrow$) as the fixed costs of production falls ($\phi \downarrow$), or as consumers' preferences shift toward good 1 ($\gamma \uparrow$).

Since the firm-consumer ratio is one, a natural welfare criterion is the aggregate expected surplus generated by a firm-consumer pair. In the equilibrium without product differentiation, this value was z and the firm appropriated it all. When there is product differentiation, each firm's profit declines to $(1 - \gamma) z$; type-$\bar{v}$ consumers enjoy a positive surplus; but, type-$\underline{v}$ consumers experience no welfare improvement. Aggregate welfare may increase or decrease depending on the model's specifics, even when search costs are small.

COROLLARY 1: *Suppose search costs become small, i.e.,* $s \to 0$. *The equilibrium with product differentiation and no price dispersion generates greater aggregate*

[16] λ_1^* is increasing in γ. Hence, as $\gamma \to 0$, then $\lambda_1^* \to 0$. Condition (2) is eventually violated.

expected surplus than the equilibrium with no product differentiation if and only if $\bar{v} \geq z + \underline{v}$.

B. *Product Differentiation and Price Dispersion*

The preceding subsection showed that when there is product differentiation but no price dispersion, some consumers benefit, others receive their reservation payoff, and firms' profits decline. The aggregate outcome might be better or worse than had the local monopoly situation prevailed. These conclusions lack universal appeal. Ideally, we would like to identify cases where all consumers and all producers are better off relative to the benchmark case, at least in an ex ante sense. Whether such a Pareto-improvement is a realistic aspiration or a Pollyannish fantasy is not obvious a priori.

Two facts hint at a possibility. First, searching consumers gain access to a wider set of producers allowing price comparisons. This trims firms' market power. Thus, ensuring that consumers gain from active search does not seem too demanding. Second, consumer search has a cross-cutting implication for firm profits. Search generates a movement of consumers and also produces a sorting of consumer types. The former effect increases sales volumes while the latter allows for more targeted pricing. Both effects counteract firms' weakened market power and will prove critical in supporting a Pareto-superior equilibrium.

The first challenge is to ensure that consumers search actively in equilibrium. To justify costly search there must exist equilibrium contracts that consumers view as particularly appealing. Promising carrots in equilibrium, however, poses a challenge for producers. The problem here is that firms have the incentive to free-ride whenever consumers search. We have in mind the following phenomenon. When consumers search for a prized contract, that contract must promise a sufficiently high reward and be offered by sufficiently many firms in order to justify the search costs incurred. When many firms offer the appealing contract, the consumer is confident his search will be worthwhile. Since search costs are incurred incrementally, a firm promising a very appealing contract has an incentive to slightly downgrade its offer. A searching consumer will still stop and buy rather than search further. The "deviating" downgraded contract does not affect a consumer's beliefs but does trap him like flypaper before he finds the anticipated deal. This points to a subtle and under-appreciated paradox present in search economies. A firm benefits from having a sufficient number of direct competitors since that induces consumers to search for its type of product. It therefore wants competitors to coordinate, but the firm itself seeks to miscoordinate with its competition in order to profit. The implied tension for an equilibrium is obvious.

To add some formalism, from Theorem 2 we know that at most four contracts—say σ_0, σ_1, $\hat{\sigma}_0$, and $\hat{\sigma}_1$—can coexist in equilibrium. For notation, we let $\sigma_k = \langle k, p_k \rangle$ and $\hat{\sigma}_k = \langle k, \hat{p}_k \rangle$, with the convention that $p_k < \hat{p}_k$. A typical firm follows a mixed strategy offering contract σ_k $(\hat{\sigma}_k)$ with probability λ_k $(\hat{\lambda}_k)$. With notation set, we can consider our first key lemma. This lemma accomplishes two tasks. First, it identifies constraints neutralizing the free-riding deviation described above. Second, coupled with Lemma 3, it defines the equilibrium pattern of consumer search.

LEMMA 2: *Consider an equilibrium with product differentiation*:

(*i*) *If good* 0 *sells at prices* p_0 *and* $\hat{p}_0$, $p_0 < \hat{p}_0$, *then* $z - \hat{p}_0 = \max\{u^*(\underline{v}), 0\}$ *and* $z - p_0 = u^*(\overline{v})$.

(*ii*) *If good* 1 *sells at prices* p_1 *and* $\hat{p}_1$, $p_1 < \hat{p}_1$, *then* $\underline{v} - p_1 = \max\{u^*(\underline{v}), 0\}$ *and* $\overline{v} - \hat{p}_1 = u^*(\overline{v})$.

We can paraphrase Lemma 2 as follows. If a particular good's prices are dispersed, the prices at which the good trades must correspond to the agents' reservation values. These reservation values are determined in equilibrium and account for both the goods' intrinsic utility and the search-and-purchase strategy adopted by the agents. Generally, of course, these reservation values will differ for different types of agents. An implication of the four equalities in Lemma 2 is that in equilibrium no firm can increase its posted price even slightly without alienating some consumers.[17] This serves as a disciplinary force holding the equilibrium together. The same intuition is present in a variety of search models (Albrecht and Axell 1984; Diamond 1987).

As noted above, we are particularly interested in equilibria with desirable welfare implications for both producers and consumers. This added requirement necessarily leads us to consider cases where both goods sell at multiple prices.

LEMMA 3: *In any equilibrium that Pareto-dominates the equilibrium of Theorem* 1:

(*i*) *All goods must trade at multiple prices.*

(*ii*) *On the equilibrium path, a type-*$\overline{v}$ *agent searches when* $\hat{\sigma}_0$ *is available to him; else, he accepts* σ_0, σ_1, *or* $\hat{\sigma}_1$. *A type-*$\underline{v}$ *agent searches when* $\hat{\sigma}_1$ *is available to him; else, he accepts* σ_0, $\hat{\sigma}_0$, *or* σ_1.

Lemma 3 summarizes a complex pattern of consumer search, illustrated in Figure 2. First, from Lemma 2 we see that two contracts endogenously emerge as the "deals" motivating search. The preferred option of a type-$\overline{v}$ consumer is σ_1 while a type-$\underline{v}$ consumer views σ_0 as best. Consumers hunt for their preferred contract, but may settle for a more expensive option or the other good should they stumble upon it.

With the pattern of search determined, the equilibrium's remaining necessary features briskly fall into place. Lemmas 2 and 3 allow us to characterize a typical consumer's (cutoff) search strategy and the quantitative relationship among equilibrium prices.

[17] Intuitively, we might say that consumers feel a smidgen of indignation and refuse to purchase any good whose price even marginally exceeds expectations.

