

Incomplete-Information Models for Industrial Organization*

Simon Loertscher[†] Leslie M. Marx[‡]

January 13, 2021

Abstract

Abstract: Recent advances have demonstrated that pertinent issues in Industrial Organization and Antitrust, such as the profit-surplus tradeoff and countervailing power, arise naturally in incomplete-information models. These model also explain, coherently, a host of issues such as why first-degree price discrimination is not possible, why uniform pricing is optimal when revenue is concave, and why opaque pricing and rationing (which may induce resale) is optimal otherwise. This chapter provides an overview of the relevant literature, explains the key mechanics of the framework, discusses applications to horizontal mergers, vertical integration, and buyer and countervailing power, and lays out a roadmap for the work ahead.

Keywords: Triple-IO, bargaining power, buyer power, productive power, merger, vertical integration, mechanism design

JEL Classification: D44, D82, L41

*We thank the audience at the Econometric Society & Bocconi University Virtual World Congress for valuable comments. Financial support from the Samuel and June Hordern Endowment and Australian Research Council Discovery Project Grant DP200103574 is gratefully acknowledged.

[†]Department of Economics, Level 4, FBE Building, 111 Barry Street, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au.

[‡]Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708, USA: Email: marx@duke.edu.

1 Introduction

Industrial Organization covers imperfectly competitive markets and incentives within firms and organizations. It derives particular real-world relevance through its application to regulation, consumer protection, antitrust, and competition policy, as well as related court cases. Naturally, it aims to analyze the effects of alternative arrangements, such horizontal and vertical integration, workers' right to unionize, and firms' ability to price discriminate, on consumer and social surplus, and the inevitable tradeoffs between them.

This applicability and real-world exposure simultaneously constrain the field and give it direction and purpose. As anyone privy to policy debates or courtroom arguments will readily attest, it is not the fanciest and theoretically most advanced arguments ("Shouldn't you be doing dynamic mechanism design here?") that win the day, partly because of the need for evidence-based decision making and the inevitable limitations of data and empirical methods, and partly because of the need for those arguments to be comprehended by intelligent, noneconomist policymakers and judges.¹ At the same time, the fundamental need for conceptual frameworks to contemplate and communicate genuine counterfactuals—changes in the economic environment from the known into the unknown—is painstakingly clear. In an ever changing world, there are simply never going to be enough takeovers of Instagram by Facebook or Time Warner by AT&T for any empiricist to make conclusive inferences, one way or another, without relying on assumptions. In other words, the field inevitably requires a combination of theory and empirics within a framework that generates a tradeoff between profits and social and consumer surplus, allows for social surplus effects of varying bargaining power, and accommodates incentives for investment and innovation.

In this chapter, we provide an introduction to what we call the *Incomplete Information Industrial Organization* framework—or *Triple-IO*. The aims of this chapter are, on the one hand, to demonstrate to readers not familiar with it the framework's tractability, discipline, and flexibility, and, on the other, to illustrate and discuss its internal consistency and external validity. The chapter first discusses the background and backdrop for Triple-IO. Then we review progress to date. The chapter concludes with an overview of outstanding questions and the path ahead.

¹Baye and Wright (2011) state: "We find that decisions involving the evaluation of complex economic evidence are significantly more likely to be appealed, and decisions of judges trained in basic economics are significantly less likely to be appealed than are decisions by their untrained counterparts. Our analysis supports the hypothesis that some antitrust cases are too complicated for generalist judges."

2 Motivation

In this section, we provide background for antitrust economics and an overview of recent developments in Antitrust and Industrial Organization, taking these as motivation for incomplete-information models.

2.1 Background

Monopoly behavior is a longstanding and ongoing focus of interest in Industrial Organization and Antitrust, and thus serves well as a lead-in to discussion of alternative models. Traditional approaches—by which we mean approaches that are not based on incomplete-information models—to monopoly pricing typically assume that a monopoly is, for reasons that are outside the model (such as resale among buyers), restricted to setting a uniform price to consumers. The deeper, although not necessarily widely publicized, reason for imposing this assumption is that it prevents the monopoly from perfectly price discriminating among the consumers, in which case there would be no social surplus rationale for regulating monopolies or preventing monopolies from forming. The goal of a monopoly that captures all of the social surplus that it creates is perfectly aligned with the social planner’s objective since it grabs all surplus. As we will see, incomplete information immediately implies a tradeoff between social surplus and profit because eliciting information about agents’ values has an incentive cost.

