

Frictions in product markets[☆]

6

Alessandro Gavazza^{a,*} and Alessandro Lizzeri^b

^aLondon School of Economics, London, United Kingdom

^bPrinceton University, Princeton, NJ, United States

*Corresponding author: e-mail address: a.gavazza@lse.ac.uk

CONTENTS

1 Introduction	434
2 Transaction costs	435
2.1 Vertical differentiation	436
2.2 Market power and secondary markets	439
2.3 Empirical research	440
2.4 Role of intermediaries in overcoming transaction costs	443
3 Asymmetric information	444
3.1 Theory	444
3.1.1 <i>Static adverse selection, exogenous ownership</i>	444
3.1.2 <i>Dynamic trading with exogenous initial ownership</i>	445
3.1.3 <i>Endogenous initial ownership</i>	446
3.2 Empirics	448
3.3 Role of intermediaries in overcoming asymmetric information	449
4 Search frictions	450
4.1 Theoretical models	451
4.1.1 <i>Information clearinghouse</i>	451
4.1.2 <i>Simultaneous search</i>	453
4.1.3 <i>Sequential search</i>	454
4.1.4 <i>Sellers' non-price behavior</i>	456
4.1.5 <i>Product differentiation</i>	457
4.2 Empirical contributions	460
4.2.1 <i>Descriptive papers</i>	460
4.2.2 <i>Estimation of simultaneous-search models</i>	461
4.2.3 <i>Estimation of sequential-search models</i>	463
4.3 The role of intermediaries in search markets.....	466
4.3.1 <i>Theoretical contributions</i>	466
4.3.2 <i>Empirical contributions</i>	468
5 Matching frictions	470
5.1 The role of intermediaries in matching markets	475

[☆] We thank three anonymous reviewers for insightful suggestions. Alessandro Gavazza gratefully acknowledges support from the European Research Council (ERC-Consolidator grant award no.771004).

References.....	477
-----------------	-----

Abstract

We focus on markets with frictions, such as transaction costs, asymmetric information, search and matching frictions. We discuss how such frictions affect allocations, favor the emergence of intermediaries or dealers, and potentially create market power. Our focus is mostly on markets with many participants rather than on transactions that are bilateral or involve a small number of players.

Keywords

Transaction costs, Asymmetric information, Search frictions, Matching frictions, Intermediaries, Market power

1 Introduction

Frictions such as transaction costs feature prominently in Ronald Coase's examination of the nature of the firm (Coase, 1937). In Coase's view, organizations are structured so as to minimize transaction costs. This view is also fundamental in Coase's exploration of the role of property rights (Coase, 1960). The so-called Coase theorem characterizes an important feature of a benchmark frictionless environment in which property rights do not matter, but much of the discussion in Coase (1960) shows that the structure of property rights is essential when efficient transactions are impeded by frictions.¹ In a similar vein, Williamson (1989), following in Coase's footsteps, coined the term "Transaction cost economics" to denote the analysis of institutions via a contractual framework (see previous Handbook chapter by Williamson, 1989). Some frictions are essential for such institutions to play a role; a classic example of Williamson's (1985) is asset specificity.

This chapter studies several types of frictions that have been analyzed in more detail since the work of Coase and Williamson. Our focus is on market frictions rather than those within a firm. We discuss how allocations are affected by such frictions, but we will also show that these frictions can affect the structure of the market by favoring the emergence of intermediaries or dealers. Our focus is mostly on markets with many participants rather than on transactions that are bilateral or involve a small number of players, which are the focus in much of the literature following Williamson.

We begin our review by considering transaction costs, which can be viewed as a reduced-form friction that impedes exchange. It is sometimes convenient to be relatively agnostic about the form of these impediments to trade; many studies assume the presence of a generic transaction cost. In Section 2, we first introduce the most commonly used model of secondary market transactions and discuss how transaction costs affect allocations and incentives in secondary markets for durable goods.

¹ For a historical discussion of Coase (1960), see Medema (2020).

We emphasize durable goods because we discuss transactions among consumers and, unless the goods are durable, such transactions would not be possible.² We then investigate a more specific form of an impediment to trade—asymmetric information—in Section 3. We continue our focus on durable goods because other chapters in this Handbook examine asymmetric information in markets such as those for credit or insurance (Einav et al., 2021). We next consider markets with search and matching frictions in Sections 4 and 5, respectively.

Much of industrial organization studies the sources and consequences of market power, which often only exists if there are few firms. This chapter focuses mostly on frictions in environments in which there are many participants, but some of the frictions we consider generate endogenous market power or can exacerbate existing market power. For instance, search costs generate some pricing power even if there are many competing firms. Furthermore, in the presence of information asymmetries, sellers have unique access to knowledge of the quality of their own goods. Even if there are many other competing sellers, such a monopoly over information can create inefficiencies. Some of the frictions we discuss negatively affect allocations, even in purely competitive environments.

In each section, we begin by discussing how a specific friction affects allocations and consumer behavior. We then discuss the effects of these frictions on the incentives of producers or sellers to manipulate product design. We also discuss the incentives of such producers to manipulate such frictions. We also consider how the presence of these frictions gives scope for third parties, such as intermediaries, to play an important role in some markets. Two types of intermediaries often help consumers and/or sellers: In some markets, intermediaries acquire ownership of the goods (i.e., dealers or retailers); in other markets, they act as information intermediaries (i.e., price-comparison websites, brokers, or professional advisors). Although, as Spulber (1996, 1999) emphasizes, the emergence of market intermediaries is in response to underlying market frictions and their main role is to identify how to clear the market, the literature shows that the welfare effect of such intermediaries is not always positive.

There are other potentially important frictions that we do not cover. For instance, switching costs can have substantial effects on market power and allocations. The chapter by Farrell and Klemperer (2007) in Volume 3 of this Handbook surveys the literature on switching costs.

2 Transaction costs

In this section, we mostly consider the literature that has examined the consequences of transaction costs in secondary markets for durable goods. Transaction costs also

² Waldman (2003) also surveys durable goods markets. One area he discusses and we do not is time inconsistency for sellers. In addition, there is some overlap with our Section 3.

feature prominently in financial markets; in particular, by limiting arbitrage (Constantinides, 1986). Our focus here is on secondary markets for durable goods. Such markets play an important allocative role for some durable goods. For instance, in the U.S., the number of used-car transactions is approximately three times larger than the number of new-car transactions. Furthermore, the dispersion of used-car prices (measured by the coefficient of variation, calculated from the National Automobile Dealers Association Car Price Guide) is approximately five times larger than the dispersion of new-car prices, which suggests that secondary markets play an important role in broadening the spectrum of goods available to consumers. Several markets for used capital equipment are also very active. For example, more than half of the trucks traded in the U.S. in 1977 sold in secondary markets (Bond, 1983). Active markets also exist for used medical equipment, construction equipment, aircraft, and other goods. Since the mid-1980s, trade in the secondary market for commercial aircraft has grown steadily, and the number of transactions in used markets today is about three times the number of purchases of new aircraft (Gavazza, 2011a). Business aircraft also have active secondary markets (Gavazza, 2016).

The amount of activity in secondary markets varies dramatically across goods, with some markets extremely active (e.g., cars, aircraft) and others much less so (e.g., household appliances, computers). The amount of activity also varies substantially across different countries for the same goods. For instance, as highlighted by Gavazza et al. (2014), the U.S. used-car market is more active than the French one. What forces are responsible for these differences? What are the consequences for prices and allocations, producers' profits and consumers' welfare? How do these differences in activity affect the extent of variety available to consumers? Can some of the observed differences in the primary markets across goods and countries be due, in part, to the underlying causes of the differences in activity in the secondary markets?³

2.1 Vertical differentiation

Some degree of heterogeneity is essential to explain substantial volumes of trade in secondary markets. At the same time, some type of friction must be present in most of these secondary markets; otherwise, trade would likely be much higher than what is observed. Understanding the volume and patterns of such trade requires a model that takes into account the nature of the gains from trade. In some cases, the gains from trade are due to changes in preferences that could be deterministic or due to shocks. For instance, industrial machinery could be available for sale on the secondary market because of bankruptcy. In other cases, secondary market transactions are due to systematic forces related to some more primitive heterogeneity. For instance, in car

³ Empirical studies of secondary markets include the following markets: cars (Porter and Sattler, 1999; Adda and Cooper, 2000; Stolyarov, 2002; Esteban and Shum, 2007; Chen et al., 2013; Schiraldi, 2011); truck tractors (Bond, 1983); commercial aircraft (Pulvino, 1998; Gavazza, 2011a,b); business aircraft (Gilligan, 2004; Gavazza, 2016); and capital equipment (Eisfeldt and Rampini, 2006).

markets, new goods are typically purchased by higher-income consumers, who sell their used goods when they wish to upgrade to the newer models.

The literature has modeled this type of heterogeneity and the corresponding motives for trade as follows. Consumers are infinitely lived and characterized by a persistent taste parameter θ distributed according to some distribution $F(\theta)$. The flow utility from consuming a good of quality q for a consumer of type θ who pays a (possibly implicit) rental price r is given by $\theta q - r$. The parameter θ is therefore the marginal value of quality for this consumer. There can be many interpretations of the source of heterogeneity underlying this θ . A common interpretation is that θ is related to the inverse of the marginal utility of wealth, so that wealthier people have higher valuation for quality. More generally, as in Gavazza et al. (2014), this taste parameter can be decomposed into a component related to income or wealth and an idiosyncratic component.

Durable goods are then modeled as effectively exhibiting a potential menu of qualities (q_1, \dots, q_n) , each associated with the vintage of the good. Goods of quality q_1 are new goods that are the highest quality; the good then gradually depreciates, either deterministically or stochastically, and quality deteriorates over time. Quality q_n , the worst quality, requires special interpretation because it may be determined endogenously by consumers' scrappage decisions, which in turn depend on market frictions.

In a static competitive market in which such a menu of qualities is exogenously given, and consumers exhibit such vertically differentiated preferences, consumers and qualities would be matched assortatively so that higher valuation consumers choose higher qualities. In the case of durable goods, such a menu of qualities can emerge endogenously in steady state due to the depreciation of a good that is newly produced at a constant flow rate y . In this case, sorting requires trade in the secondary market. In particular, all goods of qualities $q < q_1$ (all used goods) are only available if they are sold on the secondary market. If such markets are frictionless, then allocations are efficient, and trading whenever a good depreciates replicates the assortative matching described above for the static case: Consumers are matched to a vintage of the good, they sell the good the moment it depreciates, and they again buy the appropriate vintage to maintain the "correct" lifetime quality consumption. This frictionless case features 100 percent trade of every vintage; allocations are then pinned down by the assortativeness condition, so that the mass of the highest types matches the mass of available quality q_1 goods, and so on.

Prices can easily be characterized in this setting by a chain of indifference conditions and resale values implied by the price of the next vintage, beginning with the indifference of the lowest type between consuming a car of quality q_n good and not having a car (the value of this outside option is typically normalized to zero). Specifically, the price p_i of vintage i is obtained recursively as follows: The price of the vintage n good is such that

$$\theta_n q_n - p_n = 0,$$

where θ_n is the lowest type that consumes a car, so they are indifferent between consuming and not consuming a car. Then, given prices of vintages $i + 1, \dots, n$, the price of vintage i is determined so that the marginal type θ_i , which is the lowest type that consumes vintage i , is indifferent between consuming vintage i and vintage $i + 1$:

$$\theta_i q_i - p_i + \delta p_{i+1} = \theta_i q_{i+1} - p_{i+1} + \delta p_{i+2},$$

where δ is the discount factor. This indifference takes into account the relative attractiveness of the substitute vintage as well as the resale values of both vintages. An equivalent formulation is to define the rental price r_i for holding a vintage i for one period; the indifference conditions would be standard sorting equations as follows:

$$\theta_i q_i - r_i = \theta_i q_{i+1} - r_{i+1} \quad (1)$$

and, of course, the relationship between rental prices and ownership prices is $r_i = p_i - \delta p_{i+1}$.

An important feature of secondary markets, therefore, is that the menu of available qualities is dictated by the durability of the good as measured by its speed of depreciation. In this frictionless case, therefore, finding the equilibrium allocations and prices is straightforward, because the first-best assortative matching solution is also an equilibrium. If there are frictions such as transaction costs, some welfare-improving trades are impeded and matching between quality and consumers is imperfect.

Unfortunately, in the case of transaction costs, finding equilibrium allocations and computing equilibrium prices is much more difficult, and proceeding from a planning problem does not help: We do not have an easy conjecture about the equilibrium allocation such as the first-best outcome described above; the equilibrium allocation will depend in a complicated way on the distribution of types and the level of transaction costs. Therefore, we must solve a full general equilibrium problem of $n - 1$ interrelated markets. In the case of frictions, the menu of qualities that are actually consumed is coarser than in the first best, and the degree of coarseness depends on the level of transaction costs. For instance, if transaction costs are prohibitive, so that secondary markets completely shut down, effectively the only available quality is the bundle achieved by buying a new good and holding it until scrapping. However, this bundle may be different for consumers with different valuations, because those with higher valuations may choose to scrap earlier. Suppose that a consumer scrappage decision is to scrap all goods of quality inferior to q_x . Then the lifetime quality consumed on average is the “discounted” average quality of the good of conditional quality determined by the truncation determining the support $\{q_1, \dots, q_x\}$. More generally, for an intermediate range of transaction costs, different consumers optimally choose different trading strategies, with some consumers potentially keeping only one quality and others keeping a range of qualities.

We will now show that this type of model of a durable good market with secondary markets subject to frictions is useful for exploring a number of interesting issues, including the feedback between the functioning of the secondary market and the market for new goods.

2.2 Market power and secondary markets

In the vertical differentiation framework described above, several papers have adopted what is essentially a price discrimination or market segmentation approach to an old problem: How a monopolist interacts with secondary markets. In this discussion, we will assume that the monopolist can commit to the sequence of policies.⁴

An older literature associated with Swan (1970), Swan (1971), and Sieper and Swan (1973) discusses the optimal choice of durability under what turns out to be a very special technology governing the relationship between the output of new goods and durability. The prototypical example in this literature is the light bulb. In this case, durability is simply the lifetime of the good that maintains constant quality until it dies. An alternative assumption that leads to similar results is that additional units of services can restore the quality of the depreciated good. If either of these assumptions is satisfied—which amounts to assuming that quantity and quality are perfect substitutes—then the choice of the optimal speed of depreciation of the good yields a target level of flow of services of the good at the lowest possible cost. But then, as Swan (1970) points out, under constant returns to scale, the cost-minimizing durability is independent of the particular target level of these flows of services, and therefore a monopolist provides socially optimal durability. For instance, if two light bulbs that last one year are completely equivalent in the minds of consumers to one light bulb that lasts two years, a monopolistic industry provides less light than socially optimal, but light bulbs will last just as long as the ones chosen by the social planner. The monopoly distortion takes the form of reducing output, not durability.

In contrast, when quantity is not a perfect substitute for quality (an extreme example is when there is unit demand and heterogeneous valuation of the unit, as above), then there is no reason to expect that monopoly durability is efficient.⁵ Waldman (1996) explores a model with two types of consumers and a durable good that lasts two periods. Under an additional assumption on the relative frequency of the two types, he shows that durability is underprovided by a monopolist. Hendel and Lizzeri (1999b) allow for richer distributions of valuations, and show that, when holding output constant, durability is underprovided by a monopolist, thereby generalizing the case studied by Waldman (1996). However, in general, beyond the case of two types, monopoly output of a new good and its durability are simultaneously determined—and, of course, monopoly output is typically lower than socially optimal. Hendel and Lizzeri (1999b) show that optimal durability is the result of a trade-off between a resale value effect that makes consumers willing to pay more for the new good because it reduces the effective rental price (pushing toward increased durability) and a

⁴ In the setting of a durable good monopoly, an important issue was first highlighted informally by Coase (1972), and is therefore known as the Coase conjecture: The fact that, absent commitment, a monopolist provides its own competition because of the temptation to lower prices in the future. See Gul et al. (1986) for a rigorous analysis of this phenomenon; see Waldman (2003) for a survey.

⁵ In this model, the choice of durability is the choice of quality of the used good. Therefore, the logic driving the choice of durability is similar to the one driving the choice of quality for nondurable goods, as first outlined by Spence (1975).

substitution effect that threatens the market for new goods (pushing toward reduced durability). Hence, it is possible for monopoly durability to be higher than the amount that would be chosen by a social planner.⁶

Another way for a producer to affect the menu of qualities effectively available to consumers is to change transaction costs. An old intuition going back to the Alcoa case (see the discussion by Swan, 1980) is that, because the used good is a cheap substitute for the new good, increasing transaction costs—or even inducing the used-good market to shut down—may in principle be a profitable decision for a monopolist. However, there are many examples in which manufacturers seem to intervene in the opposite direction. For instance, many manufacturers allow the transfer of warranty coverage across owners. Car manufacturers encourage dealers to accept trade-ins, and they design inspection programs and warranties for “pre-owned” vehicles of their brand. Anderson and Ginsburgh (1994) show that, under certain conditions, a monopolist may indeed wish to shut down the secondary market. Hendel and Lizzeri (1999b) in contrast, show that if durability were endogenous, a monopolist would prefer to reduce durability rather than shut down the secondary market.