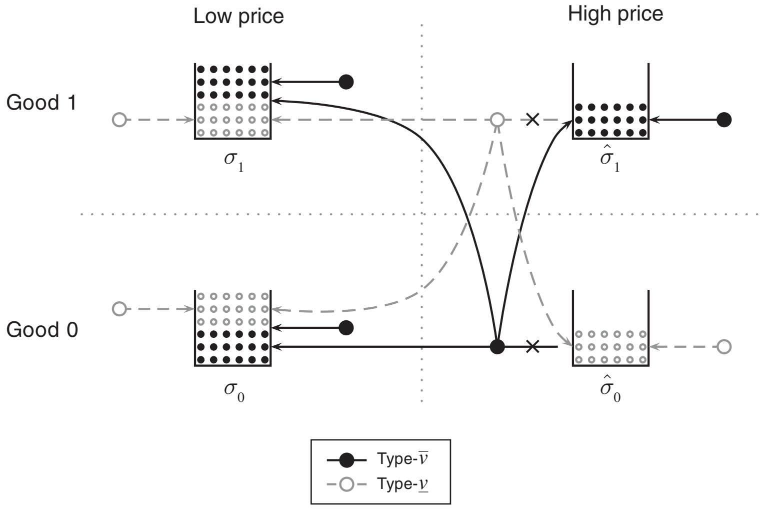

FIGURE 2

Note: The flow of consumers in an equilibrium with product differentiation where both goods trade at multiple prices.

LEMMA 4: *Consider an equilibrium that Pareto-dominates the equilibrium of Theorem* 1:

(*i*) *The cutoff values characterizing a consumer's search strategy are* $u^*(\overline{v}) = \overline{v} - p_1 - s/\lambda_1$ *and* $u^*(\underline{v}) = z - p_0 - s/\lambda_0$.

(*ii*) *The prices associated with the contracts satisfy*

$$\hat{p}_0 = p_0 + \frac{s}{\lambda_0}, \qquad p_0 = p_1 + \frac{s}{\lambda_1} + z - \overline{v},$$

$$\hat{p}_1 = p_1 + \frac{s}{\lambda_1}, \qquad p_1 = p_0 + \frac{s}{\lambda_0} + \underline{v} - z.$$

Noting the relationship among prices, an immediate corollary is that

$$\overline{v} - \underline{v} = \frac{s}{\lambda_0} + \frac{s}{\lambda_1}. \tag{4}$$

Expression (4) carries two economic implications. First, it can be satisfied only if search costs are sufficiently small. Second, when search costs rise, the prevalence of at least one inexpensive contract must also rise. This response is counterintuitive since higher search costs suggest that firms' market power should be enhanced, which favors a shift toward higher priced goods. In equilibrium, however, the observed response is necessary to maintain consumers' incentive to search.

The search pattern also determines the profits associated with each contract. The low-price contracts generate high volume sales while catering to a diverse group of consumers. High-priced contracts feature targeted prices and a homogenous clientele.

LEMMA 5: *Consider an equilibrium that Pareto-dominates the equilibrium of Theorem* 1. *The expected profits associated with each contract are*:

$$\Pi_j(\sigma_1) = \left(\frac{1-\gamma}{1-\hat{\lambda}_1} + \frac{\gamma}{1-\hat{\lambda}_0}\right)(p_1 - c) - \phi, \quad \Pi_j(\hat{\sigma}_1) = \frac{\gamma}{1-\hat{\lambda}_0}(\hat{p}_1 - c) - \phi,$$

$$\Pi_j(\sigma_0) = \left(\frac{1-\gamma}{1-\hat{\lambda}_1} + \frac{\gamma}{1-\hat{\lambda}_0}\right)p_0, \qquad \Pi_j(\hat{\sigma}_0) = \frac{1-\gamma}{1-\hat{\lambda}_1}\hat{p}_0.$$

In an equilibrium, of course, the expected profits must be the same across all offered contracts. Solving $\Pi_j(\sigma_0) = \Pi_j(\hat{\sigma}_0), \Pi_j(\sigma_1) = \Pi_j(\hat{\sigma}_1)$, and using Lemma 4 allows us to express equilibrium prices as a function of the distribution of contracts. That is,

$$(5) \qquad p_0 = \frac{1-\hat{\lambda}_0}{1-\hat{\lambda}_1}\cdot\frac{1-\gamma}{\gamma}\cdot\frac{s}{\lambda_0} \quad \text{and} \quad p_1 = c + \frac{1-\hat{\lambda}_1}{1-\hat{\lambda}_0}\cdot\frac{\gamma}{1-\gamma}\cdot\frac{s}{\lambda_1}.$$

Thus, we can characterize the equilibrium entirely as a distribution of contracts—$\lambda_0, \lambda_1, \hat{\lambda}_0, \hat{\lambda}_1$ —satisfying (4), (5), and the following four conditions:

$$(6) \qquad \Pi_j(\sigma_0) = \Pi_j(\sigma_1),$$

$$(7) \qquad p_1 = p_0 + \frac{s}{\lambda_0} + \underline{v} - z,$$

$$(8) \qquad 1 = \lambda_0 + \lambda_1 + \hat{\lambda}_0 + \hat{\lambda}_1,$$

and

$$(9) \qquad z - p_0 - \frac{s}{\lambda_0} \geq 0.$$

Expression (6) ensures the equality in profits; (7) ensures the remaining relationship among equilibrium prices is satisfied; (8) ensures the λs define a valid probability distribution; and, (9) ensures type-$\underline{v}$ agents wish to search, i.e., $u^*(\underline{v}) \geq 0$. All remaining constraints have either been eliminated or incorporated into the above expressions.

The preceding discussion confirms that conditions (4)–(9) are necessary for the equilibrium we wish to construct. They encapsulate and simplify two sets of requirements. First, the equilibrium distribution of contracts must ensure that the expected profits of each firm are the same taking as given consumers' search behavior. This requirement simultaneously restricts the frequency with which different contracts are offered and the prices associated with each good. Second, the equilibrium contract distribution must ensure that consumers' payoffs respect the constraints in

Lemma 2, thereby ensuring the required pattern of search is consistent with equilibrium prices. Complementary activity of the two sides of the market is therefore critical for the equilibrium to be sustained.

Satisfying conditions (4)–(9) is also sufficient to sustain an equilibrium where both goods trade at multiple prices (Lemma A.15 in the online Appendix). Whereas the complexity of the required conditions suggests that equilibria conforming to our desiderata are rare, this is not the case. The following example offers some guidance when we may expect such equilibria to emerge and when different types of equilibria coexist.