While the dependence of the predictions and prescriptions on these assumptions has long been recognized, the modeling device of assuming uniform pricing remains influential in undergraduate teaching, research, and practical applications. As a case in point, successive uniform pricing monopolies give rise to double marginalization (Spengler, 1950) whose elimination continues to be an influential and widely accepted efficiency-based argument in favor of vertical integration (see, e.g., the 2020 U.S. *Vertical Merger Guidelines*).

The basic tenet in the monopoly model that one side of the market sets the price for the other carries over to oligopoly models à la Cournot or Bertrand, which are used to model imperfectly competitive markets. The main effect of competition in these models is that it renders each firm’s residual demand function more price elastic and, therefore, makes it behave more aggressively. Competition does not alter the fact that one side of the market (in consumer markets, typically the supply side) sets the terms of trade. These models are therefore incapable of capturing notions like shifts in bargaining power, buyer power, or countervailing power. This creates a disconnect between the reliance on these notions

in practice,² and the models being used to analyze and quantify associated effects, which remain largely grounded in oligopoly models à la Cournot or Bertrand.³

Antitrust concerns and thinking in their current form emerged in an era of steel and rail. The regulation of *natural monopolies* that arise because of nonconvex cost functions was a central concern, and took on the form, for example, of Ramsey pricing (see, e.g., Ramsey, 1927; Wilson, 1993), which involves a convex combination of perfectly competitive prices and monopoly prices. Such approaches recognize that, because of the cost structure and network organization of the problems at hand, perfectly competitive pricing would not allow the operators to break even (and would provide poor incentives for maintaining and improving infrastructure), and that the monopoly price would have undesirable effects for consumer and social surplus. Even though it may not have been achievable, marginal cost pricing was a clear and sensible benchmark against which one could evaluate firms' behavior. As the discussion below will make clear, this is a luxury that concurrent issues in Antitrust are no longer necessarily afforded because notions of competitive pricing are often vague or vacuous.

2.2 Recent developments in Antitrust

The two salient developments in Antitrust in recent times are, on the one hand, an upsurge of concern with big tech companies and, on the other, a recognition of the need to account for bargaining and negotiation procedures that are not captured by standard oligopoly models.

The last few years have seen an upsurge of interest in and concern with competition issues in the digital age. In 2019 alone, the Australian Competition and Consumer Commission, the European Commission, the Digital Competition Expert Panel in the U.K., and the Stigler Center for Study of the Economy and the State produced four substantive reports devoted to that.⁴ In the United States and Europe, 2020 saw investigations and antitrust cases launched against Google, Facebook, and Apple, and combinations thereof. The ongoing nature of the

²The 2010 U.S. *Horizontal Merger Guidelines* (2010 *Guidelines*), which both guide and reflect antitrust practice in the United States, note in Section 6.2 on “Bargaining and Auctions” that “Buyers commonly negotiate with more than one seller, and may play sellers off against one another,” and in Section 8 on “Powerful buyers” that “Powerful buyers are often able to negotiate favorable terms with their suppliers.” Further, the *Guidelines* allow that Agencies could identify a *price discrimination market* defined by a single powerful buyer.

³For example, the reliance of competition authorities around the globe on merger screens based on the Herfindahl-Hirschman Index (HHI) has foundations in the Cournot oligopoly model (see, e.g., Werden, 2008; Werden and Froeb, 2008). A variety of well-accepted merger review tools and simulation techniques are grounded in the differentiated products Bertrand oligopoly model (see, e.g., Epstein and Rubinfeld, 2004). Much of the analysis surrounding collusive conduct has as its conceptual framework a model of repeated Cournot competition along the lines of Green and Porter (1984) (see, e.g., Ivaldi et al., 2007).