2.3 Empirical research

We now move on to contributions that discuss empirical patterns in secondary markets.

Porter and Sattler (1999) use data from title transfers for cars in Illinois to test the predictions of a model of trade in durables as a result of quality decline over time, similar to those discussed above. They exploit differences across model/years in price declines over time. They find that the rate of decline of car prices over time is negatively and significantly correlated with the length of ownership tenure, or with the likelihood that the original owner does not sell the vehicle before the end of the sample period.

Stolyarov (2002) numerically solves a model of secondary markets with transaction costs in line with the model described above. He investigates resale rates across car vintages. Specifically, he is interested in understanding why these resale rates in the U.S. used-car market feature a “double-hump” shape: Young vehicles have low resale rates; the probability of resale subsequently rises, peaks at around 3 to 5 years, then it falls; it rises then again for vehicles approximately 10 years old. The model is capable of replicating these “waves” in resale rates due to consumers’ equilibrium responses to transaction costs. Young vintages do not get traded because they are relatively close substitutes for new goods, so it is not worth incurring transaction costs to upgrade. Once goods depreciate enough, trade rises. However, this applies mostly to

⁶ Rust (1986) assumes that depreciation is stochastic. Despite having homogeneous consumers, he finds that the monopolist distorts durability. This is because the scrapping decision made by optimizing consumers is different from the one preferred by the monopolist. If the monopoly were to rent the good instead of selling, the monopoly solution would be socially optimal in his environment.

the resale of goods that were purchased new. New goods are special because, by construction, they are a vintage that gets fully traded for any level of transaction costs. In contrast, resale for slightly older vintages would mostly appeal to those who bought relatively young used goods—but we observe that very few of those were available for purchase, so trade then falls as an echo of the low level of trade in these earlier vintages. The next wave of high trade volume also takes place as an echo of the first wave, and this requires time in order to justify bearing the transaction costs.

Schiraldi (2011) obtains an estimate of the level of transaction costs in the used-car market based on 11 years of data on sales dates for individual cars over time from a small province in Italy (Isernia). He estimates a demand model with random coefficients. By comparing the share of consumers who choose to hold a given car type each period with the share of consumers who choose to purchase the same car type in that period, Schiraldi is able to identify time-varying transaction costs that are conditioned on the car model. He finds sizable levels of transaction costs that have been falling over time: The average transaction cost was €2700 in 1994, and dropped to €1700 in 2004. He attributes this drop to a reduction in taxes and interest rates. Schiraldi also finds that transaction costs are fairly disperse; for instance, in 2004 the minimum is approximately €800 and the maximum is €4700. Schiraldi then uses his structural model to study the effect of scrappage subsidies.

Chen et al. (2013) consider a similar model, but focus on the effects of the secondary market on the primary market: They consider an oligopoly model with forward-looking firms and compare the effects of secondary markets for both the case of commitment and the case in which manufacturers lack commitment. They show that smoother transactions in secondary markets harm firms' market power by as much as 35 percent. Interestingly, they also show, via a counterfactual, that, when goods become less durable, the harm of secondary markets for producers is reduced.

Gavazza et al. (2014) consider a model that is related to Chen et al. (2013) and Stolyarov (2002), but they focus on different issues. They allow for richer secondary markets in part because they want to study how the extent of trade depends on the degree of heterogeneity in preferences. The paper quantitatively investigates how the distribution of non-durable consumption contributes to this heterogeneity by calibrating the model to match the aggregate volume of trade in the used-car market. Specifically, they enrich the model to incorporate some features that are important in the data, such as multiple cars per household and the combination of exogenous and endogenous scrappage.

Gavazza et al. (2014) perform several counterfactuals to understand the functioning of secondary markets and their impact on the market for new goods. They examine the effects of transaction costs by comparing the baseline calibrated model with two polar cases: frictionless secondary markets with zero transaction costs and the complete shutdown of secondary markets with prohibitive transaction costs. We naturally expect any changes in secondary markets to affect primary markets. Thus, the supply response of new-goods producers is an important element in determining the welfare consequences of secondary markets' frictions. They consider two extreme supply scenarios that help highlight how primary markets adjust: (1) a per-

fectly elastic supply—i.e., the price of new cars does not respond to changes in transaction costs, but the quantity does; and (2) a perfectly inelastic supply—i.e., the quantity of new cars does not respond to changes in transaction costs, but the price does. These counterfactuals are useful for understanding the importance of transaction costs for manufacturers, since they indicate that either output or prices change when transaction costs change, even in an oligopolistic market for new cars. In these counterfactuals, Gavazza et al. find that the impact of transaction costs on allocations in both the secondary market and the market for new goods is large. In contrast, the effects on aggregate welfare are small, although the distribution of these effects is uneven, with low-valuation households suffering large losses from increases in transaction costs.

Barahona et al. (2020) examine the consequences of vintage-specific vehicle restrictions as a tool to reduce pollution caused by vehicle traffic. Driving restrictions are a common (often temporary) policy response to heavy pollution in cities. Some policies give drivers the incentive to buy a second, cheaper car that is therefore older and potentially more polluting than the first one. Vintage restrictions that allow freer access to less-polluting models generate the opposite incentive, by giving drivers the incentive to switch to a cleaner car. Barahona et al. build on Gavazza et al. (2014), discussed above, by incorporating an intensive margin of how much to drive. They study a wide range of environmental policies similar to those implemented in a variety of countries and analyze the transitional dynamics following the introduction of each policy (instead of focusing only on steady-state comparisons). They argue that the most efficient policies are those that target the types of cars people drive—and thus the composition of the fleet that will eventually obtain—rather than the number of miles people drive.

Gavazza and Lanteri (2021) study the equilibrium dynamics that arise in consumer durable goods markets in response to aggregate credit supply shocks. They embed the vertical structure of durable goods discussed above into a general-equilibrium model with heterogeneous households that face idiosyncratic income risk and borrowing constraints. Their model successfully matches several empirical patterns of U.S. car markets around the Great Recession. Specifically, after a negative credit shock, debt-constrained households postpone the decision to scrap and upgrade their cars. The economy experiences a period of low resale prices for used cars, which reduces wealthier unconstrained households' incentives to replace their cars, thereby decreasing new-car sales.

One important limitation of the main model that has been used as a reference point for the empirical contributions discussed so far is that it abstracts from features such as horizontal differentiation among car models, as in Berry et al. (1995). It is in principle straightforward to allow for additional vertical differentiation in the market for new goods that is potentially preserved by used vintages—e.g., a 3-year-old Mercedes would be equivalent to a new Toyota. However, it is difficult to integrate both secondary markets with vertical differentiation and rich horizontal differentiation of new (and used) car models. The focus of our discussion has been on the replacement patterns that emerge from gains from trade due to vertical differentiation among dif-

ferent ages of a given model of a durable good. Integrating both important aspects of durable goods markets is still an important, largely open, question. The work of Gillingham et al. (2019) is a recent exception. As in the literature discussed above, they allow for transaction costs and vertically differentiated preferences. They add multiple brands and individual choice-specific heterogeneity and show how to compute such a model. They then calibrate it to Danish car data and evaluate two policy counterfactuals: a registration tax and the merger of two car sellers.

2.4 Role of intermediaries in overcoming transaction costs

We saw above that sometimes manufacturers have both the incentive and the means to affect transaction costs in secondary markets. We now discuss the role of third parties such as intermediaries. Examples are pervasive: automobile dealers, aircraft lessors, various used equipment brokers, and real estate agents.

Spulber (1999) builds on Coase's fundamental insights regarding firms' role in reducing transaction costs. He provides a book-length treatment of firms' role as intermediaries between customers and suppliers.

Gavazza (2011a) discusses the role of lessors in reducing frictions in secondary markets for aircraft. He reports that, as of 2010, approximately one-third of all aircraft operated by major carriers were under an operating lease—i.e., a rental contract between a lessor and an airline for use of the aircraft for a period of 4 to 8 years. Lessors are active buyers on the primary market, and their acquisitions have increased over time. Moreover, lessors are also active participants in secondary markets, since they frequently buy used aircraft and, more importantly, lease out each aircraft several times during their useful lifetime. Gavazza develops a model of aircraft transactions to understand the role of lessors when secondary market trading is subject to frictions. The model combines five ingredients: (1) Carriers have heterogeneous stochastic productivity; (2) carriers have heterogeneous volatility in sales, and therefore need for aircraft; (3) aircraft can be bought or leased; (4) carriers incur costs to sell aircraft; and (5) lessors incur per-period costs to monitor their assets. Secondary markets allow carriers to adjust their productive capacity in response to shocks. In this environment Gavazza shows that leased aircraft change carriers more frequently and have higher utilization. These phenomena are due to two effects. First, for any given carrier characteristics, trading is less frequent for airlines that buy because transaction costs are higher; second, there is a selection effect due to higher-volatility carriers who choose to lease because they anticipate transacting more.

Gavazza (2011a) uses a rich dataset on commercial aircraft to provide evidence on the model's qualitative implications. He finds that leased aircraft have: (1) holding durations 38 percent shorter than owned aircraft and (2) flying hours 6.5 percent higher than owned aircraft. The empirical analysis shows that leased aircraft are parked inactive less frequently than owned aircraft, and that, conditional on being in use, leased aircraft have a higher capacity utilization than owned aircraft. Moreover, he finds evidence in favor of both effects highlighted by the model, but their empirical relevance is lopsided: Most of the differential patterns in trading and utilization arise because

ownership's larger transaction costs widen owners' inaction bands relative to lessees', and carriers' self-selection into leasing plays a minor role. Finally, Gavazza calibrates the model to investigate whether it is quantitatively consistent with the data. Interestingly, the calibration shows that small differences in carriers' volatilities can lead to the observed larger differences in trading and utilization between leased and owned aircraft, and confirms that self-selection of carriers does not play a dominant role.

3 Asymmetric information

3.1 Theory

3.1.1 Static adverse selection, exogenous ownership

It is natural to think that the current owner of a durable good has information about its quality superior to that of potential buyers. This logic is presumably what led Akerlof (1970) to prominently feature a used car in a leading role in his seminal paper on adverse selection. Asymmetric information is one of the main sources of potential impediments to trade. Akerlof's model can partially fit within the model outlined above in Section 2. One important difference is that in Akerlof's model, there is only a market for used goods and only one age of used goods; there is no interplay between the used-good market and the market for new goods. He assumes that the world begins with an owner of a used good who has valuation θ_L for quality and a potential buyer who has valuation $\theta_H > \theta_L$ (more specifically in his example, $\theta_L = 1$, $\theta_H = 3/2$). Therefore, it is common knowledge that there are gains from trade: In an efficient allocation, the used good should be sold regardless of its quality. Akerlof assumes that quality q is distributed uniformly on $[0, 1]$, and that the owner knows q and the buyer does not. The key insight is that the sellers' willingness to sell a car at a particular price is a signal about the quality of the car.

The equilibrium for this example can be easily characterized. First, we need to work out buyers' beliefs about quality conditional on a car being for sale at a particular price. At a price P , a seller is willing to sell if and only if $q \leq P$. Therefore, expected quality conditional on availability is given by

$$E(\theta | \theta \leq P) = \frac{P}{2} \text{ (because of uniform distribution).}$$

The utility of buying at price P is given by

$$\frac{3}{2}E(\theta | \theta \leq P) - P = \frac{3}{4}P - P = -\frac{1}{4}P \leq 0 \text{ for } P \leq 1.$$

Therefore, it is never optimal to buy. The only equilibrium is one in which there is **no trade**: i.e., a complete market breakdown. This complete market breakdown is a special case; in particular, it depends on the fact that valuations of the lowest quality are the same for the buyer and the seller (in this case, because this quality is zero). However, the idea of adverse selection as an impediment to desirable trades

is general. Note that information asymmetry is essential to this outcome. If sellers and buyers were equally informed, then trade for quality q could take place at some price $P(q)$ such that $q \leq P(q) \leq \frac{3}{2}q$. If, in contrast, sellers and buyers were equally ignorant—with both having a uniform prior over unknown quality—then trade in all qualities could take place at some price P^I such that $\frac{1}{2} \leq P^I \leq \frac{3}{4}$. In either case, the final allocation is efficient.

Adverse selection has become an important source of potential market inefficiencies. Sometimes such inefficiencies are viewed as sufficient to warrant government intervention. However, it is important to examine the tools a potential government authority has that may ameliorate allocations in this environment. In the example described above, it turns out that no improvement is possible if we expect budget balance (no subsidies) and respect agents' ability to walk away from an undesirable transaction.

Of course, an equilibrium under adverse selection requires a potentially unrealistic amount of sophistication on the part of consumers. It is tempting to expect that consumer naivete could improve outcomes by increasing consumers' willingness to pay for quality. However, this expectation is not necessarily warranted. For instance, Esponda (2008) offers an equilibrium notion for economies with adverse selection and a specific form of naivete on the part of consumers. He shows that in Akerlof-type environments the adverse selection problem is exacerbated if the buyer is naive.

3.1.2 Dynamic trading with exogenous initial ownership

Akerlof's (1970) model is static and considers a single opportunity to transact. In many possible applications, if goods are actually durable—as in the case of used cars—it is natural to think that nothing prevents market participants from attempting to trade again after an initial period. Thus buyers whose initial offer is rejected could, for instance, make additional, and potentially higher, offers. We could also consider competitive equilibria with multiple periods and changing prices over time.

Janssen and Roy (2002) and Janssen and Roy (2004) extend Akerlof's model by allowing for trade to take place in a possibly indeterminate number of subsequent periods. Akerlof's specific example with no trade survives such an extension. However, this feature is a special case. Janssen and Roy (2002) consider competitive equilibria and show that, when used-goods markets are open on many dates, more additional trades are possible over time. Prices and traded qualities increase over time, and the time to trade acts as a screening device because owners of high-quality cars are relatively more willing to wait to obtain high prices compared with owners of low-quality cars. However, a delay in trade for high-quality cars is necessary to induce incentive-compatible trade, and this leads to a welfare loss. This welfare loss does not disappear as the discount factor goes to one, or as the time between opportunities to trade shrinks to zero, because the welfare loss is intrinsic to the screening process.

Hörner and Vieille (2009) explicitly model the bargaining process in an environment similar to that of Janssen and Roy (2002) and compare equilibrium outcomes in the case in which offers are private with the case in which offers are publicly observ-

able. With public offers, in the unique equilibrium, trade is limited by the fact that only the first offer is a “serious offer” at which trade may take place. If this offer is rejected, no further serious offer is submitted on the equilibrium path. This equilibrium is somewhat paradoxical: At each time period, a buyer would submit a serious offer if he were assured to be the last to bid. Thus it is precisely the competition from future buyers that deters him from making such an offer. However, on the equilibrium path, all subsequent buyers also submit losing offers. In contrast, if offers are secret, agreement is always eventually reached, albeit with delay as in Janssen and Roy (2002).

Daley and Green (2012) extend the idea of dynamic trading by incorporating the natural possibility that public information about the quality of the good arrives over time. Specifically, they assume that the good is one of two possible qualities, and model the news process as Brownian motion with a drift that depends on the quality of the good. They then characterize the Markov equilibria of this model. The most intriguing aspect of their characterization is that equilibrium can feature the following behavior. As information about quality becomes more favorable, and beliefs about quality lie in some intermediate range, trade dries up and participants wait for news in order to resume trade. When beliefs are very optimistic, equilibrium trade is immediate and efficient. When beliefs are very pessimistic, trade is partial because the high type rejects equilibrium offers, whereas the low type randomizes, and beliefs evolve both endogenously, due to the improved selection of types, and exogenously in response to the public news.

The literature on trading in asset market has also incorporated adverse selection as an important ingredient at least since the work of Kyle (1985) and Glosten and Milgrom (1985). The 2008 financial crisis has spurred a new wave of research on potential market failures in secondary markets for financial assets and on the role of government intervention in such markets. Particularly related to the material in this chapter are Tirole (2012), Philippon and Skreta (2012), and Guerrieri and Shimer (2014). Fuchs and Skrzypacz (2015), for instance, examine whether government intervention can be beneficial in such dynamic markets. They show that restricting trading opportunities after an initial round of trade is always optimal, because this induces earlier trading.

3.1.3 Endogenous initial ownership

Let us now temporarily return to an environment in which there is a single opportunity to trade a used good, and examine the interactions between the market for used goods and the market for new goods. The same logic that suggests that owners of a used good have private information about its quality suggests that the current potential seller presumably bought the good when it was new. A natural question to ask about the ownership arrangement assumed in Akerlof (1970) is how the low-valuation consumer came to own the used car: In a model with vertical differentiation in which consumers choose whether to buy new or used cars, as we saw in Section 2.1, only high-valuation consumers would be used-car sellers. Thus, in such a model, the ownership arrangement Akerlof depicts would not arise. For instance, if new cars are

produced only once, the equilibrium would involve higher types buying new goods and holding on to used goods. Thus the equilibrium of this alternative model would also feature zero trade in the secondary market. However, this allocation would be efficient: The high-valuation consumer already owns the used good. These considerations suggest that it is important to simultaneously model new and used goods markets, and to incorporate the motives for trade—such as quality upgrading—in order to understand the impact of adverse selection in durable goods markets such as cars. Furthermore, it is of interest to understand the consequences for the new-car market of potential adverse selection in the used market.