EXAMPLE 1: *Consider an economy where* $\overline{v} = 4$, $\underline{v} = z = 1$, $\phi = 1/3$, *and* $c = 0$. *Depending on the magnitude of the search cost and the distribution of agents' preferences, this economy may feature one equilibrium or multiple equilibria. Figure* 3 *summarizes the present equilibria as a function of* (γ, s).[18] *The one-product equilibrium of Theorem* 1 *always exists. It is the unique equilibrium when* γ *is sufficiently low. Given our parameterization, a typical firm's profit is equal to* 1 *in this equilibrium.*

An equilibrium with product differentiation, but no price dispersion (*Theorem* 3), *is common once* γ *is sufficiently large* (*the gray region*). *In this equilibrium, the profits of a typical firm are* $1 - \gamma$, *which is less than in the one-product benchmark scenario.*

If γ *is sufficiently large and search costs are sufficiently small, equilibria with both product differentiation and price dispersion occur as well* (*the hatched region*). *Quite often such cases lead to profits exceeding the benchmark case. For instance, when* $s = 0.03$ *and* $\gamma = 0.3$, *a typical firm's expected profit is approximately* 1.18. *This final case, and others like it in the lower right corner of Figure* 3, *Pareto-dominates the benchmark scenario as consumers also enjoy a positive surplus as well.*

The welfare implications associated with Example 1 are its key conclusion. The incentives created by the economy's multiple goods and dispersed prices ensure that consumers search and sort among producers. While the resulting reshuffling of consumers is less than perfect (some consumers do not purchase the item that they value the most), it sufficiently improves matters to ensure that good 1 producers face adequate demand to render their production profitable. Surprisingly, producers of good 0 also gain.

Intermediate Price Dispersion.—We have thus far focused on two extreme equilibria with product differentiation. Either uniform prices prevail or all goods trade at multiple prices. Intermediate cases are possible as well, though such equilibria do not enjoy the appealing welfare and profit implications of the preceding case (see Lemma 3 above). For example, it is straightforward to construct an equilibrium where good 0 trades at price $p_0 = z$ and good 1 trades at prices $p_1 = \underline{v} < \hat{p}_1$. A type-$\underline{v}$ consumer never engages in search and a type-$\overline{v}$ consumer searches

[18] To construct the figure, we solved numerically for an equilibrium contract distribution satisfying conditions (4)–(9) over a fine grid of the parameter space.

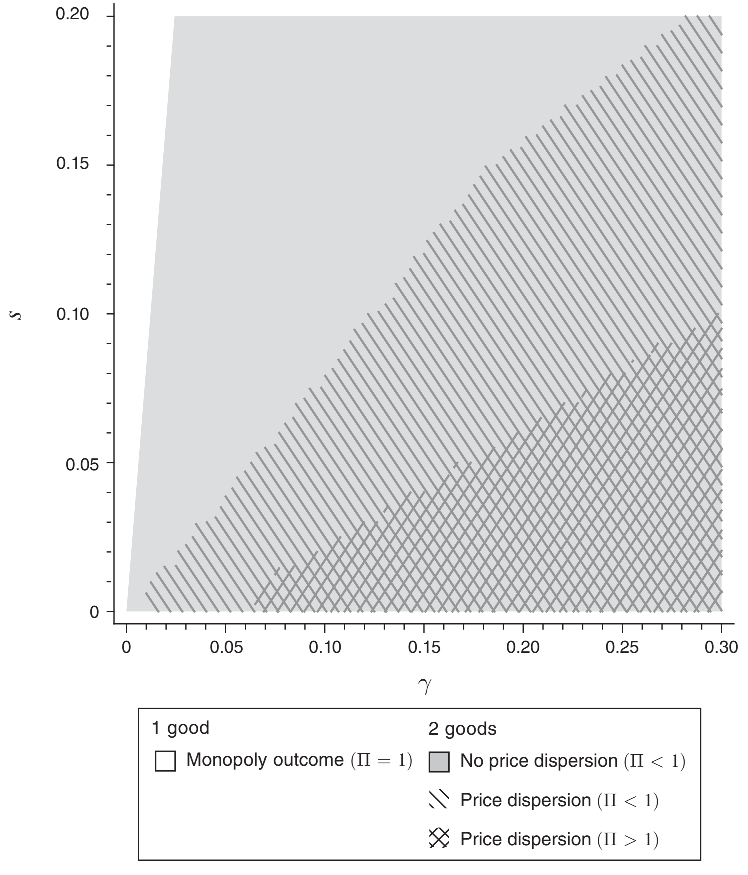

FIGURE 3. MULTIPLE EQUILIBRIA WITH PRODUCT DIFFERENTIATION AND PRICE DISPERSION

for good 1, which he buys at both high and low prices. As in the equilibrium of Theorem 3, firms' profits are $(1-\gamma)z$.[19] Similar logic lets us construct an equilibrium where good 0 trades at multiple prices as well.

V. Interpretations and Extensions

To focus on important economic effects, we kept our model spare and simple. Naturally, building on this streamlined formulation allows for many extensions and alternative interpretations. First, several technical amendments are possible. Our model extends to the case of multiple consumer types and multiple goods.[20] The

[19] Such equilibria are comparatively tractable. Notably: $\hat{p}_1 = \underline{v} + s/\lambda_1$, $\lambda_0 = \dfrac{c+\phi+(1-\gamma)z-\underline{v}}{(1-\gamma)c+\phi+(1-\gamma)(z-\underline{v})}$, $\lambda_1 = \dfrac{\gamma s}{(\phi+(1-\gamma)z)\lambda_0 - \gamma(\underline{v}-c)}$, and $\hat{\lambda}_1 = 1-\lambda_0-\lambda_1$. Given the parameters of Example 1, it can be shown that an equilibrium in this class exists whenever there exists an equilibrium with product differentiation and no price dispersion (the gray region in Figure 3). Examples exist showing that this inclusion is not true in general.

[20] In a preliminary version of our analysis, we considered a continuum of consumer types. In such a model, there is additionally an endogenous cutoff consumer type delineating who seeks out good 1 versus good 0.

importance of imperfect or incomplete sorting of consumers as a mechanism that ties an equilibrium together continues to apply.

Second, with appropriate parameter restrictions, our model applies to a variety of problems. Markets with vertically differentiated products, like our luxury/standard taxicab example, require the restrictions $z \leq \underline{v} < \overline{v}$ or $\underline{v} < \overline{v} \leq z$.[21] When offered at the same price, either good 1 or good 0 is universally preferred. To study a market with horizontally differentiated products, such as restaurants, it is sufficient to adopt the convention that $\underline{v} < z < \overline{v}$. Now consumers disagree about which good is superior. Some prefer French cuisine while others view Japanese food as more appetizing.