⁴See Australian Competition & Consumer Commission (2019), Cémer et al. (2019), Furman et al. (2019), and Stigler Center (2019).

concerns is also evident in the creation in 2019 by the U.S. Federal Trade Commission of a permanent “Technology Enforcement Division,” which is tasked with monitoring competition and investigating potential anticompetitive conduct in U.S. technology markets. Further, officials at the U.S. Department of Justice have said that: “We plan to *continue* our review of competitive practices by market-leading online platforms, and where necessary address those as well.”⁵

Over the last decade or so, there have been a large number of substantial merger cases in which bargaining—understood as a process to determine the terms of trade that is not adequately captured by off-the-shelf oligopoly models—played an important role. Recent cases include the AT&T–Time Warner merger, where a key concern was the merger’s effect on negotiations between the merged entity and video content distributors, as well as numerous other media mergers.⁶ There is a similarly long list of health care mergers, including, e.g., the merger of Hackensack Meridian Health and Englewood Healthcare, which raise issues related to negotiations between hospitals and insurers,⁷ as well as other mergers where the relevant market is most appropriately viewed as a bilateral oligopoly, such as the proposed merger of Halliburton-Baker Hughes in oilfield services. Antitrust matters related to the licensing of intellectual property, particularly among pharmaceutical and high tech firms,⁸ are reflected in recently issued antitrust guidelines for the licensing of intellectual property and require an evaluation of the relative bargaining power of the licensor and licensee for the calculation of damages.⁹

Consistent with the cases just mentioned, the 2020 *Guidelines* emphasize the importance and relevance of bargaining and auctions, stating: “In many industries, especially those involving intermediate goods and services, buyers and sellers negotiate to determine prices and other terms of trade. In that process, buyers commonly negotiate with more than one seller, and may play sellers off against one another. Some highly structured forms of such

⁵“Introductory Remarks of Deputy Attorney General at Announcement of Civil Antitrust Lawsuit Filed Against Google,” October 20, 2020, <https://www.justice.gov/opa/speech/introductory-remarks-deputy-attorney-general-announcement-civil-antitrust-lawsuit-filed> (emphasis in the original).

⁶For example, Nexstar-Tribune merger raised issues related to negotiations between broadcasters and advertisers. For additional examples, see the list on the U.S. Department of Justice’s “Antitrust Case Filings” website, <https://www.justice.gov/atr/antitrust-case-filings?>.

⁷U.S. Federal Trade Commission, “Merger Review” website, <https://www.ftc.gov/news-events/media-resources/mergers-and-competition/merger-review>.

⁸For example, *VirnetX, Inc. v. Cisco Systems* (2014) related to technology used in the FaceTime and VPN features in iOS devices and Mac computers, and *FTC v. Actavis, Inc.* (2013) related to negotiated payments by a drug patentee to a generic drug manufacturer known as “pay-for-delay.”

⁹U.S. Department of Justice and Federal Trade Commission, “Antitrust Guidelines for the Licensing of Intellectual Property,” 2017, <https://www.justice.gov/atr/IPguidelines/download>. See, e.g., Sidak (2015) on the role of bargaining power in the *Georgia-Pacific* framework for damages calculations and courts’ rejection of Nash bargaining for the quantification of damages.

competition are known as auctions. Negotiations often combine aspects of an auction with aspects of one-on-one negotiation, although pure auctions are sometimes used in government procurement and elsewhere” (2010 *Guidelines*, Section 6.2).

2.3 Recent developments in Industrial Organization

There has been an upsurge of interest in bargaining in the empirical IO literature, driven both by the advent of big data sets and the need to accommodate negotiated outcomes in counterfactual analyses.

For example, issues of market power and the possibility of countervailing market power are prominent in recent debates surrounding Internet giants such as Google. For example, Decarolis and Rovigatti (2020) show that consolidation among online advertising intermediaries has increased their buyer power, countervailing Google’s significant market power in online search, and is consistent with the theoretical result that consolidation that equalizes bargaining power between the buyer and seller sides of the market can increase welfare. Related to online negotiations through eBay’s Best Offer bargaining platform, Backus et al. (2020, p. 2) find that “the observation that bargaining very frequently ends in disagreement is consistent with the presence of incomplete information and bargaining costs,” and Backus et al. (2019) find that agents take steps to credibly communicate the strength of their bargaining positions.