Hendel and Lizzeri (1999a) consider a model of the car market that is related to the model of transaction costs discussed above. They consider a steady-state market for durable goods, and assume that there is a constant flow of new cars produced every period. New goods depreciate, and used goods have a distribution of possible qualities. Consumers are infinitely lived and have per-period valuations for quality denoted by θ , as in the model described in Section 2.1. These consumers self-select into becoming new- or used-car buyers, depending on their valuation and the endogenously determined average quality of used units. Ownership of used cars is thus endogenous. Consumers who buy new goods decide whether to sell the used units, and there is asymmetric information about quality.

Hendel and Lizzeri (1999a) find that, in contrast to the example in Akerlof (1970) described above, the used market never shuts down and the volume of trade can be quite substantial, even in cases of severe information asymmetries. The main reason for the high volume of trade is that the source of gains from trade differs from that postulated by Akerlof (1970). In this model, as in many markets for durable goods, an important reason to sell a used good is to enjoy the higher quality offered by a new one. Used-good sellers are also new-good buyers—i.e., they are the consumers most sensitive to quality and who therefore have a high opportunity cost of holding on to a lemon. Regarding welfare, the equilibrium allocation under adverse selection is generally inefficient. However, Hendel and Lizzeri give an example in which the equilibrium under adverse selection results in the first-best allocation; they also show that increasing incentives to trade used units can lead to a less efficient allocation. The reason for this finding is that, although increasing the volume of trade raises the quality of traded used units, this increase is not necessarily welfare improving because in this model (in contrast to Akerlof, 1970) the used-car buyers have lower valuations than sellers of the used cars. They conclude that taking into account the interactions between the new and used market lessens the impact of adverse selection. Hendel and Lizzeri then show that their model predicts a different relation between turnover and reliability across car brands, depending on whether the trade friction is a transaction cost or asymmetric information.

In this type of dynamic environment, at the time of purchase of the new good, there is no information asymmetry. Hence, there is interesting scope for contractual arrangements that could ameliorate the adverse selection problem. Hendel and Lizzeri (2002) and Johnson and Waldman (2003) consider the role of leasing contracts in the context of durable goods markets with adverse selection. Leasing con-

tracts specify a rental price for an initial period—say, 36 months—and an option price at which the lessee can choose to purchase the car at the end of the lease period. Because this option price is set at the initiation of the lease, before the lessee observes the quality of the used good, it can generate a degree of commitment to trade regardless of the quality the lessee will later observe. Therefore, such contracts can in principle ameliorate the adverse selection problem. Hendel and Lizzeri (2002) show that it is typically not optimal, for either a monopolist or a social planner, to set a prohibitively high option price that would lead to full trade. They also show that leasing contracts can be used to profitably segment the market. They also argue that optimal leasing contracts match some of the empirical patterns observed in the contrast between leased and purchased cars in the U.S. car market.

Although consumers are infinitely lived in Hendel and Lizzeri (1999a), the papers mentioned so far assume that goods last two periods. This assumption implies that, as in Akerlof (1970), there is only one period in which trade in a specific used good is potentially subject to adverse selection.

Hendel et al. (2005) combine the insights of Hendel and Lizzeri (1999a) and those from the literature on dynamic trading, such as Janssen and Roy (2002). As in Hendel and Lizzeri (1999a), Hendel et al. (2005) consider a model in which there is a flow of new goods every period. In contrast to those papers, they allow markets for used goods to be open at any relevant instant. In the simplest environment they consider (called the *simple depreciation model*), goods are born with a deterministic quality. Goods last an indeterminate amount of time and depreciate stochastically one step at a time. Therefore, after a sufficiently long time, a good can in principle be any of a large number of possible qualities. Hendel et al. consider the possibility that the number of previous owners of a good, or the number of times the good has been traded (they call this the *vintage* of the good), can serve as an accurate signal of the quality of the good. They show that this vintage pricing mechanism generically leads to inefficient allocations with sales contracts, in which all vintages are owned by a consumer. However, they show that there exist rental contracts (of indeterminate termination) that lead to first-best allocations. Furthermore, such efficient contracts are part of an equilibrium choice by competitive producers of the goods. The prices in these rental contracts are the same as the equivalent full-information ones defined in Eq. (1). Hendel et al. then extend the analysis to richer depreciation patterns and show that similar mechanisms achieve approximate efficiency. This paper highlights the fact that, in the case of durable goods, distortions due to adverse selection are a result of the combination of asymmetric information and some other friction that either does not allow markets to be open frequently enough or does not allow some simple efficient contracts to be written.

3.2 Empirics

There have been several attempts to document the effects of adverse selection in the used-car market. Some of the evidence is also discussed in the next subsection on the emergence of intermediaries.

The evidence is mixed. A general challenge of empirical tests of asymmetric information models is that the parties' information may not be perfectly observed by the researchers. An additional challenge faced by empirical studies on durable good markets is that the transaction prices of used goods are frequently unreported in the data or, if reported, may vary depending on parties' bargaining power (rather than quality), in contrast to menu prices in insurance or credit markets that are set by insurers or lenders (Einav et al., 2021).

Among these studies, Bond (1982) and Bond (1984) compare the frequency of maintenance of otherwise similar (age, mileage) pickup trucks purchased new and trucks purchased used. He found no significant difference for younger vehicles, but he found evidence of more frequent maintenance of older vehicles that were purchased used. He interprets the latter as evidence of adverse selection. This interpretation is reasonable, but there may be other interpretations; such patterns could in principle also arise if quality were observable. Lacko (1986) compares the owner-reported quality of cars purchased from friends or relatives with that of cars purchased through newspaper ads; his hypothesis is that friends are less likely to sell a lemon. He also finds a difference only for older vehicles.

Gilligan (2004) uses data from the used business aircraft market to test the predictions made of Hendel and Lizzeri (1999a), Porter and Sattler (1999), and Hendel and Lizzeri (2002). Gilligan finds an inverse relationship between depreciation and trading volume for less reliable brands of used business aircraft, which suggests evidence of adverse selection (according to the model of by Hendel and Lizzeri, 1999a). He also finds an increase in the direct relationship between depreciation and trading volume for aircraft models with relatively high lease rates, which is in line with the logic of Hendel and Lizzeri (2002) discussed above.

Lewis (2011) examines the consequences of opportunities for quality disclosure in the used car market. He uses data from eBay motors, which is a large used-car marketplace in which sellers are able to post photos of their vehicles. He finds evidence that the photos and text posted by the seller on the auction website positively influence demand and transaction prices. He then shows that disclosure costs affect how much information the seller decides to post, and these costs in turn lead to lower average transaction prices.

3.3 Role of intermediaries in overcoming asymmetric information

The presence of distortions due to information asymmetries can lead to the emergence of third-party institutions that have the technology to discover at least part of the information of privately informed agents and can then partially reveal it to uninformed agents. These are known as certification intermediaries. There are many examples of institutions that play this role: laboratories in markets for industrial products (a famous example is Underwriters Laboratory) auditors, credit rating agencies, etc.

Biglaiser (1993) studies a model that considers whether, in a dynamic adverse selection matching model, intermediaries can improve welfare. In his model the middleman endogenously decides whether to become an expert who can discover

information.⁷ Lizzeri (1999) explicitly models the incentive for such intermediaries to use their power and ability to manipulate their information revelation policy to extract surplus from both sides of the market. He studies a model in which two initially uninformed buyers bid simultaneously to buy a product from an informed seller. In the baseline case, in a market without an intermediary, transactions would occur at a price equal to the expected value of the product. The intermediary is assumed to be in a position to commit to a certification price and a disclosure rule. Lizzeri (1999) shows that, if sellers pay to be certified, under an assumption about the density of the quality distribution, the unique equilibrium involves the intermediary capturing all of the surplus while revealing no information. This is therefore a model of a parasitic intermediary. Lizzeri then studies alternative scenarios, such as the case in which buyers pay for the information, and the case in which there are multiple intermediaries. In these cases more information is typically revealed. Dranove and Jin (2010) survey some of the theoretical and empirical literature on this topic.

Genesove (1993) empirically studies the role of dealers in markets with asymmetric information. He argues that different types of car dealers should have different willingness to sell used cars in wholesale markets. In particular, Genesove distinguishes between dealers who exclusively sell used cars and dealers who also sell new cars. He argues that the latter should be more eager to sell older used cars wholesale. He finds evidence from used-car auctions that is consistent with this prediction.

Biglaiser et al. (2020) provide a model of used-car dealers as intermediaries who ameliorate an adverse selection problem. In the model, dealers have a superior testing technology and offer perfect warranties; they only sell cars that pass the test. The main predictions of the model are: (1) there is a dealer premium for all ages; (2) the dealer share declines in the age of the car; (3) the Dealer premium is hump-shaped in dollar terms; (4) the dealer premium is increasing over time in percentage terms; and (5) The dealer premium is higher for unreliable cars. Biglaiser et al. then provide simple tests of these predictions by using rich data from Virginia and Pennsylvania. The tests are successful in all dimensions. They also perform a separate supplementary test of asymmetric information from additional Pennsylvania data, and find that cars sold privately are resold earlier.

4 Search frictions

A large literature has studied the role of search frictions in markets, with a wide range of influential applications, such as labor markets (e.g., Mortensen and Pissarides, 1994; Rogerson et al., 2005); housing markets (e.g., Wheaton, 1990); and financial markets (e.g., Duffie et al., 2005). Moreover, the theoretical and empirical literatures on markets with search frictions have developed in a tightly connected manner, partly

⁷ Biglaiser and Friedman (1994) addresses similar issues in a moral hazard model with experience goods.

because structural analyses have a strong tradition in this area, most notably in models of the labor market (for a survey, see Eckstein and Van den Berg, 2007).

In this section, we survey some of the key theoretical models and empirical contributions in the literature on consumer product markets. This literature has been successful in accounting for a wide range of market outcomes that are at odds with benchmark models with full information: price dispersion for identical goods, trading delays and unsold inventories, and the vertical structure of industries with retailers and intermediaries. We refer the reader to Baye et al. (2006) and Anderson and Renault (2018) for more thorough surveys of the literature on consumer search markets, with a focus on theoretical contributions.

We should point out that the literature on consumer search is related to the literature on consideration sets. These two literatures share an interest in understanding the details of the process that leads consumers to purchase a specific good, although they emphasize slightly different aspects. For instance, the literature on consideration sets originated in marketing, focuses mostly on consumers, and has a relatively stronger empirical bent. In contrast, the literature on consumer search started in economics, studies mostly equilibrium (i.e., demand and supply) models, and has a stronger theoretical input. Recently, the overlaps and common themes of these two literatures have increased, partly because of greater data availability and the related increasing shift toward empirical research. Honka et al. (2019) provide an insightful survey of the relationships between these related literatures, and reviews several empirical papers.

4.1 Theoretical models

The starting point of the search literature is the observation that price dispersion is ubiquitous even in homogeneous product markets, as noted in Stigler's (1961) seminal paper. The earliest contributions, including that of Stigler, focused on buyers' search behavior in the face of a distribution of prices, but they did not explain how sellers set these prices and what forces led to such price dispersion. Later contributions developed equilibrium models with a nondegenerate distribution of prices in which buyers' and sellers' actions are mutually consistent, and emphasized the heterogeneity of buyers' information as the key determinant of price dispersion.

The literature has since incorporated additional market features. Some of the main features are (1) explanations for why *ex ante* identical consumers may endogenously obtain heterogeneous levels of information in equilibrium; (2) an analysis of the process by which consumers gather price quotes (e.g., the contrast between simultaneous and sequential search); (3) the enrichment of seller strategies by allowing for price advertising or obfuscation; and (4) incorporation of product differentiation, thereby expanding consumers' motives from simple search for the lowest price to a search for the highest idiosyncratic match value. We group these theoretical models in five main strands, though of course the boundaries between strands may not be sharp.

4.1.1 Information clearinghouse

In the first group of models, goods are homogeneous and consumers gather information about prices through a clearinghouse, such as a newspaper or a price-comparison

website (Salop and Stiglitz, 1977; Varian, 1980). In most models in this group, equilibrium price dispersion arises from some ex ante heterogeneity across consumers or firms. Specifically, in Salop and Stiglitz (1977) and Varian (1980), some consumers are informed about all prices and buy from the lowest price seller, whereas other consumers are uninformed and buy from a random seller. Given this heterogeneity, sellers can choose to either set a high price and only sell to some uninformed consumers, or set a lower price to also attract informed consumers; in equilibrium, sellers must be indifferent between these alternative pricing strategies, and thus price dispersion arises.

More formally, consider Varian's (1980) model. N firms produce a homogeneous good at zero marginal costs. A continuum of consumers of measure one each have unit demand and homogeneous valuation v .⁸ An exogenous fraction $1 - \lambda$ of consumers is informed about all prices and buys at the lowest available price; the complementary fraction λ is uninformed and buys from a random supplier.

In this setting, there can be no pure-strategy equilibrium. In order to see this, assume by way of contradiction that such an equilibrium exists. At any positive price, the traditional Bertrand logic applies and each firm has an incentive to undercut rivals and attract informed consumers. If price equals zero, then each firm makes zero profit. But a deviation could then earn strictly positive profits by increasing the price to some level $p \leq v$: at such a price, the firm will serve some uninformed customers. Thus, the only possibility for an equilibrium is one in which some firms use mixed strategies, and thereby feature price dispersion.⁹

In a symmetric mixed-strategy equilibrium, each firm chooses prices according to the cumulative distribution function $F(p)$. This distribution $F(p)$ has some important properties: (1) it does not have any step, otherwise a firm could strictly increase its profits by pricing slightly below the price with a mass point; and (2) the upper bound of its support equals the monopoly price v , because only uninformed consumers buy at the highest possible price, and thus sellers can extract all surplus from those buyers. Moreover, because firms are indifferent between any price p in the support of the equilibrium distribution $F(p)$, the following indifference condition applies:

$$pQ(p) = vQ(v),$$

where

$$Q(p) = \frac{\lambda}{N} + (1 - \lambda)(1 - F(p))^{N-1} \quad (2)$$

is each seller's demand function at price p . The first term in (2) is the demand of uninformed consumers who buy at random from one of the N sellers; the second term

⁸ To simplify exposition, we assume throughout this section that buyers have unit demand, but if the profit function is increasing in price up to the monopoly price, the results generalize to the case in which buyers have a downward-sloping demand function.

⁹ Baye et al. (1992) study symmetric and asymmetric equilibria in Varian's (1980) model, and show that there is a continuum of asymmetric equilibria in which some (but not all) firms place a mass point at the monopoly price.

is the demand of informed consumers who buy at the lowest available price, because in equilibrium a firm setting p is the cheapest seller with probability $(1 - F(p))^{N-1}$.

Using $F(v) = 1$, the equilibrium pricing strategy equals

$$F(p) = 1 - \left(\frac{\lambda}{N(1-\lambda)} \left(\frac{v}{p} - 1 \right) \right)^{\frac{1}{N-1}} \quad (3)$$

for $p \in [\underline{p}, v]$, where the lower bound \underline{p} solves $F(\underline{p}) = 0$.

It is interesting to explore the comparative statics of the equilibrium price distribution. The distribution in (3) changes in an intuitive way when the fraction λ of uninformed consumers rises; specifically, prices are higher (in a first-order stochastic dominance sense), and thus both types of consumers pay higher prices. The comparative statics when the number of sellers increases are quite subtle. In this case, firms' incentives to charge extreme prices increase, since a larger number of competitors forces each seller to set lower prices in order to attract informed consumers. However, a larger number of competitors also makes it relatively more profitable to target uninformed consumers only, and thus sellers charge high prices more frequently. These stronger incentives to set extreme prices imply that in equilibrium informed consumers pay lower expected prices and uninformed consumers pay higher expected prices when the number of sellers increases from N to $N + 1$.

4.1.2 Simultaneous search

In the second group of models, sellers and goods are still homogeneous, consumers are also homogeneous ex ante, and consumers search simultaneously by choosing at the beginning of the game the number of quotes to obtain from sellers, paying a per-quote cost.¹⁰ Many simultaneous-search models require that consumers obtain their first price quote for free in order to participate in the market and obtain non-negative surplus.

Burdett and Judd (1983) develop one of the most noteworthy models with simultaneous search. We simplify its exposition in order to connect it to Varian's (1980) model described above by assuming that all consumers are ex ante uninformed, but they can decide to become informed by paying a search costs s , which is common across consumers. The connection between Varian (1980) and Burdett and Judd (1983) makes it clear that clearinghouse models and simultaneous-search models share many features, with the main difference being that consumers pay one cost to observe all (or most) prices in clearinghouse models, whereas the cost is often per quote in simultaneous-search models.

The ex ante homogeneity of consumers enlarges the set of possible equilibria. Most notably, a monopoly equilibrium always exists in which consumers do not become informed because they expect all firms to charge the monopoly price, and all

¹⁰ The literature on consideration sets described above is particularly close to the literature on simultaneous search, as it focuses on the first-stage process of selecting the number of sellers/products before choosing one from among them.

firms charge the monopoly price because firms expect consumers to not become informed and therefore not compare prices. However, there is no equilibrium in which all consumers become informed. If that were the case, all firms would set identical prices equal to the marginal cost, undermining consumer incentives to become informed.