Our model's ability to accommodate both vertical and horizontal differentiation spans the ways in which new products surface. For example, Bower and Christensen (1995) and Christensen (1997) argue that new goods frequently enter at lower quality tiers than existing products. Nevertheless, with a careful pricing strategy targeting the appropriate market segment such new goods can be very profitable (Christensen 1997). The emergence of large discount retailers and of no-frills airlines represent prime examples. Innovation along a horizontal dimension rests on identifying products tailored to specific groups of consumers. Many innovations in the technological realm, such as the file-sharing services mentioned earlier, produce markets where multiple competitors offer modestly different capabilities that appeal differentially to alternative consumers. Our model accommodates all of these situations.

Third, labor markets are a prime example where search is paramount. A relabeling transposes our model to a labor-market context where firms search to recruit skilled workers.[22] In this case a worker may choose whether to invest ϕ to upgrade his human capital or skill. Skilled workers supply higher quality labor (good 1) while unskilled workers supply lower quality labor (good 0). Type-$\overline{v}$ employers are more eager to hire skilled workers than type-$\underline{v}$ employers. Firms search for employees and incur a per candidate interview cost, s. The economy's wage distribution, i.e., the distribution of prices, is the equilibrium outcome of firms and workers interacting in this labor market with workers choosing whether to acquire human capital.

Finally, many economically motivated extensions are also easily accessible. We briefly outline some of these below. For brevity and clarity, in each extension we confine our attention to equilibria with product differentiation but no price dispersion (as in Theorem 3), and we suppress many technical caveats. More complex equilibria can also be constructed with parallel implications to our preceding analysis.

A. *Orphaned Consumers*

We have assumed that each consumer knows the contract offer of a local producer. In practice, however, a local producer may not exist. For example, upon realizing the need to take a taxi, a consumer may observe that no taxis are around. To

[21] Wildenbeest (2011) develops a search model, along with an empirical analysis, focusing on vertically differentiated products.

[22] As documented by van Ommeren and Russo (2014), sequential search by firms appears to be an important recruitment practice. However, they qualify this conclusion noting that it depends on the common role of formal and informal search methods in the market under study.

accommodate this situation, suppose that after firms have made their production decisions, but before consumers learn any contract offers, with probability $1 - q$ each firm experiences an independent shock that prevents it from selling any goods. Consequently, an orphaned consumer's initial set of available contracts is empty and, occasionally, a searching consumer may "find" a firm that has gone out of business. Otherwise, the model remains unchanged.

The main conclusions of our original analysis continue to apply. In an equilibrium with product differentiation but no price dispersion, firms would offer contracts $\sigma_0 = \langle 0, z \rangle$ and $\sigma_1 = \langle 1, \underline{v} \rangle$. Of course, the shock impacts profitability; accordingly, the equilibrium contract distribution adjusts. The probability that a firm offers good 1 now becomes $\lambda_1^* = \frac{\gamma(\underline{v} - c)}{\phi - (1-\gamma)q(\underline{v} - c - z)}$. Also, and in common with the remaining extensions below, the sufficient and necessary conditions for this equilibrium to exist—analogous to those from Theorem 3—need to be amended appropriately. (We omit these calculations for brevity.) Formalities notwithstanding, our model's qualitative behavior remains unchanged.

B. *Multi-unit Purchases*

Consider the problem faced by a visitor looking for lodging in an unfamiliar city. If the traveler intends a long stay, he has a greater incentive to invest in an extensive search for the right accommodations. The returns to a thorough search are far less for a one-time, one-night visit. Within a broad range of prices, almost any hotel will do.

We address this situation by comparing one- and multi-unit purchasers. Some consumers buy only one unit of a good while others acquire $\tau > 1$ units. The fraction γ of consumers are multi-unit purchasers. We continue to assume that good 0 has uniform appeal. Each consumer's willingness to pay is z. Preferences for good 1 are heterogenous. Multi-unit purchasers are willing to pay $\bar{v}$ per unit. Its value to one-time buyers is $\underline{v}$.

This formalization relates simply to our hotel example. Good 0 is a standard hotel located in the city center that caters to short-term business travelers. Good 1 is a long-term hotel that offers amenities that are more valuable to long-term guests (for example, kitchenettes), but which are not important for one-night visitors.[23]

As usual, there exists an equilibrium where only $\langle 0, z \rangle$ is provided. However, if type-$\bar{v}$ consumers search, providers of good 1 will spring up. Again consider an equilibrium with product differentiation but no price dispersion. Even though a fraction of consumers wishes to purchase multiple copies of a good, a moment of reflection suggests that the usual two contracts will be on offer: $\sigma_0 = \langle 0, z \rangle$ and $\sigma_1 = \langle 1, \underline{v} \rangle$. As before, the equilibrium distribution of contracts must adjust. Now, $\lambda_1^* = \frac{\gamma\tau(\underline{v} - c)}{\phi - (1-\gamma)(\underline{v} - c - z)}$, which generalizes our original result.

[23] As $\bar{v} < z$ is possible, the long-term hotel may be located a couple of bus stops away from the city center. Given the costs of understanding the bus system, it may be objectively inferior to the standard hotel for one-night stayers.

C. *Increasing Search Costs*

Although a standard assumption in the literature, a consumer's search costs are rarely constant in practice. Rather, search costs typically rise over time. The increasing marginal cost of time, and the limiting case of deadlines, provide reasons why this is so. A more subtle reason search costs drift upward is that a person's search technology deteriorates as he substitutes toward more costly (or less effective) search methods once easy options get exhausted.[24]

In Section B of the online Appendix, we generalize our model by assuming that search costs rise over time. Under suitable conditions, we show that there exists an equilibrium paralleling our results from Theorem 3. Firms offer $\sigma_0 = \langle 0, z \rangle$ or $\sigma_1 = \langle 1, \underline{v} \rangle$ and a type-$\bar{v}$ consumer searches for σ_1. However, a searching consumer settles for σ_0 if his search fails to yield a positive result sufficiently quickly. As above, type-$\underline{v}$ consumers forgo search entirely.

A qualitatively similar behavioral pattern emerges when a searching consumer has a constant search cost but is learning about the distribution of contracts as he goes along.[25] The associated intuition is simple. A consumer embarks on his quest with an air of optimism believing his desired product is likely around the corner. Since he is uncertain regarding the actual distribution of contracts, each unfavorable draw dulls his confidence. Eventually pessimism takes over and he settles for what is available.