Outside the realm of big Internet firms, Larsen (2020) estimates a model of incomplete-information bargaining in the context of negotiations in the wholesale used-car industry, documenting, empirically, the efficiency effects of private information. Byrne et al. (2019) find evidence of bargaining power effects in negotiations between simulated customers and call centers for providers of retail electricity, where customers demonstrating a greater willingness and ability to search and bargain obtained better prices. Sweeting et al. (2018) examine data on U.S. beer sales and show that signaling effects that arise with incomplete information can cause merger simulations that assume complete information to significantly underpredict post-merger price rises. Related to the licensing of intellectual property, Caves et al. (1983) find evidence that patent owners cannot fully appropriate a licensee’s maximum rent due, in part, to incomplete information, and Kankanhalli and Kwan (2019) provide empirical evidence that royalty rates in intellectual property licensing agreements are affected by bargaining power.

In contrast to the research just mentioned, to date, the more standard approach to the empirical analysis of markets with negotiated prices has been to use the complete-information

models of Nash bargaining and Nash-in-Nash bargaining.¹⁰ The Nash bargaining solution has the feature that parties divide a fixed pie, with no tradeoff between the division of the pie and its overall size. For example, one way that greater bargaining power might express itself is through more aggressive offers, with the result that bargaining is more likely to breakdown. But because bargaining breakdown does not occur in Nash bargaining on the equilibrium path, that model is a poor fit for settings where there is systematic bargaining breakdown on the equilibrium path. For example, Backus et al. (2020) find a breakdown probability of roughly 55 percent in data covering 25 million observations of bilateral negotiations on eBay.¹¹

The Nash-in-Nash model, introduced by Horn and Wolinsky (1988), extends Nash bargaining to accommodate multiple simultaneous Nash bargains.¹² In that model, the bargaining equilibrium is the Nash equilibrium of a game in which players participate in potentially many Nash bargains, but with each player essentially sending separate representatives to each Nash bargain, with no coordination across the representatives of a player. As noted by Crawford and Yurukoglu (2012, p. 659), “This setup does not allow for advantages due to informational asymmetries. ... We view this absence of informational asymmetries as a weakness of the bargaining model. In return, however, we gain tractability in determining how the threat of unilateral disagreement determines input costs in a bilaterally oligopolistic setting.” Thus, although the Nash-in-Nash model constrains the market interaction in a way that raises concerns, it is viewed as offering a benefit in terms of tractability.

The large literature making use of Nash-based models is testament to the relevance of negotiated outcomes, bargaining power, and countervailing power in Industrial Organization. On the effects of hospital mergers on bargaining power and negotiated outcomes, see, e.g., Capps et al. (2003); Sorensen (2003); Capps and Dranove (2004); Dafny (2009); Gaynor and Town (2012); Cutler and Scott Morton (2013); Gowrisankaran et al. (2015); Ho and Lee (2017, 2019); Craig et al. (2020); Prager and Schmitt (forth.).¹³ For example, Ho and Lee (2017) apply the Nash-in-Nash framework to the question of countervailing power by

¹⁰The remarks of Nevo (2014) describe the important role that the theory of Nash bargaining plays in antitrust practice, and Shapiro (2018) uses the Nash bargaining model to analyze and quantify competitive harms associated with the AT&T-Time Warner merger.

¹¹As described by Crawford (2014), there are regular blackouts of broadcast television stations on cable and satellite distribution platforms due to the breakdown of negotiations over the terms for retransmission of the broadcast signal.

¹²Collard-Wexler et al. (2019) and Rey and Vergé (2019) provide theoretical foundations for the Nash-in-Nash model.

¹³On bargaining between hospitals and suppliers of medical devices, see Grennan (2013, 2014); Grennan and Swanson (2020). In work on insurer-hospital bargaining, Lewis and Pflum (2015) find that bargaining power (the share of generated surplus captured by an agent) is a greater determinant of post-merger markups than bargaining positions (the agent’s threat point).

insurers when negotiating with hospitals and find evidence that consolidation among insurers improves their bargaining position vis-à-vis hospitals. Craig et al. (2020) find evidence that following mergers, hospitals are able to negotiate somewhat lower prices for medical supplies, consistent with mergers involving countervailing effects of improved buyer power and managerial disruption.¹⁴ For media mergers, where negotiations play a role in determining what content is distributed where and at what prices and what rates are paid by advertisers, see, e.g., Gal-Or and Dukes (2003, 2006); Crawford et al. (2018a). Related to bundling and vertical integration, see Crawford and Yurukoglu (2012) and Crawford et al. (2018b). Related to intellectual property, for Nash-bargaining based models of patent hold-up and litigation, see, e.g., Katz and Shapiro (1985); Bessen and Meurer (2006); Shapiro (2010); Ganglmair et al. (2012); Lemus and Temnyalov (2017); Choi and Gerlach (2017).