Burdett and Judd (1983) show that, if the search cost s is not too large, symmetric mixed-strategy equilibria exist that display features similar to those in Varian (1980): Sellers are indifferent between a range of prices chosen according to distribution (3), and consumers are indifferent between becoming informed or not, with a fraction λ choosing not to become informed and the complementary fraction $1 - \lambda$ choosing to become informed. The critical fraction λ satisfies

$$p_U(\lambda) - p_I(\lambda) = s, \quad (4)$$

where $p_U(\lambda)$ and $p_I(\lambda)$ are the expected prices uninformed and informed consumers pay, respectively. The difference $p_U(\lambda) - p_I(\lambda)$ satisfies $p_U(\lambda = 0) - p_I(\lambda = 0) = p_U(\lambda = 1) - p_I(\lambda = 1) = 0$, and it is strictly increasing and then strictly decreasing as λ increases from 0 to 1. Thus, if s is not too large, up to two symmetric mixed-equilibria exist, with the one with the lower λ being a stable equilibrium and the one with the higher λ unstable.

Janssen and Moraga-González (2004) build on Burdett and Judd (1983) to study the interactions between consumers' endogenous incentives to search (i.e., the number of sellers they sample) and oligopolistic pricing. This paper shows that there are at most three equilibria with price dispersion, characterized by low, moderate, and high search intensity. It further studies the effects of an increase in the number of firms on consumers' search behavior and expected prices, and notably determines the conditions under which the equilibrium expected price increases in the number of firms.

Armstrong and Vickers (2021) provide a general treatment of simultaneous-search models by studying oligopolistic pricing in markets in which consumers may differ in their consideration sets. They take as given the distribution of consideration sets among consumers and study how the nature of this distribution affects the nature of competitive interactions. A key object in their model is the degree of overlap among consideration sets, because this overlap affects the extent of competition among firms. Armstrong and Vickers show that an important measure of the overlap between any particular set of firms is the ratio of the reach of the set—i.e., the fraction of consumers who consider all firms in the set—to the product of reaches of the individual firms within the set.

4.1.3 Sequential search

Models with simultaneous, fixed-sample-size search suffer from the potential drawback that in some cases it may not be optimal for a consumer to obtain a fixed number of quotes, as Rothschild (1973) recognized. If a consumer finds a low price, the additional expected benefits of further quotes may be lower than the costs of gathering

them. Hence, the literature has developed models in which consumers search sequentially. Morgan and Manning (1985) provide conditions for either simultaneous or sequential search to be optimal, and emphasize that simultaneous search allows the decision maker to gather information more quickly than sequential search, but may result in overinvestment in information.

Diamond (1971) is one of the first papers to consider sequential search in a market with homogeneous buyers and homogeneous sellers. This paper obtains a result that is known as the Diamond paradox: Regardless of the number of firms, the only equilibrium features all sellers setting the monopoly price and consumers not searching at all. The logic of this result is simple. Assume that consumers observe one price quote for free, but they have to pay a search cost to learn each additional price. If firms charge prices below the monopoly price, each firm setting the lowest available price could increase its profits by increasing its price by a small amount—lower than the search cost—because this increase would not trigger further consumer search. The monopoly price is the only one without such incentives, and thus the only equilibrium. This outcome features no price dispersion and no consumer search.

The literature has attempted to overcome the Diamond paradox by introducing heterogeneity among sellers and/or buyers. A critical step in these models, developed first by Kohn and Shavell (1974), is a characterization of consumers' optimal decisions in sequential search, later generalized by Weitzman (1979) as discussed below. Consider n symmetric options, with associated utility u_i distributed independently and identically according to the cumulative distribution function $F(\cdot)$ with support $[\underline{u}, \bar{u}]$. Consumers have an outside option of u_0 and sample each new option at cost s , but can recall a previously sampled option for free. In this sequential-search environment, we can use backward induction to show that consumer behavior is characterized by a stationary reservation utility a :

$$E(\max\{u - a, 0\}) = \int_a^{\bar{u}} (u - a)dF(u) = s \quad (5)$$

that equates the benefit of an additional search to its search cost s . This reservation utility determines the optimal stopping rule: If $a \leq u_0$, consumers do not search; if $a > u_0$, they search until they find an option with utility of at least a , or, if they do not find an option with utility greater than a , they pick the option with the highest utility if it is greater than u_0 .

Stahl (1989) embeds this optimal decision in an equilibrium search model with search cost heterogeneity among buyers. Stahl shows that a fraction of consumers with zero search cost can account for both price dispersion and search in equilibrium. More precisely, consumers with zero search costs visit all sellers and buy from the cheapest one, whereas consumers with positive search costs follow the optimal sequential search (5) with reservation price r . Hence, the reservation price r becomes the upper bound of sellers' pricing strategy, because consumers with positive search costs are the only consumers who purchase at the highest price. Therefore, sellers face the same trade-off as that in Varian's (1980) model between serving all consumers with a low price and targeting consumers with positive search costs with a

high price. In equilibrium, sellers are indifferent between a range of prices and the mixed-strategy equilibrium equals the price distribution (3) with upper bound r . In equilibrium, consumers with positive search costs search one seller only and buy from it.

4.1.4 Sellers' non-price behavior

A further strand of the literature considers the consequences of allowing sellers to expand the scope of their activities. Starting with Butters (1977), some authors have developed models in which sellers convey price information to consumers; for example through advertisements. Consumers then choose from the set of firms they become aware of. More recently, several papers go in the opposite direction and argue that sellers strategically make their prices complex for consumers, thereby increasing consumers' costs of evaluating different products and firms' market power (e.g., Carlin, 2009; Wilson, 2010; Ellison and Wolitzky, 2012).

To understand the role of price advertising, we discuss Stahl's (1994) generalization of Butters's (1977) seminal model. N symmetric firms produce a homogeneous good and can each send independent price ads to a fraction α of homogeneous consumers with valuation v at a convex cost $c(\alpha)$. Firms simultaneously choose their ad intensity α and price p .¹¹

In a symmetric equilibrium, all firms choose a pure strategy with the same advertising intensity α and a mixed strategy for prices, selecting them according to the cumulative distribution function $F(p)$. Firm revenues equal

$$p\alpha(1 - \alpha + \alpha(1 - F(p)))^{N-1} = p\alpha(1 - \alpha F(p))^{N-1}. \quad (6)$$

This expression combines effects from the by now familiar two groups of customers: uninformed consumers, who receive ads from one firm only, and informed customers, who receive at least one other ad from a rival firm at a price higher than p . Evaluating firm revenues/profits (6) at the monopoly price v and exploiting indifference between a generic equilibrium price p and the monopoly price v , the equilibrium distribution of prices equals

$$F(p) = \frac{1}{\alpha} \left(1 - (1 - \alpha) \left(\frac{v}{p} \right)^{\frac{1}{N-1}} \right). \quad (7)$$

In turn, the optimal ad intensity α_i of firm i maximizes its profits, given the ad intensity α of other firms:

$$\max_{\alpha_i} v\alpha_i(1 - \alpha)^{N-1} - c(\alpha_i).$$

Thus, in a symmetric equilibrium $\alpha_i = \alpha$ satisfies:

$$v(1 - \alpha)^{N-1} = c'(\alpha). \quad (8)$$

¹¹ McAfee (1994) and Ireland (1993) develop models in which firms choose ad intensity in a first stage and prices in a second stage.

The equilibrium has some intuitive comparative statics. As the cost of advertising drops, buyers become better informed through more ads, and thus equilibrium prices also drop and converge to the Bertrand outcome; as advertising becomes more expensive, fewer buyers become informed, prices rise, and the support of the distribution of prices shrinks toward the monopoly price. More nuanced are the effects of a larger number of firms on the equilibrium, which depend on advertising costs: Eq. (8) implies that, when the marginal cost of advertising increases slowly, a larger number of firms decreases both firm and industry advertising; thus, more consumers remain uninformed and out of the market, and high prices are more likely, with lower expected consumer surplus.

Robert and Stahl (1993) extend the above setting by allowing consumers to engage in sequential search if they do not receive an ad with a suitably low price. The equilibrium exhibits price dispersion as firms randomize between charging lower advertising prices and a high unadvertised price. A noteworthy result of their model is that lowering the costs of the two information channels (i.e., firm advertising and consumer search) has different effects on firm pricing: When advertising costs tend to zero, consumers become informed almost surely and prices tend to marginal costs; however, when search costs tend to zero, consumers remain uninformed and prices are above marginal costs.

In contrast to these papers, in which sellers use ads to inform consumers, Carlin (2009), Wilson (2010) and Ellison and Wolitzky (2012), among others, develop models in which sellers practice obfuscation by increasing consumers' search costs, and thus making it more difficult for consumers to evaluate different products.¹² Specifically, Ellison and Wolitzky (2012) modifies Stahl's (1989) sequential search model (reviewed above) so that search costs are convex rather than linear in search time; hence, when a firm raises its search cost, it makes a consumer less willing to search further. More precisely, they assume that a fraction $1 - \lambda$ of consumers have zero search costs and sample all sellers, whereas the other fraction λ incur costs $g(t)$ to spend total time t searching, with $g(\cdot)$ increasing and weakly convex. Evaluation of firm i 's price requires time $\tau + t_i$, where τ is exogenous and t_i is the obfuscation level chosen by firm i . If obfuscation is costless to firms and τ is not too large, they show that it is individually rational for firms to obfuscate to and unilaterally raise consumer search costs such that consumers sample one firm only, and the equilibrium price distribution equals (3) with upper bound v (the monopoly price).

4.1.5 Product differentiation

The last group of models we discuss introduces product differentiation and consumer taste heterogeneity, thus changing the gains from trade and the search motives relative to the models described above. In homogeneous goods markets, consumers only motive to search is to find low prices. When goods are differentiated, they search for high match values, sometimes in addition to low prices. This group of models has

¹² Fisher Ellison (2016) surveys many theoretical and empirical papers on obfuscation.

developed in several directions by incorporating an increasingly rich set of features. An important distinction among these models is the type of assumption they make about consumers' ex ante knowledge (i.e., before they search) regarding sellers' differentiation.

In the simplest cases, consumers view sellers as ex ante identical, search sequentially, and thus sample sellers in random order. Wolinsky (1986) and Anderson and Renault (1999) develop workhorse equilibrium sequential-search models with horizontally differentiated products and ex ante identical sellers. They adapt the consumer sequential search behavior we discuss above to the case with horizontally differentiated products. They show that, although in equilibrium firms set identical prices, consumer search behavior disciplines sellers' prices. Consumers search to find a high match value, and this search behavior disciplines firms' pricing choices and reduces equilibrium prices. Moreover, an increase in product differentiation increases consumer search, which in turn increases competition among sellers, and thus may lower prices. This interesting result is in stark contrast to an influential literature that shows that, without search costs, an increase in product differentiation unambiguously increases prices (Perloff and Salop, 1985).

Several papers build on the framework of Wolinsky (1986) and Anderson and Renault (1999). Bar-Isaac et al. (2012) study sellers' product design, and thereby endogenize product differentiation. They find that a reduction in search costs can simultaneously explain the increase in superstar products with very broad appeal and fringe goods in the "long tail." Zhou (2014) introduces multiproduct firms and shows that in an environment with consumer search, reducing the price of one product can increase demand for a firm's other products, thereby causing products to be complements even when valuations are independent. When search costs decrease, this complementarity becomes weaker, and thus prices may increase in equilibrium.

Alternatively, consumers may be aware of some ex ante differences among sellers, because of different distributions of prices or match values, in which case their sequential search depends on this ex ante differentiation, and therefore their search is ordered or directed.¹³ Weitzman's (1979) seminal paper analyzes consumers' optimal sequential search behavior with ex ante differences in the distributions of differentiated products, thereby generalizing Kohn and Shavell's (1974) case of identically distributed products described above. Specifically, consumers know in advance the distribution $F_i(\cdot)$ of the utility of product i . Thus, generalization of the reservation value a_i (5) becomes

$$\int_{a_i}^{\bar{u}_i} (u - a_i) dF_i(u) = s. \quad (9)$$

Weitzman (1979) shows that consumers' optimal behavior is to search these products in decreasing order of values a_i , stop when the next product has a reservation

¹³ Wright et al. (2021) provide a thorough survey of the directed search literature, with a focus on markets with congestion and applications mostly outside IO (e.g., labor, housing, and monetary economics). Armstrong (2017) also surveys ordered search models in consumer product markets.

value below the best realization so far, and purchase the product with the highest realized utility. An important effect is that consumers may choose not to start their search with the product that has the highest expected value. This is because the option value of a high realization also matters, which implies that higher variance is beneficial. Moreover, consumers may optimally recall and purchase a product they evaluated in previous searches if subsequent searches prove disappointing.

Studying sellers' pricing and market equilibrium in Weitzman's (1979) general setting is complicated by the fact that many search paths could lead to purchase of the same good, and thus describing the demand of each seller becomes cumbersome. The literature has made progress by starting from simplified cases. Armstrong et al. (2009) study product prices in a market with a prominent seller that all consumers sample first before choosing to search other sellers randomly. They show that a prominent firm charges lower prices than rivals because it faces more elastic demand, and thus it is optimal for consumers to search this prominent seller first. Nevertheless, prominence decreases consumer surplus and welfare relative to a market with completely random search, because consumers search less when a firm is prominent, which increases non-prominent firms' prices and lowers aggregate output. Armstrong (2017) shows that often models with ordered search exhibit multiple equilibria: The consumer search order depends on which sellers choose lower prices, and the prices sellers choose depend on their position in the search order; moreover, random consumer search and symmetric prices, as in Wolinsky (1986) and Anderson and Renault (1999), is one equilibrium.

Armstrong and Vickers (2015) are able to characterize sellers' demand in a duopoly market in which consumers observe prices *ex ante* and search sequentially to discover match values by reformulating it as a discrete choice problem without search frictions. Armstrong (2017) and Choi et al. (2018) generalize this insight in the same setting, and thereby provide a powerful analytical tool that greatly improves the tractability of the general ordered-search model. This general reformulation removes the need to keep track of different purchase paths described above, and thus allows researchers to apply existing results on equilibrium in markets with a discrete-choice demand. Moreover, Choi et al. (2018) provide sufficient conditions that guarantee the existence and uniqueness of the market equilibrium, and show that sufficient consumer heterogeneity is crucial for the existence of a pure-strategy equilibrium. An interesting result of these models with observable prices is that equilibrium prices are usually lower when search costs are higher, because firms have greater incentives to be sampled first when search costs are higher.

Haan et al. (2018) emphasize that price observability has important consequences for market outcomes. They build a directed consumer search model in which consumers partially observe firm attributes, and show that prices and profits are lower when consumers observe prices before search than when they do not.

While the above papers consider sequential search models with differentiated products, Chade and Smith (2006) characterize the optimal portfolio choice in a simultaneous-search model with ranked stochastic options. Moraga-González et al. (2021) develop an equilibrium model with simultaneous search and differentiated

oligopolists; they show that, although prominent firms charge higher prices than non-prominent firms, both prices are below the symmetric equilibrium prices. Thus prominence increases consumer surplus, in contrast to Armstrong et al.'s (2009) results (reviewed above) in a sequential-search setting.

4.2 Empirical contributions

A large series of papers empirically study consumer search. Development of the internet and online markets, with the associated wealth of new data on the search process, has increased the attention economists have devoted to these topics. There is also a large number of contributions by researchers in information systems and marketing.

An important challenge the empirical search literature faces is distinguishing heterogeneous consumer information and/or search costs from unobserved heterogeneity in consumer preferences for differentiated sellers (for example, because of their different locations or reputations), even when they sell homogeneous items. This challenge is common to the literature on state dependence, inertia, and switching costs, as recognized at least since Heckman (1981) and Heckman and Singer (1984). The theoretical restrictions implied by sellers' optimality conditions or the institutional features of some markets have often been crucial in separately identifying consumer information and consumer preferences, as some papers reviewed below illustrate.

4.2.1 Descriptive papers

Analogous to the theoretical literature, price dispersion has been a central focus of many early empirical papers. These papers document price dispersion and examine how it varies with consumer, seller, and product characteristics. Sorensen (2000) documents large price dispersion in retail markets for prescription drugs. Interestingly, he shows that the degree of dispersion decreases for drugs that consumers purchase more frequently—a reasonable proxy for consumers' incentives to become informed. Calem and Mester (1995) present empirical evidence of consumers' limited search in the credit card market and argue that they account for the high interest rates, high profits, and imperfect competition in the industry that Ausubel (1991) previously documented. Similarly, Stango and Zinman (2016) find that heterogeneous shopping/search behavior across consumers drives the large cross-individual dispersion of interest rates among borrowers with similar observable characteristics in their credit card data. Lach (2002) analyzes the persistence of price dispersion across stores over time for four homogeneous products in Israel. He documents a large intra-distribution mobility of stores in the distribution, which is broadly consistent with the mixed-strategy equilibria of the search models discussed above. Chandra and Tappata (2011) use gasoline price data to show that retail price dispersion decreases as wholesale costs increase, which is consistent with the fact that consumers search more when prices are high, as reported by Lewis and Marvel (2011). Kaplan and Menzio (2015) use the Kilts-Nielsen Consumer Panel dataset to study the shape and structure of the distribution of prices of identical goods sold in a given market and time period; they

find that consumers who pay lower prices visit a larger number of stores instead of shopping more frequently at the same store.