D. *Adverse Selection and Insurance*

As a final extension, consider the case of adverse selection.[26] As a specific example, consider health insurance.[27] Individuals who are willing to pay more for health coverage tend to be those who consume the most medical resources. Hence, insuring them is more costly. Insurance markets are also complex and thoughtfully searching among alternative policies is often difficult. Even in cases where the direct costs of search are low, such as Medigap insurance, where policies are standardized and quotes are posted on the internet, the perception of high search costs likely deters many consumers from seriously investigating their options.[28] Search frictions facilitate higher insurance premiums by reinforcing insurers' market power (Cebul et al. 2011). Samuelson and Zeckhauser (1988) document the high persistence in health insurance plan choice among Harvard employees. Search costs, of course, provide one explanation for this observation.

Consider an insurance market and assume that both types of consumers have a willingness to pay of z for a "standard" policy (good 0). The "premium" policy

[24] Osberg (1993) recognizes the substitution among search methods in an empirical study of job search.

[25] See McCall (1970) or Rothschild (1974) for formal analyses.

[26] Recent studies of search economies with adverse selection include Inderst (2005); Guerrieri, Shimer, and Wright (2010); and Lauermann and Wolinsky (2013). Related studies focusing on insurance markets include Mathewson (1983), Schlesinger and von der Schulenburg (1991), and Seog (2002).

[27] Cutler and Zeckhauser (2000) summarize this market's key features. Our model does not address moral hazard which, along with adverse selection, is an important caveat in studying insurance markets.

[28] Lin and Wildenbeest (2013) and Kim (2014) both find the presence of substantial search frictions in the market for Medigap insurance.

(good 1) offers enhanced coverage and elicits a willingness to pay of $\underline{v}$ and $\bar{v}$ from the two types of consumers. The fixed cost of supplying the premium policy (ϕ) can be interpreted as an extra administrative cost, or the cost of establishing a network of higher quality health care providers.

To accommodate adverse selection, we depart from our baseline model in two ways. The first departure is substantive. Suppose that the cost of serving a customer depends on the customer's type. Type-$\bar{v}$ consumers are at high health risk and the marginal cost of providing them coverage is c_0 when they enroll in a standard policy and $c_1 > c_0$ when they enroll in a premium policy. Type-$\underline{v}$ consumers have a low health risk. The cost of insuring a type-$\underline{v}$ consumer is normalized to zero, regardless of the policy that they purchase. Ex ante, an insurer cannot distinguish between the two types of consumers.

The second departure is technical. For the sake of argument suppose that

$$\text{(A-1}'\text{)} \quad \underline{v} - \gamma(c_1 - c_0) - \phi < (1-\gamma)z \quad \text{and} \quad \gamma(\bar{v} - c_1 + c_0) - \phi < z.$$

Note that (A-1′) gently modifies Assumption A-1 on account of the novel cost structure. Furthermore, to focus on a case where adverse selection impedes the market's operation, suppose $z < c_0 < c_1 < \underline{v} < \bar{v}$.

An equilibrium exists where no one searches. The no-search equilibrium may take one of two forms. First, if

$$\pi_0 = \gamma(z - c_0) + (1-\gamma)z < 0,$$

then insurance provision is not profitable and no coverage will be available. In this case, adverse selection effectively destroys the market. The second possibility is less ominous. It occurs when $\pi_0 = \gamma(z - c_0) + (1-\gamma)z \geq 0$. Now in equilibrium all firms offer the basic policy $\sigma_0 = \langle 0, z\rangle$ and all consumers accept it immediately. Since $z < c_0$, the presence of high-risk consumers reduces the profitability of a standard policy. In both cases, despite the presence of many consumers who could be well-served by an enhanced policy, no firm will offer it.

Consider the case where consumers search for insurance. Now two types of insurance contracts will be offered: $\sigma_0 = \langle 0, z\rangle$ and $\sigma_1 = \langle 1, \underline{v}\rangle$. Type-$\bar{v}$ consumers will search to find an enhanced policy while type-$\underline{v}$ consumers will default into either option without bothering to search. The consequences of this behavior are surprisingly beneficial. A firm's equilibrium profit increases to $(1-\gamma)z$ and consumers gain as well. Type-$\underline{v}$ consumers are no worse off, while type-$\bar{v}$ consumers now enjoy access to their preferred product. Search behavior and the costly self-sorting of consumers has led to a Pareto improvement. In this regard, coordinating on an outcome where it is common practice for consumers to "shop around," even if the direct costs of search remain relatively high, can prove beneficial.

Given the contrast between the no-search and search equilibria, it is instructive to investigate policies that can nudge a market toward the Pareto-superior outcome. As in the original model, several conditions must be satisfied. First, type-$\bar{v}$ consumers must consider search worthwhile. A type-$\bar{v}$ consumer begins to search when the fraction of firms offering σ_1, λ_1, is large enough. More formally, $u^*(\overline{v}) = \overline{v} - \underline{v} - s/\lambda_1 \geq 0$, or equivalently

$$(10) \qquad \lambda_1 \geq \overline{\lambda}_1 \equiv \frac{s}{\overline{v} - \underline{v}}.$$

When the fraction of firms offering σ_1 exceeds $\overline{\lambda}_1$, it is sufficiently easy for type-$\overline{v}$ consumers to find their desired contract. Second, it is straightforward to show that when $\sigma_0 = \langle 0, z \rangle$ and $\sigma_1 = \langle 1, \underline{v} \rangle$ are the available policies and type-$\overline{v}$ consumers search for σ_1, fraction

$$(11) \qquad \lambda_1^* = \frac{\gamma(\underline{v} - c_1)}{\phi - (1 - \gamma)(\underline{v} - z)}$$

of firms must offer σ_1 in equilibrium. Expression (11) is the direct analogue of (1) from Section IVA. It ensures that σ_0 and σ_1 are equally profitable for the firms given the pattern of consumer search.[29] Together, (10) and (11) imply that

$$(12) \qquad \lambda_1^* \geq \overline{\lambda}_1$$

is a necessary condition for an equilibrium with product differentiation. Finally, in direct parallel to (2) in Theorem 3, search costs cannot be too high, i.e.,

$$(13) \qquad s \leq \frac{(\lambda_1^*)^2(1 - \gamma)\underline{v}}{\gamma}.$$

If they were, a premium-policy firm would be tempted to slightly raise its price while still catering to high-risk consumers who happen to find its offer first.

The policy implications of (11), (12), and (13) are straightforward. First, if no consumers are engaged in search, then policies that only reduce costs per search (s) fail to generate benefits. Consumers have no incentive to utilize a superior or cheaper search technology if there is nothing to find in the first place. By contrast, a policy that encourages the supply of σ_1 to rise has the potential to be more effective in practice. For example, a lump-sum subsidy of κ paid to any firm that offers $\sigma_1 = \langle 1, \underline{v} \rangle$ can help offset its fixed cost. The required magnitude of the subsidy depends on the market's default state. If in the absence of consumer search σ_0 is a profitable contract, i.e., $\pi_0 \geq 0$, a subsidy of $\kappa = \pi_0 - \pi_1$, where

$$\pi_1 = \gamma(\underline{v} - c_1) + (1 - \gamma)\underline{v} - \phi,$$

is sufficient to ensure that a firm is indifferent between offering σ_0 or σ_1.[30] Crucially, the subsidy need only be taken up by fraction $\overline{\lambda}_1$ of firms. Once that threshold is met, consumers become confident that their search activity will be rewarded, the market

[29] Using (A-1′) it can be shown that $0 < \lambda_1^* < 1$.