2.4 The case for incomplete-information IO models

As mentioned, in models of complete information bargaining, the only thing to be bargained over is the distribution of the pie; whether or not there is a pie to share and, if so, its size is common knowledge before bargaining even starts. This not only makes it difficult for complete information bargaining models to account for the salient real-world phenomenon of bargaining breakdown, but also deprives bargaining of some of its most important features.¹⁵ While complete information bargaining models have been criticized by some economists,¹⁶ (and, of course, been embraced by many others), some skepticism vis-à-vis complete information bargaining has also been uttered by practitioners, perhaps most vocally by Judge Leon in the AT&T–Time Warner merger case. Upon hearing the government’s Nash bargaining model, he wrote: “I wondered on the record whether its complexity made it seem like a Rube Goldberg contraption. ... But in fairness to Mr. Goldberg, at least his contraptions would normally move a pea from one side of a room to another. By contrast, the evidence at trial showed that [the] model lacks both ‘reliability and factual credibility,’ and thus fails to generate probative predictions of future harm associated with the Government’s increased-leverage theory” (Leon, 2018, p. 32).

¹⁴They conclude that price savings seem “modest relative to the cross-sectional price variation across hospitals and claims of potential savings via increased ‘buyer power’” (Craig et al., 2020, p. 41). Related to buyer power more generally, see, e.g., Snyder (1996); Caprice and Rey (2015).

¹⁵Or as Fudenberg and Tirole (1991, p. 399) put it: “bargaining derives much of its interest from incomplete information.”

¹⁶For example, Holmström and Myerson (1983, p. 1809) note: “Some economists, following Coase have ... argued that we should expect to observe efficient allocations in any economy where there is complete information and bargaining costs are small. However, this positive aspect of efficiency does not extend to economies with incomplete information.” For a similar view, see Samuelson (1985).

Incomplete-information bargaining

The modeling backbone for incomplete-information bargaining that we work with is the bilateral trade model of Myerson and Satterthwaite (1983), which we revisit in the next section, and its generalizations such as Williams (1987), who extended it to allow the traders to have different bargaining weights, and Gresik and Satterthwaite (1989), who generalized it to multiple traders on each side of the market. As observed by Ausubel et al. (2002, p. 1934), in light of the results of Myerson and Satterthwaite (1983) insisting on efficiency in incomplete-information bargaining is “fruitless.” The methodology we employ is mechanism design, and empowered and encouraged by the revelation principle (see below), our focus is on direct mechanisms. It is our view, and our aim to convince the readers, that the time is ripe to do so.¹⁷

Earlier literature on incomplete information bargaining has paid considerable attention to the extensive-form game in which bargaining takes place. Part of the reason for this was theoretical. Ausubel et al. (2002, p. 1908) note that, despite its virtues, the mechanism-design approach has the weaknesses of relying on common knowledge of agents’ beliefs and utility functions and of assuming too much commitment. They write: “In practice bargainers use simple trading rules—such as a sequence of offers and counteroffers—that do not depend on beliefs or utility functions. And bargainers may be unable to walk away from known gains from trade.” As will become clear from the analysis below, it is precisely by varying an agent’s ability to walk away from known gains from trade that the mechanism-design approach is able to capture notions like bargaining and buyer power, which have social surplus effects. So we view the ability of the mechanism-design approach to endow agents with such commitment power as a merit rather than as a shortcoming. Of course, we maintain the assumption of common knowledge of type distributions and payoff functions throughout. In defense of this assumption, we merely note that every model has to start somewhere and stipulating that agents share a common prior has proved a productive assumption in economics. Further, the assumption is not specific to incomplete-information models, although it is almost exclusively raised in the context of incomplete-information models, perhaps because it is, somewhat ironically, more glaring there. With complete information, agents do not only have common knowledge of distributions and utility functions, but they also happen to know that every agent’s distribution is degenerate with the lower bound of support equal to its upper bound. Put differently, progress toward models that dispense with