Other papers test the implications of search theory for the behavior of searching agents. Among these, Genesove (1995) uses data on both accepted and rejected offers received by car sellers to study their decisions to accept or reject an offer as the outcome of a search reservation strategy, as described by Eq. (5) above.

4.2.2 Estimation of simultaneous-search models

After some of the early descriptive studies on price dispersion, several researchers have attempted to estimate search models to infer search costs, and often find relatively high monetary values of these costs. As in the theoretical literature, the earliest contributions focused on the simplest cases with homogeneous goods, and later contributions incorporated richer features, such as product differentiation and/or price advertising.

Homogeneous goods/sellers—Hong and Shum (2006) recover the distribution of search costs by estimating simultaneous- and sequential-search models with homogeneous goods. One important contribution is showing that the distribution of search costs can be recovered from price data only by exploiting the theoretical restriction whereby the observed price distribution $F_p(p)$ is the mixed strategy played by homogeneous sellers. Specifically, in their simultaneous-search model based on Burdett and Judd (1983), sellers' profits equal

$$(p - c) \sum_{k=1}^{+\infty} q_k (1 - F_p(p))^{k-1},$$

for all prices p observed in the data, where c is the marginal cost and the q_k are the proportions of consumers with k price quotes. Critically, indifference requires that all n sample prices p_i for $i = 1, \dots, n$ of the empirical distribution $\hat{F}_p(p_i)$ satisfy

$$(\bar{p} - c) q_1 = (p_i - c) \sum_{k=1}^K q_k (1 - \hat{F}_p(p_i))^{k-1}, \quad (10)$$

where \bar{p} is the highest sample price and K is the maximum number of firms from which a consumer obtains a price quote. The indifference condition (10) holds at each sample price $p_i < \bar{p}$, thereby yielding a system of $n - 1$ equations that can be solved to estimate the $K \leq n - 1$ unknowns $\{r, q_1, \dots, q_{K-1}\}$.

In turn, the proportions q_k satisfy

$$\begin{aligned} q_1 &= 1 - F_s(\Delta_1), \\ q_k &= F_s(\Delta_{k-1}) - F_s(\Delta_k) \text{ for } k > 1, \end{aligned}$$

where $F_s(\cdot)$ is the distribution of search costs and Δ_k equates the search cost with the expected savings of an additional search:

$$\Delta_k = Ep_{1,k} - Ep_{1,k+1}, \quad (11)$$

where $Ep_{1,k}$ denotes the lowest price out of k draws from the equilibrium price distribution $F_p(p)$. Eq. (11) generalizes Eq. (4) to the case of heterogeneous search costs: Consumers with search costs above Δ_1 obtain only one quote, and consumers with search costs between Δ_{k-1} and Δ_k obtain exactly k quotes. Hence, the distribution $F_s(\cdot)$ can be nonparametrically estimated by using the estimated values of $\{q_1, \dots, q_{K-1}\}$.¹⁴

Hong and Shum (2006) illustrate their methodology by estimating the model using online prices for popular textbooks, and find that, for several textbooks, more than 60 percent of consumers search one seller only, and thus have nontrivial search costs.

Heterogeneous goods/sellers—Wildenbeest (2011) extends the simultaneous-search framework by introducing vertically differentiated sellers, and shows that the model is identified from price data only. The model is estimated on U.K. data from grocery items. The estimates suggest that seller heterogeneity accounts for approximately 60 percent the observed price variation in these data, and search frictions account for the remaining 40 percent.

De los Santos et al. (2012) exploit data on web browsing and consumer purchases to argue that simultaneous, fixed-sample-size search provides a more accurate description of search behavior than sequential search, based on two patterns in their data: (1) consumers frequently recall a product that they previously sampled; and (2) consumers' decision to continue searching does not depend on the realized prices of their previous searches, whereas consumers searching sequentially should continue searching when they sample a relatively high price. Hence, they use the simultaneous, fixed-sample-size search model with heterogeneous sellers to estimate demand elasticities of online bookstores, and find larger elasticities than in standard models that assume consumers consider all firms.

Woodward and Hall (2012) document a large price dispersion in broker fees in the U.S. mortgage market. They consider a framework in which each borrower seeks initial quotes from N brokers with heterogeneous costs and uses these quotes to negotiate until the process ends with one quote no other broker is willing to beat, similar to a second-price auction. They argue that most borrowers overpay for brokers' services because they seek only a few quotes.

Honka (2014) estimates a simultaneous-search model of the auto insurance industry by exploiting data on individual-specific consideration sets, and finds that large search costs account for the high retention rate in the industry.

Sellers' non-pricing behavior—Ellison and Ellison (2009) provide extensive empirical evidence on obfuscation among online retailers. They show that although the internet generally facilitates search, it also allows firms to make consumer search more difficult. They argue that this sellers' practice is a main reason for why the internet had a smaller effect on the reduction of transaction prices and of their dispersion than initially predicted.

¹⁴ Moraga-González et al. (2013) and Sanches et al. (2018) extend this methodology.

Galenianos and Gavazza (2020) develop and calibrate a nonsequential search model of the credit card market with differentiated products that incorporates the main feature of Butters's (1977) seminal contribution (reviewed above)—i.e., sellers sending price offers to prospective buyers. They document that many households receive many preapproved credit card offers, but nevertheless borrow at high and very dispersed interest rates. Their calibration suggests that most borrowers examine few of the offers they receive, thereby forgoing cards with low interest rates and high non-price benefits, and account for an inelastic demand and high lender profits.

4.2.3 Estimation of sequential-search models

Random search—Sirri and Tufano (1998) argue that search costs are an important determinant of investors' mutual fund choice. Hortaçsu and Syverson (2004) build on this insight and estimate a sequential, random-search model with vertically differentiated goods using data on S&P 500 index mutual funds. They document large fee dispersion across seemingly homogeneous financial products such as index funds. They aim to understand the sources of this dispersion, distinguishing between heterogeneous search costs and vertical differences among funds.

In their model, all consumers combine fees p_j and fund characteristics W_j , which are observable to consumers and the econometrician, and fund characteristics ξ_j , which are observable to consumers but not to the econometrician, of mutual fund j according to the following utility function:

$$u_j = W_j\beta - p_j + \xi_j, \quad (12)$$

where the parameters β are common across all consumers.

The heterogeneity across consumers lies in their search costs c , distributed according to the cumulative distribution function $G(\cdot)$. Consumers search funds sequentially, sampling fund k with probability ρ_k on each search. In the typical random search model, funds have equal sampling probabilities, $\rho_k = \frac{1}{N}$, but Hortaçsu and Syverson also consider the case in which these probabilities can vary across funds, depending on some observable fund characteristics Z_k , according to the functional form $\rho_k = \frac{Z_k^\alpha}{\sum_{j=1}^N Z_j^\alpha}$. Hence, consumers stop searching when they find a fund yielding utility that exceeds their reservation utility, as in (5). This search rule yields the lowest search costs c_j of any consumer who purchases fund j :

$$c_j = \sum_{k=j}^N \rho_k (u_k - u_j), \quad (13)$$

where funds are ordered according to their utilities, $u_1 < \dots < u_N$, and consumers know the empirical cumulative distribution function of utilities $\frac{1}{N} \sum_{j=1}^N \mathbf{I}[u_j \leq u]$.

Based on these cutoff values (13), fund market shares can be defined recursively:

$$q_1 = \rho_1 (1 - G(c_1)), \quad (14)$$

$$q_j = \rho_j \left[1 + \frac{\rho_1 G(c_1)}{1 - \rho_1} + \sum_{k=2}^{j-1} \frac{\rho_k G(c_k)}{(1 - \rho_1 - \dots - \rho_{k-1})(1 - \rho_1 - \dots - \rho_k)} - \frac{G(c_j)}{(1 - \rho_1 - \dots - \rho_{j-1})} \right] \text{ for } j > 1. \quad (15)$$

To understand the market share q_1 in (14), notice that the only consumers who buy the lowest-utility fund u_1 are those who draw it first, which happens with probability ρ_1 , and have search costs higher than c_1 . The market share q_j of good j in (15) follows a similar, though slightly more complex, intuition.

Hortaçsu and Syverson further specify a supply side with optimal pricing. Given market size S , fund j profits equal:

$$\Pi_j = S q_j(p, W)(p_j - mc_j).$$

Hence, optimal fund fees satisfy:

$$q_j(p, W) + (p_j - mc_j) \frac{\partial q_j(p, W)}{\partial p_j} = 0. \quad (16)$$

Hortaçsu and Syverson estimate several versions of their model, depending on whether sampling probabilities ρ_j vary across funds and whether funds' nonfinancial characteristics affect consumers' utility—a model with both unequal sampling probabilities and fund nonfinancial characteristics is not identified with their data.

In the version with heterogeneous sampling probabilities and without nonfinancial characteristics, funds can be ranked based on their prices. With parametric assumptions on the search cost distribution $G(\cdot)$, the $N - 1$ market share equations (15) and the N pricing equations (16) can be combined to estimate the parameters of the sampling probabilities and of the search cost distribution, as well as marginal costs, using nonlinear least squares. The estimation of this version fits the data very well, and reveal that substantial asymmetries in funds' sampling probabilities are necessary to explain the data and that search costs are small: The median search cost equals \$5 per \$10,000 invested.

In the version with equal sampling probabilities and with fund nonfinancial characteristics, funds can be ranked based on their market shares. Assuming that marginal costs are known, the market share equations (15) and the pricing equations (16) nonparametrically identify the distribution of search costs and fund utilities. In turn, using (12), consumers' preference parameters about fund nonfinancial characteristics can be estimated. Estimation of this version reveals that these nonfinancial characteristics of funds are quite important for consumers, and search costs are again small.

Ordered search—Several empirical papers use an ordered-search framework to study markets for consumer products. Some of these papers focus only on the consumer side by exploiting data on internet searches. Among these, Kim et al. (2010)

and Kim et al. (2017) map Amazon.com's view rank data—which contain information on which other products consumers viewed in the same session—onto a sequential-search model. Jolivet and Turon (2019) exploit inequalities in Weitzman's (1979) search-and-purchase problem to set-identify preferences and search costs using data from CDs sold on a French internet platform; their estimates imply quite heterogeneous and sometimes large search costs.

Ellison and Ellison (2018) focus on sellers' pricing in markets in which match value is important, without rationalizing the consumer search process. They develop and estimate a model in which sellers with heterogeneous customer arrival rates set prices to sell to high match-value customers. They show that improvements in internet search technologies may increase price dispersion, seller profits, and consumer surplus in the market for used books.

Other papers consider markets in equilibrium (i.e., demand and supply). Allen et al. (2019) develop and estimate a sequential, ordered-search model of the Canadian mortgage market. They document that most mortgage borrowers exhibit loyalty to the bank with which they have an established relationship through a deposit account, and they incorporate this relevant starting condition in their sequential-search model. They estimate substantial consumer search costs and brand loyalty, which confer to banks a quasi-monopoly over their customers, with sizable effects on their profits. Using similar Canadian mortgage data, Allen et al. (2014) document that bank mergers reduce the dispersion of prices because they reduce the bargaining leverage of consumers with the lowest search costs.

Moraga-González et al. (2018) exploit the theoretical insights of Armstrong (2017) and Choi et al. (2018), which reformulate the optimal sequential consumer search as a discrete-choice demand, to study market equilibrium in the Dutch automobile industry. Their estimates of search costs are quite large and indicate that consumers consider a few sellers only. Moreover, they show that own-price elasticities are lower and markups are higher than if they assume that consumers have full information.

Models with additional features—Galenianos and Gavazza (2017) estimate a dynamic equilibrium search model of the retail market for illicit drugs that builds on Galenianos et al. (2012) to understand how the observability of quality at the time of the transaction, as well as repeat purchases, account for the dispersion of gains from trade across transactions. Buyers search for high-quality drugs, but they determine drugs' quality/purity only after consuming them (i.e., they are experience goods). Hence, sellers cheat first-time buyers or offer higher-quality drugs to induce buyers to purchase from them again. In equilibrium, a distribution of qualities persists, because of traditional search frictions and the imperfect observability of drug quality. Their estimated model implies that if drug quality/purity were observable, the average purity of drugs would increase and the dispersion would decrease substantially, which implies that imperfect observability of quality is the main reason for its dispersion.

Coey et al. (2020) build a model of consumer search in which consumers are willing to pay more as a deadline approaches, which demonstrates that buyers' reser-

vation prices increase over time, and thereby account for price dispersion and the coexistence of different selling mechanisms, such as auctions and posted prices. They estimate their model using data on online retail sales from eBay and show how, in the presence of deadlines, an increase in platform fees can increase the share of listings with posted prices relative to those with auctions, consistent with the trend documented by Einav et al. (2018).

4.3 The role of intermediaries in search markets

As in the case of other frictions discussed previously, the presence of search frictions allows for the possibility that third-party agents intermediate transactions between buyers and sellers. Such intermediaries may provide value to consumers by reducing search costs. Intermediaries may take several different forms, depending on the exact type of frictions and the precise nature of the gains from trade—i.e., whether consumers are searching for prices or for high match values. Nevertheless, an important distinction is whether intermediaries take ownership of the goods and build inventories—like dealers and retailers—or serve as information intermediaries that help consumers in their search process—e.g., brokers, real estate agents, financial advisors, platforms, or recommender systems.¹⁵

In some markets, the distinction between the two types of intermediaries is often not clear cut. For example, in the case of over-the-counter bond markets, most dealers perform both principal trades and agency trades, with no obvious difference between the two types of trades for buyers and sellers. In other markets, information intermediaries may have greater experience and knowledge than some buyers about specific characteristics of different products and/or sellers, in which case intermediation is frequently bundled with some advice, which may lead to conflicts of interest.

Our review of intermediaries in search markets is selective, and we refer interested readers to the chapter by Jullien et al. (2021) in this Handbook for an insightful discussion of the literature on platform intermediaries.

4.3.1 Theoretical contributions

Physical intermediaries—Rubinstein and Wolinsky (1987) provide a framework to analyze the activity of intermediaries in search markets by building a dynamic two-sided search model in which intermediaries shorten the time agents have to wait to find a trading counterparty. The paper is noteworthy because, among other reasons, it highlights that intermediation may not benefit buyers and sellers, even if the intermediaries have a trading technology superior to that of buyers and sellers.

Specifically, consider a market for an indivisible homogeneous good, with three groups of agents: buyers with valuation one, sellers with valuation zero, and interme-

¹⁵ Hagiu and Wright (2015b) and Hagiu and Wright (2015a) analyze intermediaries' organizational choice. Similarly, Rust and Hall (2003) develop a model in which the microstructure of trade in an asset is endogenously determined, and Spulber (2002) studies how the organization of the intermediation market affects buyers' and sellers' incentives to invest.

diaries with valuation zero. Trades take place over time, with all agents discounting the future at the common rate δ . A constant flow of e buyers and e sellers enter the market in every period, and leave the market after they trade; intermediaries stay in the market forever and, for simplicity, each intermediary can have at most one unit of inventory.

A steady-state equilibrium features constant measures L_B of buyers, L_S of sellers, L_N of intermediaries with the good, and L_M of intermediaries without the good. In each period, the pairwise meeting process between agents is described by probabilities $\alpha_B(L_B, L_S)$, $\alpha_S(L_S, L_B)$, $\beta_B(L_B, L_M)$, $\beta_S(L_S, L_N)$, $\gamma_N(L_N, L_S)$, and $\gamma_M(L_M, L_B)$, where, for example, α_B is the probability that an individual buyer directly meets a seller, and similarly for the other probabilities. In a symmetric case, $L_B = L_S = L$ and $L_M = L_N = K$, and the above probabilities simplify $\alpha = \alpha_S = \alpha_B$, $\beta = \beta_B = \beta_S$, and $\gamma = \gamma_M = \gamma_N$; in steady state, these matching probabilities require

$$\alpha(L, L)L + \gamma(K, L)K = e,$$

where $\alpha(L, L)L$ is the aggregate number of direct trades between buyers and sellers, $\gamma(K, L)K$ is the aggregate number of intermediated trades, and the equality follows from the steady-state requirement that the number of buyers or sellers exiting the market must equal the number of entrants e .

When agents meet, they trade at the price P_{ij} , paid by agent j to agent i , determined by the Nash (1950) bargaining solution with equal weights.