[30] (A-1′) ensures that $\kappa = \pi_0 - \pi_1 \geq 0$. The firm is indifferent between σ_0 and σ_1 conditional on no consumer engaging in search. In expectations, this is the most pessimistic case from a firm's perspective.

for σ_1 thickens rapidly, and σ_1 becomes profitable on its own merits. Similarly, if $\pi_0 < 0$ and no insurance is available, subsidizing fraction $\overline{\lambda}_1$ of firms to offer the premium σ_1 policy would allow the market to surface, giving all consumers a choice between deluxe *and* basic insurance policies. Thus, the targeted support of one new product can encourage close substitutes to ride into the market on its coattails. In this case, a subsidy of $\kappa = \pi_1$ would be sufficient.[31] A simple cost-benefit calculation can determine if and when and if these policies are worthwhile. Generally, their effectiveness is greater whenever search costs are relatively small and the main obstacle to overcome is the coordination problem between consumers and firms.[32]

VI. Concluding Remarks

Our simple model highlights a market-creation externality that emerges when consumers search. Search leads to both product differentiation and heightened competition. The effects we emphasized operate through the enhanced sorting that costly search behavior necessarily spawns. On many occasions, equilibria with active search produce greater welfare and higher profits. This conclusion persists despite the deadweight of search costs (which consumers must bear), and despite firms' curtailed market power (which dampens their profits).

Before concluding, we would like to explain how our model contrasts with previous analyses that emphasize search externalities or that employ similar terminology. Shimer and Smith (2001) discuss both the thick-market and congestion externalities tied to search behavior. Our setting revolves around distinct, albeit related, effects.

First, regarding thick markets, we have argued that search supports market expansion and creation. However, search behavior alone cannot foster this expansion. Rather, it requires market entry by producers as well as consumer search. By coordinating around an outcome with superior sorting, local markets for particular goods thicken. On the one hand, this effect ensures that it is safe for producers to offer new products. On the other hand, consumers become sufficiently confident in their prospects for a successful search to incur the associated costs.

A congestion externality, as described by Shimer and Smith (2001), is not present in our environment. In our setting, a firm can match with an unlimited number of consumers, and thus consumers do not elbow out one another. That said, our environment exhibits a related pecuniary externality. As more firms produce a specific good, they steal some consumers from their direct competitors. This effect exists in any market, so is to be expected.

[31] This subsidy level ensures that a firm is indifferent between offering σ_1 or no policy conditional on no consumer searching.

[32] Assuming the direct policy implementation costs are negligible, the subsidy is worthwhile from a cost-benefit perspective if the gain in consumer welfare and firm profits exceeds the subsidy. For example, when $\pi_0 > 0$, then the proposed subsidy is worthwhile if

$$\underbrace{\gamma\left(\overline{v} + z - s(1 + 1/\lambda_1^*)\right)}_{[1]} + \underbrace{\gamma(c_0 - z)}_{[2]} \geq \underbrace{\overline{\lambda}_1 \kappa}_{[3]}. \tag{14}$$

Term [1] is the per capita change in ex ante consumer welfare. Term [2] is the per firm change in firm profits. And term [3] is the per firm subsidy cost. Since $\overline{\lambda}_1 = s/(\overline{v} - \underline{v})$, (14) will be satisfied when s is sufficiently small. A similar analysis applies to the $\pi_0 < 0$ case.

Our conclusions concerning consumer sorting are distinct from the sorting externalities studied by Burdett and Coles (1997). Derived in the context of a dynamic "marriage market" setting, their sorting externality refers to the change in the distribution of available partners as agents match and subsequently exit the market (Burdett and Coles 1999). These specific effects play no role in our setting. In practical terms, we allow producers to practice polygamy. Thus, consumers experience no direct effect from the search activity of other agents. The indirect effects can be substantial, as we have explained above.

Our model was kept spare to plainly exhibit the effects of interest. Many extensions beyond those proposed here may be pursued. For example, our study focused on the case of competitive firms. Allowing for oligopolistic firms or a monopolist, who can internalize some of the highlighted externalities, would complement our model.[33] Our conclusions should carry over to these alternative market structures.

Similarly, we have abstracted from the role of advertising in our analysis (Robert and Stahl 1993).[34] A natural way to introduce advertising into our model is to allow consumers to direct their search toward producers offering products they like, without wasting effort visiting suppliers of the other good.[35] Thus, a consumer's search would focus on a subset of firms, but he still incurs a cost to learn the details of a firm's offer. In this model, the added information helps consumers, but it may fail to bring benefits to firms absent a robust search effort among all consumers. For example, if *only* type-$\overline{v}$ consumers search for good 1 and type-$\overline{v}$ consumers find both goods acceptable, as in the equilibrium of Section IVA, all firms' equilibrium profits remain $(1-\gamma)z$. Though consumers benefit in comparison to the single-good, monopoly price outcome, firms' profits end up lower. Other forms of advertising, particularly those that serve as a coordination device on the search-intensive equilibrium, may help sustain more appealing market-wide outcomes.[36]

A natural interpretation of our model concerns comparisons across different markets, with some markets exhibiting search-intensive equilibria, and others characterized by buyers' ready acceptance of sellers. Alternatively, the model can also be viewed as a description of the same market operating at different moments in time. The sale of existing homes offers an illustration. While homes are bought and sold throughout the year, it is well known that existing home sales cluster in the spring and summer. Patient sellers are often advised not to list their house in the fall or winter, since few people are searching for homes then. (Stale houses tend to lose their value, and it is also somewhat costly in terms of convenience to have one's house on the market.) In the spring many people are searching, so houses are listed on the market in vastly disproportionate numbers. Throughout the year this market moves

[33] A monopolist could produce many versions of the same product or offer multiple goods.

[34] Bagwell (2007) offers a review of models of advertising in economic analysis, including search economies.

[35] In this sense search becomes partially directed (Menzio 2007; Goldberg 2007). For example, consumers have a list of all firms offering good 0 and good 1 and can use this list to guide their store visits.