¹⁷Fifteen to twenty years ago, there was something of a breakthrough of auction theory into “mainstream” economics, with Paul Klemperer explaining “why every economist should learn some auction theory” (Klemperer, 2002); see also Krishna (2002), Klemperer (2004), and Milgrom (2004). With that in mind, the view and belief expressed here is that now is the time to move to another level of abstraction and to embrace the mechanism-design approach.

the common-knowledge assumption would add value, but lack thereof is nothing that sets incomplete-information models at any disadvantage relative to complete-information models because, whatever criticism applies to the former, a fortiori applies to the latter.

Another part of the motivation for extensive-form representations of bargaining seems to have been empirical. For example, the structural estimation approach in auctions makes use of knowledge of the auction format—for example, first-price auction in Guerre et al. (2000)—and assumptions about type distributions and equilibrium—in this case, identical distributions and Bayes Nash equilibrium—to estimate the distributions of types based on bidding data. From a practical modeling perspective, the problem of relying on extensive forms is that they permit tractability only in special cases. For example, the first-price auction is tractable with identical distributions but rarely otherwise.¹⁸ Likewise, the k -double-auction of Chatterjee and Samuelson (1983) is tractable for two bidders and uniform distributions, but not otherwise. This severely impedes the use of these models for, say, merger review because the distributions cannot sensibly be stipulated to be the same across all bidders before and after two agents merge. These problems are compounded in an empirical context, where it is arguably more often than not the case that the outside observers cannot be sure that they observe the precise bargaining game that is being played.¹⁹

Fortunately, reliance on knowledge of the extensive form is not required empirically either. As Larsen and Zhang (2018) show, incentive compatibility and the revelation principle can be used, together with the hypothesis that agents’ actions constitute a Bayes Nash equilibrium, to infer type distributions from limited observations of agents’ actions, payments, and allocations. We illustrate their approach in Section 4.5 below.

From extensive forms to direct mechanisms

Real-world bargaining can take a variety of forms. Its description might be relatively straightforward, such as when a seller posts a take-it-or-leave-it price, which a buyer can then either accept or reject. Or, the price-determination process might involve significant back and forth, with the buyer and seller engaging in a set of unstructured offers and counteroffers

¹⁸Of course, the first-price auction is detail-free in the sense of Wilson (1987), which makes it appealing for practical use. But that does not make it a suitable theoretical device to model price formation any more than the great theoretical properties of the second-price (or Vickrey) auction make it appealing for practical use.

¹⁹As Fudenberg and Tirole (1991, pp. 398–399) note: “Even with complete information, any split of the pie can be obtained by changing the extensive form for bargaining. This is worrisome because we, as outside observers of a bargaining process, usually have little information about which extensive form is being played, and furthermore the extensive form is likely to vary from one situation to the next. Of course, in any application of game theory the conclusions can vary with the extensive form chosen, but the issue seems more serious here than in other contexts, in which we may be able to limit attention to a smaller set of extensive forms.”

and other communication over time, eventually settling on terms of trade. The process might be partially structured, with an initial auction phase in which suppliers submit bids to an online platform, but then the buyer might continue to negotiate with one or more of the auction participants, perhaps asking for additional concessions or asking for quantity discounts, after the auction phase is over. The specification of the details of a bargaining process is referred to as its *extensive form*. Extensive forms have in common that they take as inputs the characteristics and private information of the parties, particularly as related to the good being traded, and, through the application of agents' optimal strategies given their private information, produce as outputs a specification of what trades occur and what payments are made. These inputs and outputs are also the information required to evaluate the contribution of the bargaining process to consumer, producer, and social surplus. This raises the question: to understand how a particular set of inputs (agent characteristics and private information) affect outcomes such as consumer, producer, and social surplus, do we need to allow for all the possible extensive-form bargaining games that could be played and how they might be played. Happily, the answer to this is no. The *revelation principle* implies that any extensive-form game that agents with independent private information might play can be reformulated as a *direct mechanism* that maps agents reports of their private information to allocations and transfers. Somewhat more precisely, given an extensive-form game and optimal strategies for all the players as a function of their independent private information, there exists a direct mechanism such that it is optimal for each agent to truthfully report its private information to the mechanism and such that the direct mechanism replicates the outcome associated with the agents' optimal strategies in the extensive-form game (see, e.g., Myerson, 1981, Lemma 1; Börgers, 2015, Proposition 2.1). This powerful result simplifies our task greatly because it implies that without loss of generality, we can focus on direct mechanisms.