Two possible cases arise in equilibrium:

1. If intermediaries are not active, i.e., $K = 0$, then $\gamma(K, L) \leq \alpha(L, L)$ and $P_{SB} = 1/2$.
2. If intermediaries are active, i.e., $K > 0$, then $\gamma(K, L) \geq \alpha(L, L)$ and

$$P_{SB} = (1 + V_S - V_B) / 2,$$

$$P_{SN} = (V_M + V_S - V_N) / 2,$$

$$P_{MB} = (1 - V_N - V_B + V_M) / 2,$$

where the values V_i are the nonnegative solutions to the system:

$$V_S = \delta (\alpha (1 + V_S - V_B) / 2 + \beta (V_M + V_S - V_N) / 2 + (1 - \alpha - \beta) V_S), \quad (17)$$

$$V_B = \delta (\alpha (1 - V_S + V_B) / 2 + \beta (1 + V_N + V_B - V_M) / 2 + (1 - \alpha - \beta) V_B), \quad (18)$$

$$V_M = \delta (\gamma (1 + V_N - V_B + V_M) / 2 + (1 - \gamma) V_M), \quad (19)$$

$$V_N = \delta (\gamma (V_M - V_S + V_N) / 2 + (1 - \gamma) V_N). \quad (20)$$

To understand these equations, consider that the value of an unmatched seller in (17) equals the expected discounted value of meeting a buyer, which happens with probability α , and trading at price P_{SB} ; or meeting with an intermediary, which happens with probability β , and trading at price P_{SN} ; or remaining unmatched, which will happen with probability $1 - \alpha - \beta$. Similarly for the other Eqs. (18)–(20).

The two possible cases characterized above make it clear that a market with active intermediation requires that intermediaries are at least as efficient as regular buyers and sellers in meeting counterparties. This is intuitive, since the intermediaries can profit only if they reduce the trading delays of buyers and sellers, which is the only friction/inefficiency in the model.

The equilibrium with intermediaries displays some interesting features. First, the presence of intermediaries does not necessarily benefit buyers and sellers. Specifically, the difference between $V_B + V_S$ in a steady-state equilibrium with intermediaries and its value in a state-state equilibrium without intermediaries depends on the specifics of the meeting probabilities, and thus intermediaries may reduce buyers' and sellers' joint welfare.

Second, an asymmetry between buyers and sellers arises when prices are determined by the Nash bargaining solution, as in the above case; notably $V_S < V_B$ and $P_{SB} = 1/2$ —i.e., the equilibrium outcome favors buyers. Moreover, the difference $V_S - V_B$ grows larger as β , the probability that buyers or sellers trade with an intermediary, grows. The reason for this asymmetry is that the price the intermediary paid to the seller is sunk when the intermediary resells the good to the buyer, and thus buyers can extract a share of the whole surplus larger than that of sellers in the bargaining. This effect is stronger the larger the share of intermediated trades. Of course, different pricing protocols may restore symmetry between buyers and sellers.

Information intermediaries—Baye and Morgan (2001) extend Varian's (1980) framework to study the interaction between the information clearinghouse's fee-setting decision and the competitiveness of the homogeneous product market it serves. The clearinghouse charges fees to firms to post prices on its site, and to consumers to access the list of posted prices. The paper shows that introducing this market for information may not increase welfare.

Several other papers analyze how platforms and recommender systems, such as price-comparison websites, affect the ease of consumer search on the platform and market efficiency, and highlight several reasons these intermediaries may not offer the best matches (e.g., Hagi and Jullien, 2011; De Cornière, 2016; Wang and Wright, 2020). More generally, a related literature considers the effects of advisors who have information superior to that of buyers and/or sellers. A common theme is that conflicts of interest, often because of commissions, may distort brokers' advice away from the first-best. These conflicts have been analyzed extensively in the case of financial advice, but they arise in many other professional services, such as real estate brokers, doctors and health care specialists, and legal services, among others. Inderst and Ottaviani (2012) provide an insightful overview of the main economic issues emphasized by the literature, as well as related policy interventions.

4.3.2 Empirical contributions

The empirical literature has considered many markets, each with its own regulatory and institutional characteristics. The richness of these studies enables useful comparisons. Several empirical studies support the theoretical finding that intermediaries

do not always increase overall welfare, because they may exploit their expertise to capture surplus.

Physical intermediaries—Gavazza (2016) develops a search-and-bargaining model of a decentralized market that builds on Rubinstein and Wolinsky's (1987) framework described above, and estimates it using business-aircraft data. The estimates imply that trading frictions create large welfare losses and that dealers play an important role in reducing frictions. However, dealers reduce aggregate welfare because they extract almost all surplus that they create, and they impose a negative externality on buyers and sellers by decreasing the number of direct transactions among them.

Similarly, Farboodi et al. (2019) argue that rent-extraction motives may account for intermediation in asset/durable goods markets, and thereby lower social welfare.

Information intermediaries—Several papers document the positive effects of information intermediaries such as price-comparison websites. Brown and Goolsbee (2002) show that the emergence of internet comparison shopping sites reduced the transaction prices of life insurance policies by reducing search costs for consumers. Baye et al. (2004) document that price dispersion is greater when few sellers list products on a price comparison site than when many sellers do. Grennan and Swanson (2020) find that hospitals pay lower prices for medical devices after joining a price-benchmarking database. Brown (2019) exploits the introduction of a price-transparency website that informed a subset of consumers of medical imaging services. He builds an empirical model of demand and supply that incorporates consumers' limited information about prices, and finds that equilibrium prices would be lower if all consumers were fully informed.

Salz (2021) finds positive effects of search brokers in business-to-business transactions. He develops a search model in which intermediaries affect buyer welfare both directly, by reducing the expenses of buyers with high search cost, and indirectly, through a search externality that affects the prices paid by buyers who do not use intermediaries, consistent with Varian's (1980) model reviewed above. He estimates the model using data from New York City's trade waste market, in which buyers can either search by themselves or through brokers, and shows that intermediaries benefit buyers in both the brokered market and the search market.

Donna et al. (2021) develop an equilibrium model with costly consumer search and two distribution channels—vertically integrated manufacturers and intermediaries—in which intermediaries reduce search frictions and provide consulting services. They estimate their model using data from the outdoor advertising industry and find that intermediaries increase welfare.

In contrast to the previous papers that document the positive effects of intermediaries, several papers demonstrate that differential expertise between brokers and consumers, as well as brokers' commissions, generate conflicts of interest that reduce consumers' benefits from intermediaries. For example, Levitt and Syverson (2008) find that real estate agents convince their clients to sell their houses too quickly, which leads them to accept low prices. Bergstresser et al. (2008) analyze the role of financial advisers in the mutual fund industry and find that they do not deliver

substantial tangible benefits to their clients. Egan (2019) studies how brokers distort household investment decisions in the convertible bond market. He documents that households often purchase expensive bonds when otherwise identical cheap bonds are also available, and that brokers earn greater commissions on more expensive bonds. He develops and estimates a search model with intermediaries, and finds that costly search is a key friction whose effects are compounded by brokers' incentives to recommend high-fee inferior products to consumers. Robles-Garcia (2019) provides a more nuanced analysis of the effects of brokers in the U.K. mortgage market: She documents that, although the commissions lenders pay to brokers distort brokers' advice to borrowers, brokers do increase the sales of lower-cost lenders, and thereby increase the overall competitiveness of the market.

Finally, a few papers study platforms' and information intermediaries' trade-offs in helping consumers. Dinerstein et al. (2018) examine the design of e-commerce platforms that guide consumers who search for products with both high match values and low prices. They exploit a change in eBay's marketplace design that narrowed the number of listings displayed to consumers and prioritized low-price sellers. They find that the impact of the platform design varies considerably across product categories, depending on their differentiation, with more homogeneous products experiencing the largest declines in posted and transaction prices.

Bar-Isaac and Gavazza (2015) studies how parties design contracts to provide intermediaries with incentives in markets in which match values are important by empirically analyzing contractual agreements between landlords and real estate brokers in the Manhattan residential rental market.

5 Matching frictions

Closely connected to search frictions are matching frictions that potentially make it difficult or altogether prevent profitable meetings between partners who would otherwise be able to exploit gains from trade. The chapter by Agarwal and Budish (2021) in Volume 5 of the Handbook discusses studies that explore the role of centralized matching processes, as in education or residency programs. There is a large literature in the labor/macro field that considers how matching frictions in decentralized markets may affect allocations, and then discusses how such frictions should inform labor market policies (for a survey, see Petrongolo and Pissarides, 2001). Most of this literature focuses on macro aggregates for the labor market and assumes a reduced-form process that is encoded in an aggregate matching function that specifies how the number of matches depends on the number of participants on both sides of the market. In the labor application, the matching function relates the flow of new hires to the stocks of vacancies and unemployment (the number of searching workers). The matching function is therefore a central ingredient of search and matching models (Mortensen and Pissarides, 1994). Such matching functions are typically assumed to be homogeneous of degree one—i.e., constant returns to scale. There have been conceptual micro foundations of specific matching functions (e.g., Butters, 1977; Burdett

et al., 2001). However, matching frictions remain a somewhat abstract concept in much of this literature. Petrongolo and Pissarides (2001) justifies the reduced-form representation in matching functions as follows:

Frictions derive from information imperfections about potential trading partners, heterogeneities, the absence of perfect insurance markets, slow mobility, congestion from large numbers, and other similar factors. Modeling each one of these explicitly would introduce intractable complexities in macroeconomic models.

We focus on a recent literature that has attempted to make progress in such explicit modeling by focusing on specific markets and spatial frictions in order to understand in more detail the role of matching frictions. This literature also seeks to understand how matching frictions affect the organization of these markets and the emergence of platforms or intermediaries. Although the markets studied are quite specific, it is valuable to have instances of more microfounded matching processes that can be useful starting points for analyses with broader scope. Lagos (2000), for instance, states that

The results of policy experiments based on models that assume an exogenous meeting process are likely to be misleading if the random search assumption is not a good characterization of the agents' underlying search behavior. If agents are able to direct their search, then the matching function is an equilibrium object and is sensitive to policy.

Furthermore, as we will see below, there are policy experiments that explicitly target the matching process to attempt to reduce frictions that are hard to conceptualize through an aggregate exogenous matching function.

Specifically, this literature focuses on spatial search frictions. The need to coordinate trade in a physical space is a major source of search frictions. In equilibrium, the organization of the market reflects the nature of spatial frictions. The literature attempts to answer the following questions. How are supply and demand determined in spatial search markets? What are the welfare consequences of search frictions? What can be done to reduce these frictions?

Taxi and ride-hailing markets (and other transportation markets) are interesting prototypes in which one can study the importance of these frictions, partly because they feature many independent “contractors” and few natural entry barriers.

Because of regulations, the taxi market traditionally features limited time variation in supply and no spatial flexibility. This inflexibility, together with the limited supply of medallions and the presence of matching (search) frictions, implies that regular (and predictable) patterns of variation in demand for rides during the day (e.g., rush hours) lead to large fluctuations in costly delays for matches between passengers and taxis. Drivers' earnings and the number of active taxis vary during the day depending on how long drivers need to spend searching for passengers. Fréchet et al. (2019) report that in New York City, the fraction of time taxis drive empty while searching for passengers ranges between 30 percent and 70 percent of active time, depending on time of day. Passengers also wait to obtain a taxi, and this wait time varies

during the day. Although in textbook markets, prices serve to balance demand and supply, in the taxi market this balancing process is a function of passenger wait times and taxi search times. In contrast to prices (which are a transfer), these service delays are wasteful. Part of taxis' idle time is due to drivers' uncertainty about where they might find a waiting passenger. Determination of the matching process between taxis and passengers is, therefore, a key feature of the analysis. Fréchette et al. also report direct evidence of search frictions. They examine the degree of predictability of the order in which empty taxis find passengers in a given area: Absent search frictions, one would expect the pickup order of taxis in a specific area to be highly correlated with their drop-off order. In other words, focusing on a narrow enough area with uniform demand conditions, the i -th cab to drop off a passenger in this area should also be the i -th to subsequently pick up the next passenger. Although the drop-off sequence and pickup sequence are positively correlated (the correlation coefficient is 0.51), a large amount of dispersion is still present. Fréchette et al. also compare the time empty taxis spent traveling between a drop-off point and their subsequent pickup location with the time reported on Google maps for the same time/day. They find that taxis spend on average 220 percent more time traveling between these points, which suggest that part of the travel time is wasted in searching for customers.

Lagos (2000) develops a model of a taxi-like market that provides a micro foundation for a (constant returns) matching function. In his model there are n geographic locations. In each location an exogenous number of passengers who wish to travel to other locations. Revenues for travel between locations are heterogeneous, distinguishing them in terms of their inherent desirability to taxis. Fares are fixed do not respond to the extent of excess demand/supply. Taxis choose where to locate depending on exogenous revenues and on endogenous matching probabilities. If d_i is the number of potential passengers in location i and c_i is the number of taxis in that location, Lagos assumes that the matching probability is $\min\left\{1, \frac{d_i}{c_i}\right\}$. This matching process therefore assumes that there is no friction in matching within a location. The aggregate matching process can generate frictions, in the sense that there are unrealized matches. The degree of frictions depends on the aggregate tightness of the market and becomes larger, roughly, the larger the degree of excess demand for taxis. The reason for this asymmetry between demand and supply is that, in this model, taxis are the active searching agents, whereas passengers are passive.

Lagos (2003) calibrates a model based on Lagos (2000) to data from taxicabs in New York City. He uses the model to empirically analyze the effect of taxi fares and medallion counts on matching rates and medallion prices in Manhattan. Using the model, he quantifies the impact of policies that increase fares and the number of medallions. Partly because of data limitations, he assumes that all medallions are active throughout the day, and thus does not model the labor-supply decision or allows demand respond to wait time.

Lagos (2003) does not have data on hourly or daily decisions by taxi drivers, and instead relies on some aggregate statistics that were available at the time. Detailed data on the New York City taxi market became publicly available in 2013 following a Freedom of Information Act request to the Taxi and Limousine Commission. These

data include every trip of the yellow cab fleet in this time span. The data entry of a trip includes the fare, tip, distance, and duration, as well as the geospatial start and end points of the trip. These data have triggered a burst of attention to this market for a variety of applications. Many interesting papers focus on individual decisions by taxi drivers (e.g., Thakral and Tô, 2021). Our focus here is on the market equilibrium.

Buchholz (2021) draws elements from Lagos's (2003) search model, but makes several changes to reflect the real-world search and matching process. Specifically, he adds nonstationary dynamics, a more realistic and flexible spatial structure, stochastic and price-sensitive demand, and heterogeneity in the matching process across different locations. Regarding the matching process, in contrast to Lagos (2003), Buchholz also allows for frictions at the neighborhood level. He divides Manhattan (plus airports) into 39 locations that together cover 94 percent of all rides. He considers the day shift between 7 am and 4 pm and divides it into 5-minute intervals. In every location i and time t , there is a Poisson arrival rate of customers λ_i^t . Empty taxis choose locations by taking into account the possibly different continuation values of those locations. He specifies a location-specific matching process using Butters's (1977) urn-ball matching function (see also Burdett et al., 2001) as follows: In a location i , given a number of passengers u_i , and vacant taxis v_i , the number of matches is given by $v_i \left(1 - \left(1 - \frac{1}{v_i}\right)^{u_i}\right)$. This matching function generates approximately constant returns except for low levels of inputs. Therefore, Buchholz specifies a matching process that is random within a neighborhood, but directed across neighborhoods. The key counterfactual Buchholz considers is designed to quantify the degree to which allocations can be improved by more flexible fares that condition on locations and time of day and are allowed to be nonlinear in distance. Note that such fares would fully restore efficiency in Lagos's (2003) model. He finds that such flexible fares lead to an increase in surplus. These flexible fares do not respond to unpredicted changes in demand. Therefore, they are not equivalent to fully flexible pricing that also allows for surge pricing to cope with unexpected demand. However, they capture many of the benefits from such surge pricing, partly because, in this market, supply cannot respond very quickly to fully unexpected demand shocks.

Fréchette et al. (2019) build a model of the supply side (taxi activity), the demand side (passengers), and the matching process between passengers and taxis in order to evaluate a number of frictions in this market. Most taxi markets, before the entry of companies such as Uber and Lyft, entailed heavy regulations on entry. Given such restrictions, if no other frictions were present, one might therefore expect all taxis to be active at least during the daytime. However, in New York City, activity is often well below capacity, which highlights the importance of understanding the labor-supply decisions of taxi drivers. Labor supply cannot instantaneously adjust to market conditions, because drivers operate on a two-shift system, which leads to shift indivisibilities in labor supply.

In the model, drivers make daily entry (starting) and hourly stopping decisions. Licenses to operate a taxi (medallions) are scarce, so entry is only possible for inactive medallions. Hourly profits are determined by the number of matches between searching taxis and waiting passengers. *Ceteris paribus*, increasing the number of

taxi increases the search time for a driver to find the next passenger and reduces expected hourly earnings. The number of active taxis is determined endogenously as part of the competitive equilibrium in this market.

The model delivers responses to earnings shocks that strongly depend on their timing during the day. Specifically, Fréchet et al. compare the response to a uniform shock throughout the day with the response to an hour-specific shock for each hour. The former results in an elasticity estimate of 1.8, and the latter results in estimates that vary from 0.9 to 2.6 depending on the hour, with beginning- and end-of-shift elasticities that are lower than those for the middle of the shift.