[36] There is evidence that certain markets may display the qualities of a "search-intensive" equilibrium. For example, Couture (2012) documents the distances traveled in search for restaurant meals in US cities. Consumers often forgo close options in favor of establishments further afield.

between a period of collectively focused search among many types of buyers and a period when search is far less intense. Despite the fact that disproportionately many other sellers are listing their homes at the same time, a patient seller often gains by waiting for the search-intensive period in the spring to sell. An extension examining the transition among equilibria would naturally capture this dynamic. As suggested by our brief analysis in Section VD, such transitions would undoubtably feature a critical threshold of buyers and sellers beyond which the market would tip—in the sense of Schelling (1969, 1971)—toward an equilibrium with a significantly higher search intensity.

The economic effect we have emphasized—a search-generated and coordination-supported market enhancement externality—is central in many markets. Capitalizing on this effect should be a policy goal. While reducing search costs directly is certainly beneficial, as a policy it is unlikely to be as effective on the margin in markets where few consumers undertake active search. The fundamental coordination problem between consumers and producers may remain unresolved. Even if search costs are consequential, the welfare gains from a search-intensive equilibrium can be substantial as new markets and new business opportunities arise due to this search activity. The recent emergence of ride-share and taxi-like services like Uber and Lyft provides but one example where enhanced search (in this case due to the enormous capabilities of the Internet) enables a multi-product equilibrium to supplant one with a single product. In many taxi markets, these platforms have added a dash of price dispersion combined with a dollop of product differentiation to motivate the required change in consumer behavior. Curiously, our analysis suggests that traditional taxi services may ultimately *benefit* from these developments if consumers reduce their use of personal vehicles sufficiently and substitute more cars for hire—sometimes taxis, sometimes car services. Better integrating this observation into applied studies in industrial organization and market design would be a valuable next step. A successful market thrives on the symbiotic relationship between producers and consumers. Absent a coordination mechanism beyond personal incentives, appropriate complementary actions must emerge and be sustained on the two sides of the market.

REFERENCES

Acemoglu, Daron, and Robert Shimer. 2000. “Wage and Technology Dispersion.” *Review of Economic Studies* 67 (4): 585–607.

Albrecht, James W., and Bo Axell. 1984. “An Equilibrium Model of Search Unemployment.” *Journal of Political Economy* 92 (5): 824–40.

Anderson, Simon P., and Regis Renault. 1999. “Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlain-Diamond Model.” *RAND Journal of Economics* 30 (4): 719–35.

Anderson, Simon P., and Régis Renault. 2000. “Consumer Information and Firm Pricing: Negative Externalities from Improved Information.” *International Economic Review* 41 (3): 721–42.

Armstrong, Mark, John Vickers, and Jidong Zhou. 2009. “Prominence and consumer search.” *RAND Journal of Economics* 40 (2): 209–33.

Bagwell, Kyle. 2007. “The Economic Analysis of Advertising.” In *Handbook of Industrial Organization,* Vol. 3, edited by Mark Armstrong and Robert H. Porter, 1701–1844. Amsterdam: North-Holland.

Baye, Michael R., and John Morgan. 2001. “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets.” *American Economic Review* 91 (3): 454–74.

Baye, Michael R., John Morgan, and Patrick Scholten. 2006. "Information, Search, and Price Dispersion." In *Economics and Information Systems*. Vol. 1, *Handbooks in Information Systems*, edited by Terrence Hendershott, 323–75. Bingley, UK: Emerald.

Borenstein, Severin, and Nancy L. Rose. 1994. "Competition and Price Dispersion in the U.S. Airline Industry." *Journal of Political Economy* 102 (4): 653–83.

Bower, Joseph L., and Clayton M. Christensen. 1995. "Disruptive Technologies: Catching the Wave." *Harvard Business Review* 73 (1): 43–53.

Burdett, Ken, and Melvyn G. Coles. 1997. "Marriage and Class." *Quarterly Journal of Economics* 112 (1): 141–68.

Burdett, Kenneth, and Melvyn G. Coles. 1999. "Long-Term Partnership Formation: Marriage and Employment." *Economic Journal* 109 (456): 307–34.

Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–69.

Burdett, Kenneth, and David A. Malueg. 1981. "The Theory of Search for Several Goods." *Journal of Economic Theory* 24 (3): 362–76.

Cachon, Gérard P., Christian Terwiesch, and Yi Xu. 2008. "On the Effects of Consumer Search and Firm Entry in a Multiproduct Competitive Market." *Marketing Science* 27 (3): 461–73.

Camera, Gabriele, Robert R. Reed, and Christopher J. Waller. 2003. "Jack of All Trades or a Master of One? Specialization, Trade, and Money." *International Economic Review* 44 (4): 1275–94.

Cebul, Randall D., James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba. 2011. "Unhealthy Insurance Markets: Search Frictions and the Cost and Quality of Health Insurance." *American Economic Review* 101 (5): 1842–71.

Christensen, Clayton M. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail.* Boston: Harvard Business Review Press.

Couture, Victor. 2012. "Valuing the Consumption Benefits of Urban Density." http://www.hec.ca/iea/seminaires/130111_victor_couture.pdf.

Curtis, Elisabeth, and Randall Wright. 2004. "Price setting, price dispersion, and the value of money: Or, the law of two prices." *Journal of Monetary Economics* 51 (8): 1599–1621.

Cutler, David M., and Richard J. Zeckhauser. 2000. "The Anatomy of Health Insurance." In *Handbook of Health Economics*, Vol. 1A, edited by Anthony J. Culyer and Joseph P. Newhouse, 563–643. Amsterdam: North-Holland.

Cutler, David M., and Richard Zeckhauser. 2004. "Extending the Theory to Meet the Practice of Insurance." *Brookings-Wharton Papers on Financial Services* 2004: 1–53.

DeGroot, Morris H. 1970. *Optimal Statistical Decisions.* Hoboken, NJ: John Wiley and Sons.

Diamond, Peter A. 1971. "A Model of Price Adjustment." *Journal of Economic Theory* 3 (2): 156–68.

Diamond, Peter A. 1982. "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy* 90 (5): 881–94.

Diamond, Peter. 1987. "Consumer Differences and Prices in a Search Model." *Quarterly Journal of Economics* 102 (2): 429–36.

Dudey, Marc. 1990. "Competition by Choice: The Effect of Consumer Search on Firm Location Decisions." *American Economic Review* 80 (5): 1092–1104.

Ellison, Glenn, and Alexander Wolitzky. 2012. "A search cost model of obfuscation." *RAND Journal of Economics* 43 (3): 417–41.

Fishman, Arthur, and Nadav Levy. 2012. "Search Costs and Risky Investment in Quality." https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2302603_code349193.pdf?abstractid=1752066&mirid=1.