2.5 Of maps and models

Models, like maps, are abstractions. As such, they deliberately violate certain facts and data in order to shed light on others. Just as the usefulness of a map depends on its purpose, so does the value of a specific model. In Industrial Organization and Antitrust, the purpose of a model is to serve as a conceptual framework for problems that involve tradeoffs between social surplus and profit and to permit the conceptualization of policy counterfactuals. Using a flat map does not make a traveller a flat-worlder, nor does adhering to a model with rational, risk neutral agents with a common prior mean that the economists who do so believe that these assumptions are factual descriptions. No user of a flat map would ever be accused of

holding the erroneous and demonstrably wrong belief that buildings are two-dimensional. Those who use these kinds of maps and models find them useful because they serve the purpose of guiding them from here to there.

It is also worth remembering that new scientific theories often precede the arrival of data substantiating them. This was true for Copernicus and Galilei, who stipulated—without evidence (for centuries to come)—that the sun rather than the earth is the center of our solar system, for Darwin’s theory of evolution, and Einstein’s theory of relativity (and, if you like, the independent private values model). It was not notions of plausibility and realism of the assumptions that won the day for these new approaches.²⁰ Indeed, it is difficult to argue that the idea that we and everything on Earth move at all times at speeds faster than bullets fly is more plausible or realistic than stipulating that Earth stands still. This seems worth bearing in mind when discussing and evaluating the pros and cons of the incomplete information approach to Industrial Organization.

3 Methodology

In this section, we first describe the setup, and then explore implications of incentive compatibility and individual rationality. Then we examine optimal and efficient mechanisms. We conclude with a discussion of the uniqueness of the properties of the independent private values setting.

3.1 Setup

We consider a setup with multiple buyers and sellers, although in some cases we specialize to the case of bilateral trade, where there is just one buyer and just one seller, or even to the case of a single agent interacting with a mechanism designer. In the baseline setup, there are n buyers and m sellers.

We assume that agents have independent private types. Each buyer i has private value v_i for one unit of a good, where v_i is drawn independently from the distribution F with support $[\underline{v}, \bar{v}]$. Each seller j can produce one unit of the good at private cost c_j drawn independently from distribution G with support $[\underline{c}, \bar{c}]$. We assume that $\underline{c} < \bar{v}$ so that gains from trade are possible.

We assume that agents are risk neutral, implying that the expected payoff of buyer i with value v_i that trades expected quantity q and makes an expected payment to the mechanism

²⁰See, for example, Kuhn (1970).

of m is $qv_i - m$, and the expected payoff of seller i with cost c_i that trades expected quantity q and receives expected payment from the mechanism of m is $m - qc_i$.

3.2 Incentive compatibility, individual rationality, and their implications

To explore the powerful implications of incentive compatibility, focus temporarily on a single agent, which could be either a buyer or a seller. To embrace both possibilities, we denote the agent's type by $x \in [\underline{x}, \bar{x}]$, which would be $v \in [\underline{v}, \bar{v}]$ for a buyer and $c \in [\underline{c}, \bar{c}]$ for a seller. Take as given a direct mechanism $\langle q, m \rangle$, where $q : [\underline{x}, \bar{x}] \rightarrow [0, 1]$ for a buyer and $q : [\underline{x}, \bar{x}] \rightarrow [-1, 0]$ for a seller, and where $m : [\underline{x}, \bar{x}] \rightarrow \mathbb{R}$. The mechanism maps the agent's reported type $z \in [\underline{x}, \bar{x}]$ onto an expected quantity traded $q(z)$ and expected payment to the mechanism $m(z)$.²¹ As mentioned above, by the *revelation principle* (see, e.g., Myerson, 1981, Lemma 1; B6rgers, 2015, Proposition 2.1), a focus on direct mechanisms is without loss of generality.