On the demand side, Fréchet et al. face a challenge. Although they observe the number of matches, neither passengers' wait time nor the number of hailing passengers is directly observable in the data. However, they are able to recover these variables by using information on other observables and the nature of the matching process. An innovation of this study is its development of an explicit microfounded model of the matching process. They build a description of the geographical nature of the matching process and then recover, via simulation, a numerical representation of the matching function. They use this information to deduce both the number of passengers and the wait time. They allow for heterogeneity among neighborhoods by dividing Manhattan into eight areas that differ (mostly exogenously) according to the arrival rate of passengers and the desired destinations that determine cab transitions across areas. Therefore, although their approach allows them to account for some of the heterogeneity in locations, they do not consider directed search, as in Buchholz (2021). Interestingly, the matching process obtained by Fréchet et al. displays increasing returns to scale: If the number of passengers and the number of taxis double, the number of matches more than doubles, and search time and wait time both drop. They also find that these returns depend on the level of activity. At low levels of activity, such as during the nighttime, returns to scale are substantial. However, for daytime levels of activity, returns to scale become essentially constant. As we discuss below, this feature of the matching process is important for understanding the consequences of entry by platforms such as Uber.

Brancaccio et al. (2020a) examine the shipping market, which constitutes 80 percent of global trade by volume. They argue that bulk shipping displays some characteristics that turn out to be surprisingly similar to those of the taxi market. One important difference is that prices are flexible and largely determined through bilateral agreements. They document large price differentials across space—e.g. the shipping price from Australia to China is about 30 percent more expensive than vice versa. Similar to the taxi market, a large number of ships is “unemployed” while searching for their next cargo. In particular, they find that 42 percent of ships currently in transit are without cargo. In their model, spatial equilibrium determines world trade, and trade costs are endogenous and determined jointly with trade flows. In contrast, standard trade models assume that trade costs are exogenous and determine trade flows. Brancaccio et al. nonparametrically estimate a matching function that is assumed to be constant returns. They key counterfactual they study in the paper is the opening of the Northwest Passage. They find that the primary beneficiary

is North America, whereas East Asian countries are mostly unaffected. Interestingly, other distant countries that are not directly affected can suffer from such a change because of an endogenous rebalancing of shipping and a resulting endogenous increase in shipping costs.

Brancaccio et al. (2020b) develop this framework further and obtain theoretical results regarding the nature of inefficiencies in this type of market. They argue that in this environment with dynamic spatial frictions, there are two types of inefficiencies. The first has to do with market thickness and is a familiar aspect of the matching literature that led to the Hosios (1990) condition for efficiency: When choosing whether to search, agents do not take into account the fact that their presence affects the matching probabilities faced by other agents, in both the same and opposite sides of the market. The second friction is a pooling friction: When customers and ships match and move on to a new destination, customers may not take into account the impact of their destination choice on the distribution of carriers over space.¹⁶ They use the same data as Brancaccio et al. (2020a) to quantify these frictions. The equilibrium allocation leads to a loss of 6 percent relative to the constrained-efficient benchmark (in which the planner faces the same frictions but resolves the misallocations due to externalities), and 14 percent relative to the frictionless benchmark.

5.1 The role of intermediaries in matching markets

As described above, many matching markets, such as education, are characterized by centralized processes, in which the matching algorithms can be thought of as intermediaries. Such centralized approaches are meant to remedy market failures in decentralized matching that have been the focus of some previous work. The chapter by Agarwal and Budish (2021) in Volume 5 of the Handbook further discusses these issues.

Some of the work on spatial matching discussed above also focuses on the role of intermediaries or platforms in these markets. Fréchette et al. (2019) (discussed above) evaluate counterfactuals that concern the magnitude of matching frictions and possible ways to reduce them. They consider an alternative matching technology in line with the dispatch system of the Uber platform and other ride-hailing services. The first extreme case they consider is one in which there is a centralized dispatcher for the entire fleet. They show large gains for both sides of the market due to reductions in wait times for both passengers and taxis. They then move on to the case in which the dispatch platform covers only part of the market, with the rest functioning according to the traditional street-hailing system. They show that market segmentation on different platforms creates an inefficiency due to a reduction in market thickness for both platforms. Partial coverage by a dispatcher has two effects compared with a market with no dispatcher: On the one hand, as in the first counterfactual, the partial dispatcher is a more effective platform for the taxis and passengers served by it; on the

¹⁶ This type of friction arises also in labor markets with heterogeneous jobs, as Acemoglu (2001) highlights.

other hand, segmentation of the market makes both segments thinner, with the consequence of longer average distances between a random taxi and a random passenger. In the case in which there is an equal number of potential taxis divided between centralized dispatch and decentralized platforms, they find that the second effect dominates and therefore aggregate outcomes are worse than in the baseline case. Interestingly, the effects are quite different during the daytime relative to nighttime hours, which reflects the importance of the initial thickness of the baseline market environment. Finally, an advantage of their framework is that space is meaningful. They use this aspect to study the effects of density. They simulate a city that is otherwise identical to Manhattan, but one-third as dense; that is, the same number of potential passengers is spread over a larger territory covered by the same number of potential taxis. They find dramatic losses in efficiency due to lower density. However, these inefficiencies are substantially alleviated by a dispatch platform whose performance is (comparatively) much better in a less dense environment.

Castillo et al. (2018) build on earlier work by Arnott (1996) to argue that surge pricing has the virtue of solving a potential matching failure that they call the wild goose chase. Castillo et al. state:

Absent surge pricing, when drivers are scarce relative to demand, drivers are quickly occupied and thus idle drivers are spread thinly throughout a geographic region, forcing matches between drivers and passengers that are on average far away from each other. Cars are thus sent on a wild goose chase [...] to pick up distant customers, wasting drivers' time and reducing earnings. This decreases the number of available cars both directly by occupying cars and indirectly as cars exit in the face of reduced earnings, exacerbating the problem.

Castillo et al. also argue that this failure is more likely to plague ride-hailing platforms than taxis. This phenomenon is particularly interesting, because it depends on a combination of the absence of dynamic pricing and the matching process specific to ride-hailing taxis (i.e., not street-hailing).

Castillo (2020) develops an empirical model of ride hailing to evaluate the distributional effects of surge pricing. As in Fréchette et al. (2019), the model incorporates a demand side that responds to (expected) wait times and prices, a supply side that responds to expected earnings, and a matching process that is specific to the ride-hail platforms and that, given the number of passengers and drivers, determines the wait time for pickup. Castillo estimates the model using 2017 data from Houston when Uber was the only ride-hailing platform. He finds that surge pricing increases overall surplus by a modest amount relative to uniform pricing. However, these gains reflect a combination of a sizable gain for riders and small losses for both drivers and Uber itself. These changes partly reflect the estimated asymmetry in the value of time between riders and drivers, especially at times of high demand. They are also partly due to the assumption that Uber maximizes rider surplus, as Castillo comprehensively discusses.

Rosaia (2021) builds on Castillo (2020), but considers data from New York City where, in contrast to the previous study, multiple platforms (e.g., Uber and Lyft) com-

pete. His model tackles the challenge of incorporating dynamic competition between the two platforms. Two important externalities these competitors face is their effects on traffic volumes, as well as the missed economies of density from platform competition (as in Fréchette et al., 2019). Rosaia uses the estimated model to simulate the effects of a merger. He finds that a merger would lead to improvements in efficiency that are captured by the monopolist, leaving both drivers and passengers worse off. A merger would also substantially reduce the number of vehicles on the street, partly because of a reduction in matching frictions due to an improvement in platform density.

References

- Acemoglu, D., 2001. Good jobs versus bad jobs. *Journal of Labor Economics* 19 (1), 1–21.
- Adda, J., Cooper, R., 2000. Balladurette and Juppette: a discrete analysis of scrapping subsidies. *Journal of Political Economy* 108 (4), 778–806.
- Agarwal, N., Budish, E., 2021. Market design. In: Ho, K., Hortaçsu, A., Lizzeri, A. (Eds.), *Handbook of Industrial Organization*, Vol. 5.
- Akerlof, G.A., 1970. The market for “lemons”: quality uncertainty and the market mechanism. *The Quarterly Journal of Economics* 84 (3), 488–500.
- Allen, J., Clark, R., Houde, J.-F., 2014. The effect of mergers in search markets: evidence from the Canadian mortgage industry. *The American Economic Review* 104 (10), 3365–3396.
- Allen, J., Clark, R., Houde, J.-F., 2019. Search frictions and market power in negotiated price markets. *Journal of Political Economy* 127 (4), 1550–1598.
- Anderson, S.P., Ginsburgh, V.A., 1994. Price discrimination via second-hand markets. *European Economic Review* 38 (1), 23–44.
- Anderson, S.P., Renault, R., 1999. Pricing, product diversity, and search costs: a Bertrand-Chamberlin-Diamond model. *The Rand Journal of Economics* 30 (4), 719–735.
- Anderson, S.P., Renault, R., 2018. Firm pricing with consumer search. In: Corchón, L.C., Marini, M.A. (Eds.), *Handbook of Game Theory and Industrial Organization*, Volume II. Edward Elgar Publishing.
- Armstrong, M., 2017. Ordered consumer search. *Journal of the European Economic Association* 15 (5), 989–1024.
- Armstrong, M., Vickers, J., 2015. Which demand systems can be generated by discrete choice? *Journal of Economic Theory* 158, 293–307.
- Armstrong, M., Vickers, J., 2021. Patterns of competitive interaction. *Econometrica*. Forthcoming.
- Armstrong, M., Vickers, J., Zhou, J., 2009. Prominence and consumer search. *The Rand Journal of Economics* 40 (2), 209–233.
- Arnott, R., 1996. Taxi travel should be subsidized. *Journal of Urban Economics* 40 (3), 316–333.
- Ausubel, L.M., 1991. The failure of competition in the credit card market. *The American Economic Review* 81 (1), 50–81.
- Bar-Isaac, H., Caruana, G., Cufiàt, V., 2012. Search, design, and market structure. *The American Economic Review* 102 (2), 1140–1160.
- Bar-Isaac, H., Gavazza, A., 2015. Brokers’ contractual arrangements in the Manhattan residential rental market. *Journal of Urban Economics* 86, 73–82.

- Barahona, N., Gallego, F.A., Montero, J.-P., 2020. Vintage-specific driving restrictions. *The Review of Economic Studies* 87 (4), 1646–1682.
- Baye, M.R., Kovenock, D., De Vries, C.G., 1992. It takes two to tango: equilibria in a model of sales. *Games and Economic Behavior* 4 (4), 493–510.
- Baye, M.R., Morgan, J., 2001. Information gatekeepers on the Internet and the competitiveness of homogeneous product markets. *The American Economic Review* 91 (3), 454–474.
- Baye, M.R., Morgan, J., Scholten, P., 2004. Price dispersion in the small and in the large: evidence from an Internet price comparison site. *Journal of Industrial Economics* 52 (4), 463–496.
- Baye, M.R., Morgan, J., Scholten, P., 2006. Information, search, and price dispersion. In: Hendershott, T. (Ed.), *Economics and Information Systems*, Vol. 1. Citeseer, pp. 323–375.
- Bergstresser, D., Chalmers, J.M., Tufano, P., 2008. Assessing the costs and benefits of brokers in the mutual fund industry. *The Review of Financial Studies* 22 (10), 4129–4156.
- Berry, S., Levinsohn, J., Pakes, A., 1995. Automobile prices in market equilibrium. *Econometrica* 63 (4), 841–890.
- Biglaiser, G., 1993. Middlemen as experts. *The Rand Journal of Economics* 24 (2), 212–223.
- Biglaiser, G., Friedman, J., 1994. Middlemen as guarantors of quality. *International Journal of Industrial Organization* 12 (4), 509–531.
- Biglaiser, G., Li, F., Murry, C., Zhou, Y., 2020. Middlemen as information intermediaries: evidence from used car markets. *The Rand Journal of Economics* 51 (3), 905–933.
- Bond, E.W., 1982. A direct test of the “lemons” model: the market for used pickup trucks. *The American Economic Review* 72 (4), 836–840.
- Bond, E.W., 1983. Trade in used equipment with heterogeneous firms. *Journal of Political Economy* 91 (4), 688–705.
- Bond, E.W., 1984. Test of the lemons model: reply. *The American Economic Review* 74 (4), 801–804.
- Brancaccio, G., Kalouptsi, M., Papageorgiou, T., 2020a. Geography, search frictions and endogenous trade costs. *Econometrica* 88 (2), 657–691.
- Brancaccio, G., Kalouptsi, M., Papageorgiou, T., Rosaia, N., 2020b. Search Frictions and Efficiency in Decentralized Transport Markets. Working Paper 27300. National Bureau of Economic Research.
- Brown, J.R., Goolsbee, A., 2002. Does the Internet make markets more competitive? Evidence from the life insurance industry. *Journal of Political Economy* 110 (3), 481–507.
- Brown, Z.Y., 2019. An Empirical Model of Price Transparency and Markups in Health Care. Mimeo. University of Michigan.
- Buchholz, N., 2021. Spatial equilibrium, search frictions and dynamic efficiency in the taxi industry. *The Review of Economic Studies*. Forthcoming.
- Burdett, K., Judd, K.L., 1983. Equilibrium price dispersion. *Econometrica* 51 (4), 955–969.
- Burdett, K., Shi, S., Wright, R., 2001. Pricing and matching with frictions. *Journal of Political Economy* 109 (5), 1060–1085.
- Butters, G.R., 1977. Equilibrium distributions of sales and advertising prices. *The Review of Economic Studies* 44 (3), 465–491.
- Calem, P.S., Mester, L.J., 1995. Consumer behavior and the stickiness of credit-card interest rates. *The American Economic Review* 85 (5), 1327–1336.
- Carlin, B.I., 2009. Strategic price complexity in retail financial markets. *Journal of Financial Economics* 91 (3), 278–287.
- Castillo, J.C., 2020. Who Benefits from Surge Pricing? Mimeo. University of Pennsylvania.
- Castillo, J.C., Knoepfle, D.T., Weyl, E.G., 2018. Surge Pricing Solves the Wild Goose Chase. Mimeo. University of Pennsylvania.

- Chade, H., Smith, L., 2006. Simultaneous search. *Econometrica* 74 (5), 1293–1307.
- Chandra, A., Tappata, M., 2011. Consumer search and dynamic price dispersion: an application to gasoline markets. *The Rand Journal of Economics* 42 (4), 681–704.
- Chen, J., Esteban, S., Shum, M., 2013. When do secondary markets harm firms? *The American Economic Review* 103 (7), 2911–2934.
- Choi, M., Dai, A.Y., Kim, K., 2018. Consumer search and price competition. *Econometrica* 86 (4), 1257–1281.
- Coase, R., 1937. The nature of the firm. *Economica* 4 (16), 386–405.
- Coase, R., 1960. The problem of social cost. *The Journal of Law & Economics* 3, 1–44.
- Coase, R., 1972. Durability and monopoly. *The Journal of Law & Economics* 15 (1), 143–149.
- Coey, D., Larsen, B.J., Platt, B.C., 2020. Discounts and deadlines in consumer search. *The American Economic Review* 110 (12), 3748–3785.
- Constantinides, G.M., 1986. Capital market equilibrium with transaction costs. *Journal of Political Economy* 94 (4), 842–862.
- Daley, B., Green, B., 2012. Waiting for news in the market for lemons. *Econometrica* 80 (4), 1433–1504.
- De Cornière, A., 2016. Search advertising. *American Economic Journal: Microeconomics* 8 (3), 156–188.
- De los Santos, B., Hortaçsu, A., Wildenbeest, M.R., 2012. Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review* 102 (6), 2955–2980.
- Diamond, P.A., 1971. A model of price adjustment. *Journal of Economic Theory* 3 (2), 156–168.
- Dinerstein, M., Einav, L., Levin, J., Sundaresan, N., 2018. Consumer price search and platform design in Internet commerce. *The American Economic Review* 108 (7), 1820–1859.
- Donna, J., Pereira, P., Pires, T., Trindade, A., 2021. Measuring the welfare of intermediaries. *Management Science*. Forthcoming.
- Dranove, D., Jin, G.Z., 2010. Quality disclosure and certification: theory and practice. *Journal of Economic Literature* 48 (4), 935–963.
- Duffie, D., Gârleanu, N., Pedersen, L.H., 2005. Over-the-counter markets. *Econometrica* 73 (6), 1815–1847.
- Eckstein, Z., Van den Berg, G.J., 2007. Empirical labor search: a survey. *Journal of Econometrics* 136 (2), 531–564.
- Egan, M., 2019. Brokers versus retail investors: conflicting interests and dominated products. *The Journal of Finance* 74 (3), 1217–1260.
- Einav, L., Farronato, C., Levin, J., Sundaresan, N., 2018. Auctions versus posted prices in online markets. *Journal of Political Economy* 126 (1), 178–215.
- Einav, L., Finkelstein, A., Mahoney, N., 2021. The IO of selection markets. In: Ho, K., Hortaçsu, A., Lizzeri, A. (Eds.), *Handbook of Industrial Organization*, Vol. 5.
- Eisfeldt, A.L., Rampini, A.A., 2006. Capital reallocation and liquidity. *Journal of Monetary Economics* 53 (3), 369–399.
- Ellison, G., Ellison, S.F., 2009. Search, obfuscation, and price elasticities on the Internet. *Econometrica* 77 (2), 427–452.
- Ellison, G., Ellison, S.F., 2018. Match Quality, Search, and the Internet Market for Used Books. Working Paper 24197. National Bureau of Economic Research.
- Ellison, G., Wolitzky, A., 2012. A search cost model of obfuscation. *The Rand Journal of Economics* 43 (3), 417–441.
- Esponda, I., 2008. Behavioral equilibrium in economies with adverse selection. *The American Economic Review* 98 (4), 1269–1291.