Fong, Yuk-fai, and Balázs Szentes. 2005. "Compensation for Quality Difference in a Search Model of Money." *International Economic Review* 46 (3): 957–71.

Gatti, J. Rupert J. 2000. "Equilibrium Price Dispersion with Sequential Search." http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.198.7466&rep=rep1&type=pdf.

Gerardi, Kristopher S., and Adam Hale Shapiro. 2009. "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry." *Journal of Political Economy* 117 (1): 1–37.

Goldberg, Dror. 2007. "Money with partially directed search." *Journal of Monetary Economics* 54 (4): 979–93.

Guerrieri, Veronica, Robert Shimer, and Randall Wright. 2010. "Adverse Selection in Competitive Search Equilibrium." *Econometrica* 78 (6): 1823–62.

Inderst, Roman. 2005. "Matching Markets with Adverse Selection." *Journal of Economic Theory* 121 (2): 145–66.

Kaplan, Greg, and Guido Menzio. 2014. "The Morphology of Price Dispersion." National Bureau of Economic Research (NBER) Working Paper 19877.

Kim, You Suk. 2014. "Consumer Search Frictions, Competition and Adverse Selection in Health Insurance Markets: Evidence from Medigap." PhD diss. University of Pennsylvania.

Kiyotaki, Nobuhiro, and Randall Wright. 1989. "On Money as a Medium of Exchange." *Journal of Political Economy* 97 (4): 927–54.

Larson, Nathan. 2013. "Niche products, generic products, and consumers search." *Economic Theory* 52 (2): 793–832.

Lauermann, Stephan, and Asher Wolinsky. 2013. "Search and Adverse Selection." Unpublished.

Lester, Benjamin, Andrew Postlewaite, and Randall Wright. 2012. "Information, Liquidity, Asset Prices, and Monetary Policy." *Review of Economic Studies* 79 (3): 1209–38.

Lin, Haizhen, and Matthijs R. Wildenbeest. 2013. "Search and Prices in the Medigap Insurance Market." http://kelley.iu.edu/riharbau/RePEc/iuk/wpaper/bepp2013-15-lin-wildenbeest.pdf.

Mathewson, G. F. 1983. "Information, Search, and Price Variability of Individual Life Insurance Contracts." *Journal of Industrial Economics* 32 (2): 131–48.

McAfee, R. Preston. 1995. "Multiproduct Equilibrium Price Dispersion." *Journal of Economic Theory* 67 (1): 83–105.

McCall, J. J. 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics* 84 (1): 113–26.

McCall, Brian P., and John H. McCall. 2008. *The Economics of Search.* Oxford: Routledge.

Menzio, Guido. 2007. "A Theory of Partially Directed Search." *Journal of Political Economy* 115 (5): 748–69.

Osberg, Lars. 1993. "Fishing in Different Pools: Job-Search Strategies and Job-Finding Success in Canada in the Early 1980s." *Journal of Labor Economics* 11 (2): 348–86.

Perloff, Jeffrey M., and Steven C. Salop. 1985. "Equilibrium with Product Differentiation." *Review of Economic Studies* 52 (1): 107–20.

Pratt, John W., David A. Wise, and Richard Zeckhauser. 1979. "Price Differences in Almost Competitive Markets." *Quarterly Journal of Economics* 93 (2): 189–211.

Reinganum, Jennifer F. 1979. "A Simple Model of Equilibrium Price Dispersion." *Journal of Political Economy* 87 (4): 851–58.

Rhodes, Andrew. 2014. "Multiproduct Retailing." *Review of Economic Studies* 82 (1): 360–90.

Robert, Jacques, and Dale O. Stahl, II. 1993. "Informative Price Advertising in a Sequential Search Model." *Econometrica* 61 (3): 657–86.

Rocheteau, Guillaume, and Randall Wright. 2005. "Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium." *Econometrica* 73 (1): 175–202.

Rothschild, Michael. 1973. "Models of Market Organization with Imperfect Information: A Survey." *Journal of Political Economy* 81 (6): 1283–1308.

Rothschild, Michael. 1974. "Searching for the Lowest Price When the Distribution of Prices Is Unknown." *Journal of Political Economy* 82 (4): 689–711.

Salop, S., and J. E. Stiglitz. 1982. "The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents." *American Economic Review* 72 (5): 1121–30.

Samuelson, William, and Richard Zeckhauser. 1988. "Status quo bias in decision making." *Journal of Risk and Uncertainty* 1 (1): 7–59.

Schelling, Thomas C. 1969. "Models of Segregation." *American Economic Review* 59 (2): 488–93.

Schelling, Thomas C. 1971. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1 (2): 143–86.

Schlesinger, Harris, and J.-Matthias Graf von der Schulenburg. 1991. "Search Costs, Switching Costs and Product Heterogeneity in an Insurance Market." *Journal of Risk and Insurance* 58 (1): 109–19.

Sengupta, Anirban, and Steven N. Wiggins. 2014. "Airline Pricing, Price Dispersion, and Ticket Characteristics on and off the Internet." *American Economic Journal: Economic Policy* 6 (1): 272–307.

Seog, S. Hun. 2002. "Equilibrium Price Dispersion in the Insurance Market." *Journal of Risk and Insurance* 69 (4): 517–36.

Shimer, Robert, and Lones Smith. 2001. "Matching, Search, and Heterogeneity." *B. E. Journal of Macroeconomics* 1 (1): Article 6.

Stahl, Dale O, II. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review* 79 (4): 700–712.

Stahl, Dale O. 1996. "Oligopolistic pricing with heterogeneous consumer search." *International Journal of Industrial Organization* 14 (2): 243–68.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3): 213–25.

Stiglitz, Joseph E. 1989. "Imperfect information in the product market." In *Handbook of Industrial Organization*, Vol. 1, edited by Richard Schmalensee and Robert Willig, 769–847. Amsterdam: North-Holland.

van Ommeren, Jos, and Giovanni Russo. 2014. "Firm Recruitment Behaviour: Sequential or Non-sequential Search?" *Oxford Bulletin of Economics and Statistics* 76 (3): 432–55.

Varian, Hal R. 1980. "A Model of Sales." *American Economic Review* 70 (4): 651–59.

Wildenbeest, Matthijs R. 2011. "An empirical model of search with vertically differentiated products." *RAND Journal of Economics* 42 (4): 729–57.

Wolinsky, Asher. 1984. "Product Differentiation with Imperfect Information." *Review of Economic Studies* 51 (1): 53–61.

Wolinsky, Asher. 1986. "True Monopolistic Competition as a Result of Imperfect Information." *Quarterly Journal of Economics* 101 (3): 493–512.

Zhou, Jidong. 2012. "Multiproduct Search." http://economics.mit.edu/files/8755.