- Esteban, S., Shum, M., 2007. Durable-goods oligopoly with secondary markets: the case of automobiles. *The Rand Journal of Economics* 38 (2), 332–354.
- Farboodi, M., Jarosch, G., Menzio, G., Wiriadinata, U., 2019. Intermediation as Rent Extraction. Mimeo. MIT.
- Farrell, J., Klemperer, P., 2007. Coordination and lock-in: competition with switching costs and network effects. In: Armstrong, M., Porter, R. (Eds.), *Handbook of Industrial Organization*, Vol. 3, pp. 1967–2072.
- Fisher Ellison, S., 2016. Price search and obfuscation: an overview of the theory and empirics. In: Basker, E. (Ed.), *Handbook on the Economics of Retailing and Distribution*, pp. 287–305.
- Fréchet, G.R., Lizzeri, A., Salz, T., 2019. Frictions in a competitive, regulated market: evidence from taxis. *The American Economic Review* 109 (8), 2954–2992.
- Fuchs, W., Skrzypacz, A., 2015. Government interventions in a dynamic market with adverse selection. *Journal of Economic Theory* 158, 371–406.
- Galenianos, M., Gavazza, A., 2017. A structural model of the retail market for illicit drugs. *The American Economic Review* 107 (3), 858–896.
- Galenianos, M., Gavazza, A., 2020. Regulatory Interventions in Consumer Financial Markets: the Case of Credit Cards. Working Paper DP13807. Centre for Economic Policy Research.
- Galenianos, M., Pacula, R.L., Persico, N., 2012. A search-theoretic model of the retail market for illicit drugs. *The Review of Economic Studies* 79 (3), 1239–1269.
- Gavazza, A., 2011a. Leasing and secondary markets: theory and evidence from commercial aircraft. *Journal of Political Economy* 119 (2), 325–377.
- Gavazza, A., 2011b. The role of trading frictions in real asset markets. *The American Economic Review* 101 (4), 1106–1143.
- Gavazza, A., 2016. An empirical equilibrium model of a decentralized asset market. *Econometrica* 84 (5), 1755–1798.
- Gavazza, A., Lanteri, A., 2021. Credit shocks and equilibrium dynamics in consumer durable goods markets. *The Review of Economic Studies* 88 (6), 2935–2969.
- Gavazza, A., Lizzeri, A., Roketskiy, N., 2014. A quantitative analysis of the used-car market. *The American Economic Review* 104 (11), 3668–3700.
- Genesove, D., 1993. Adverse selection in the wholesale used car market. *Journal of Political Economy* 101 (4), 644–665.
- Genesove, D., 1995. Search at wholesale auto auctions. *The Quarterly Journal of Economics* 110 (1), 23–49.
- Gilligan, T., 2004. Lemons and leases in the used business aircraft market. *Journal of Political Economy* 112 (5), 1157–1180.
- Gillingham, K., Iskhakov, F., Munk-Nielsen, A., Rust, J., Schjerning, B., 2019. Equilibrium Trade in Automobile Markets. Working Paper 25840. National Bureau of Economic Research.
- Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14 (1), 71–100.
- Grennan, M., Swanson, A., 2020. Transparency and negotiated prices: the value of information in hospital-supplier bargaining. *Journal of Political Economy* 128 (4), 1234–1268.
- Guerrieri, V., Shimer, R., 2014. Dynamic adverse selection: a theory of illiquidity, fire sales, and flight to quality. *The American Economic Review* 104 (7), 1875–1908.
- Gul, F., Sonnenschein, H., Wilson, R., 1986. Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory* 39 (1), 155–190.
- Haan, M.A., Moraga-González, J.L., Petrikaité, V., 2018. A model of directed consumer search. *International Journal of Industrial Organization* 61, 223–255.

- Hagiu, A., Jullien, B., 2011. Why do intermediaries divert search? *The Rand Journal of Economics* 42 (2), 337–362.
- Hagiu, A., Wright, J., 2015a. Marketplace or reseller? *Management Science* 61 (1), 184–203.
- Hagiu, A., Wright, J., 2015b. Multi-sided platforms. *International Journal of Industrial Organization* 43, 162–174.
- Heckman, J.J., 1981. Heterogeneity and state dependence. In: Rosen, S. (Ed.), *Studies in Labor Markets*. University of Chicago Press, pp. 91–140.
- Heckman, J.J., Singer, B., 1984. A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica* 52 (2), 271–320.
- Hendel, I., Lizzeri, A., 1999a. Adverse selection in durable goods markets. *The American Economic Review* 89 (5), 1097–1115.
- Hendel, I., Lizzeri, A., 1999b. Interfering with secondary markets. *The Rand Journal of Economics* 30 (1), 1–21.
- Hendel, I., Lizzeri, A., 2002. The role of leasing under adverse selection. *Journal of Political Economy* 110 (1), 113–143.
- Hendel, I., Lizzeri, A., Siniscalchi, M., 2005. Efficient sorting in a dynamic adverse-selection model. *The Review of Economic Studies* 72 (2), 467–497.
- Hong, H., Shum, M., 2006. Using price distributions to estimate search costs. *The Rand Journal of Economics* 37 (2), 257–275.
- Honka, E., 2014. Quantifying search and switching costs in the US auto insurance industry. *The Rand Journal of Economics* 45 (4), 847–884.
- Honka, E., Hortaçsu, A., Wildenbeest, M., 2019. Empirical search and consideration sets. In: Dubé, J.-P., Rossi, P. (Eds.), *Handbook of the Economics of Marketing*. Elsevier, pp. 193–257.
- Hörner, J., Vieille, N., 2009. Public vs. private offers in the market for lemons. *Econometrica* 77 (1), 29–69.
- Hortaçsu, A., Syverson, C., 2004. Search costs, product differentiation, and welfare effects of entry: a case study of S&P 500 index funds. *The Quarterly Journal of Economics* 119 (4), 403–456.
- Hosios, A.J., 1990. On the efficiency of matching and related models of search and unemployment. *The Review of Economic Studies* 57 (2), 279–298.
- Inderst, R., Ottaviani, M., 2012. Financial advice. *Journal of Economic Literature* 50 (2), 494–512.
- Ireland, N.J., 1993. The provision of information in a Bertrand oligopoly. *Journal of Industrial Economics* 41 (1), 61–76.
- Janssen, M., Moraga-González, J.L., 2004. Strategic pricing, consumer search and the number of firms. *The Review of Economic Studies* 71 (4), 1089–1118.
- Janssen, M., Roy, S., 2002. Dynamic trading in a durable good market with asymmetric information. *International Economic Review* 43 (1), 257–282.
- Janssen, M., Roy, S., 2004. On durable goods markets with entry and adverse selection. *Canadian Journal of Economics* 37 (3), 552–589.
- Johnson, J., Waldman, M., 2003. Leasing, lemons, and buybacks. *The Rand Journal of Economics* 34 (2), 247–265.
- Jolivet, G., Turon, H., 2019. Consumer search costs and preferences on the Internet. *The Review of Economic Studies* 86 (3), 1258–1300.
- Jullien, B., Pavan, A., Rysman, M., 2021. Two-sided markets, pricing, and network effects. In: Ho, K., Hortaçsu, A., Lizzeri, A. (Eds.), *Handbook of Industrial Organization*, Vol. 4. Chapter 7 (this volume).

- Kaplan, G., Menzio, G., 2015. The morphology of price dispersion. *International Economic Review* 56 (4), 1165–1206.
- Kim, J.B., Albuquerque, P., Bronnenberg, B.J., 2010. Online demand under limited consumer search. *Marketing Science* 29 (6), 1001–1023.
- Kim, J.B., Albuquerque, P., Bronnenberg, B.J., 2017. The probit choice model under sequential search with an application to online retailing. *Management Science* 63 (11), 3911–3929.
- Kohn, M.G., Shavell, S., 1974. The theory of search. *Journal of Economic Theory* 9 (2), 93–123.
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica*, 1315–1335.
- Lach, S., 2002. Existence and persistence of price dispersion: an empirical analysis. *Review of Economics and Statistics* 84 (3), 433–444.
- Lacko, J.M., 1986. Product quality and information in the used car market. Technical report. Federal Trade Commission, Bureau of Economics Washington.
- Lagos, R., 2000. An alternative approach to search frictions. *Journal of Political Economy* 108 (5), 851–873.
- Lagos, R., 2003. An analysis of the market for taxicab rides in New York City. *International Economic Review* 44 (2), 423–434.
- Levitt, S.D., Syverson, C., 2008. Market distortions when agents are better informed: the value of information in real estate transactions. *Review of Economics and Statistics* 90 (4), 599–611.
- Lewis, G., 2011. Asymmetric information, adverse selection and online disclosure: the case of eBay motors. *The American Economic Review* 101 (4), 1535–1546.
- Lewis, M.S., Marvel, H.P., 2011. When do consumers search? *Journal of Industrial Economics* 59 (3), 457–483.
- Lizzeri, A., 1999. Information revelation and certification intermediaries. *The Rand Journal of Economics* 30 (2), 214–231.
- McAfee, R.P., 1994. Endogenous availability, cartels, and merger in an equilibrium price dispersion. *Journal of Economic Theory* 62 (1), 24–47.
- Medema, S., 2020. The Coase theorem at sixty. *Journal of Economic Literature* 58 (4), 1045–1128.
- Moraga-González, J.L., Sándor, Z., Wildenbeest, M., 2013. Semi-nonparametric estimation of consumer search costs. *Journal of Applied Econometrics* 28 (7), 1205–1223.
- Moraga-González, J.L., Sándor, Z., Wildenbeest, M., 2018. Consumer Search and Prices in the Automobile Market. Mimeo. Tinbergen Institute.
- Moraga-González, J.L., Sándor, Z., Wildenbeest, M., 2021. Simultaneous search for differentiated products: the impact of search costs and firm prominence. *The Economic Journal* 131 (635), 1308–1330.
- Morgan, P., Manning, R., 1985. Optimal search. *Econometrica* 53 (4), 923–944.
- Mortensen, D., Pissarides, C., 1994. Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies* 61 (3), 397–415.
- Nash, J.F., 1950. The bargaining problem. *Econometrica* 18 (2), 155–162.
- Perloff, J., Salop, S., 1985. Equilibrium with product differentiation. *The Review of Economic Studies* 52 (1), 107–120.
- Petrongolo, B., Pissarides, C., 2001. Looking into the black box: a survey of the matching function. *Journal of Economic Literature* 39 (2), 390–431.
- Philippon, T., Skreta, V., 2012. Optimal interventions in markets with adverse selection. *The American Economic Review* 102 (1), 1–28.
- Porter, R., Sattler, P., 1999. Patterns of Trade in the Market for Used Durables: Theory and Evidence. Working Paper 7149. National Bureau of Economic Research.

- Pulvino, T., 1998. Do asset fire sales exist? An empirical investigation of commercial aircraft transactions. *The Journal of Finance* 53 (3), 939–978.
- Robert, J., Stahl, D.O., 1993. Informative price advertising in a sequential search model. *Econometrica* 61 (3), 657–686.
- Robles-Garcia, C., 2019. Competition and Incentives in Mortgage Markets: the Role of Brokers. Mimeo. Stanford University.
- Rogerson, R., Shimer, R., Wright, R., 2005. Search-theoretic models of the labor market: a survey. *Journal of Economic Literature* 43 (4), 959–988.
- Rosaia, N., 2021. Competing Platforms and Transport Equilibrium: Evidence from New York City. Mimeo. Columbia University.
- Rothschild, M., 1973. Models of market organization with imperfect information: a survey. *Journal of Political Economy* 81 (6), 1283–1308.
- Rubinstein, A., Wolinsky, A., 1987. Middlemen. *The Quarterly Journal of Economics* 102 (3), 581–593.
- Rust, J., 1986. When is it optimal to kill off the market for used durable goods? *Econometrica* 54 (1), 65–86.
- Rust, J., Hall, G., 2003. Middlemen versus market makers: a theory of competitive exchange. *Journal of Political Economy* 111 (2), 353–403.
- Salop, S., Stiglitz, J., 1977. Bargains and ripoffs: a model of monopolistically competitive price dispersion. *The Review of Economic Studies* 44 (3), 493–510.
- Salz, T., 2021. Intermediation and competition in search markets: an empirical case study. *Journal of Political Economy*. Forthcoming.
- Sanches, F., Silva Junior, D., Srisuma, S., 2018. Minimum distance estimation of search costs using price distribution. *Journal of Business & Economic Statistics* 36 (4), 658–671.
- Schiraldi, P., 2011. Automobile replacement: a dynamic structural approach. *The Rand Journal of Economics* 42 (2), 266–291.
- Sieper, E., Swan, P.L., 1973. Monopoly and competition in the market for durable goods. *The Review of Economic Studies* 40 (3), 333–351.
- Sirri, E.R., Tufano, P., 1998. Costly search and mutual fund flows. *The Journal of Finance* 53 (5), 1589–1622.
- Sorensen, A.T., 2000. Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy* 108 (4), 833–850.
- Spence, A.M., 1975. Monopoly, quality, and regulation. *Bell Journal of Economics* 6 (2), 417–429.
- Spulber, D., 1996. Market microstructure and intermediation. *The Journal of Economic Perspectives* 10 (3), 135–152.
- Spulber, D., 1999. *Market Microstructure: Intermediaries and the Theory of the Firm*. Cambridge University Press.
- Spulber, D., 2002. Market microstructure and incentives to invest. *Journal of Political Economy* 110 (2), 352–381.
- Stahl, D.O., 1989. Oligopolistic pricing with sequential consumer search. *The American Economic Review* 79 (4), 700–712.
- Stahl, D.O., 1994. Oligopolistic pricing and advertising. *Journal of Economic Theory* 64 (1), 162–177.
- Stango, V., Zinman, J., 2016. Borrowing high versus borrowing higher: price dispersion and shopping behavior in the US credit card market. *The Review of Financial Studies* 29 (4), 979–1006.
- Stigler, G.J., 1961. The economics of information. *Journal of Political Economy* 69 (3), 213–225.

- Stolyarov, D., 2002. Turnover of used durables in a stationary equilibrium: are older goods traded more? *Journal of Political Economy* 110 (6), 1390–1413.
- Swan, P.L., 1970. Durability of consumption goods. *The American Economic Review* 60 (5), 884–894.
- Swan, P.L., 1971. The durability of goods and regulation of monopoly. *The Bell Journal of Economics and Management Science* 2 (1), 347–357.
- Swan, P.L., 1980. Alcoa: the influence of recycling on monopoly power. *Journal of Political Economy* 88 (1), 76–99.
- Thakral, N., Tô, L.T., 2021. Daily labor supply and adaptive reference points. *The American Economic Review* 111 (8), 2417–2443.
- Tirole, J., 2012. Overcoming adverse selection: how public intervention can restore market functioning. *The American Economic Review* 102 (1), 29–59.
- Varian, H., 1980. A model of sales. *The American Economic Review* 70 (4), 651–659.
- Waldman, M., 1996. Planned obsolescence and the R&D decision. *The Rand Journal of Economics* 27 (3), 583–595.
- Waldman, M., 2003. Durable goods theory for real world markets. *The Journal of Economic Perspectives* 17 (1), 131–154.
- Wang, C., Wright, J., 2020. Search platforms: showrooming and price parity clauses. *The Rand Journal of Economics* 51 (1), 32–58.
- Weitzman, M.L., 1979. Optimal search for the best alternative. *Econometrica* 47 (3), 641–654.
- Wheaton, W.C., 1990. Vacancy, search, and prices in a housing market matching model. *Journal of Political Economy* 98 (6), 1270–1292.
- Wildenbeest, M.R., 2011. An empirical model of search with vertically differentiated products. *The Rand Journal of Economics* 42 (4), 729–757.
- Williamson, O., 1985. *The Economic Institutions of Capitalism*. Simon and Schuster.
- Williamson, O., 1989. Transaction cost economics. In: Schmalensee, R., Willig, R. (Eds.), *Handbook of Industrial Organization*, Vol. 1. Elsevier, pp. 135–182.
- Wilson, C.M., 2010. Ordered search and equilibrium obfuscation. *International Journal of Industrial Organization* 28 (5), 496–506.
- Wolinsky, A., 1986. True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics* 101 (3), 493–511.
- Woodward, S.E., Hall, R.E., 2012. Diagnosing consumer confusion and sub-optimal shopping effort: theory and mortgage-market evidence. *The American Economic Review* 102 (7), 3249–3276.
- Wright, R., Kircher, P., Julien, B., Guerrieri, V., 2021. Directed search and competitive search: a guided tour. *Journal of Economic Literature* 59 (1), 90–148.
- Zhou, J., 2014. Multiproduct search and the joint search effect. *The American Economic Review* 104 (9), 2918–2939.