The payoff of an agent with type x that reports z is

$$U(x, z) \equiv q(z)x - m(z),$$

and the payoff of an agent with type x under truthful reporting is

$$U(x) \equiv q(x)x - m(x). \tag{1}$$

The mechanism is *incentive compatible* if an agent's payoff from truthful reporting is greater than or equal to its payoff from reporting any other type.²² That is, for all $x, z \in [\underline{x}, \bar{x}]$,

$$U(x) \geq q(z)x - m(z) = U(z) + q(z)(x - z). \tag{2}$$

Individual rationality is satisfied if for all $x \in [\underline{x}, \bar{x}]$, $U(x) \geq 0$.

For example, if $q(z) > 0$ and $m(z) > 0$, then the agent is a buyer that receives $q(z)$ units of a good for which it has a per-unit value of x in exchange for a payment of $m(z)$. If instead

²¹With multiple agents, define $q(z)$ and $m(z)$ to be the agent's expected allocation and transfer when its report is z , with the expectation taken over the other agents' types.

²²With multiple agents, one can distinguish between incentive compatibility as we have defined it, i.e., evaluated at the interim stage when an agent knows its own type, but not the types of the other agents, and ex post incentive compatibility, which requires that an agent's payoff is maximized by truthful reporting for all possible realizations of the other agents' types. In our independent private values setting, the outcome of any interim (Bayesian) incentive compatible and interim individually rational mechanism can be implemented with a mechanism that satisfies ex post incentive compatibility and ex post individual rationality (Gershkov et al., 2013).

$q(z) < 0$ and $m(z) < 0$, then the agent is a seller that produces $|q(z)|$ units of a good at a per-unit cost of x and receives a payment of $m(z)$ from the mechanism.

Following standard arguments (see, e.g., Börger, 2015, Chapter 2), we can write the definition of incentive compatibility in (2) as: for all $x \in [\underline{x}, \bar{x}]$,

$$U(x) = \max_{z \in [\underline{x}, \bar{x}]} q(z)x - m(z).$$

This implies that U is a maximum of a family of affine functions, which implies that U is convex and so absolutely continuous and differentiable almost everywhere in the interior of its domain.²³

Further, the definition of incentive compatibility in (2) is equivalent to: for all $x, z \in [\underline{x}, \bar{x}]$,

$$U(x) - U(z) \geq q(z)(x - z).$$

Dividing both sides by $x - z$, one obtains, for $x > z$, $\frac{U(x) - U(z)}{x - z} \geq q(z)$ and, for $x < z$, $\frac{U(x) - U(z)}{x - z} \leq q(z)$. Letting $z \rightarrow x$ thus yields $U'(x) \geq q(x) \geq U'(x)$ at every point at which U is differentiable, forcing the conclusion that at every such point, $U'(x) = q(x)$. In other words, for all x , $q(x)$ is the slope of a line that supports the function U at the point x . This is illustrated in Figure 1.




Figure 1: Illustration of the convexity of U (Krishna, 2010, Figure 5.2)

Because U is convex, it follows that q is nondecreasing. Further, because every absolutely

²³A function $h : [\underline{x}, \bar{x}] \rightarrow \mathbb{R}$ is absolutely continuous if for all $\varepsilon > 0$ there exists $\delta > 0$ such that whenever a finite sequence of pairwise disjoint sub-intervals (x_k, x'_k) of $[\underline{x}, \bar{x}]$ satisfies $\sum_k (x'_k - x_k) < \delta$, then $\sum_k |h(x'_k) - h(x_k)| < \varepsilon$. One can show that absolute continuity on compact interval $[a, b]$ implies that h has a derivative h' almost everywhere, the derivative is Lebesgue integrable, and that $h(x) = h(a) + \int_a^x h'(t)dt$ for all $x \in [a, b]$ (see, e.g., Theorem 18 in Chapter 6 of Royden and Fitzpatrick, 2010).

continuous function is the definite integral of its derivative, for any $x, z \in [\underline{x}, \bar{x}]$, we have

$$U(x) = U(z) + \int_z^x q(t)dt. \quad (3)$$

Equation (3) is often referred to as the *payoff equivalence theorem* (Börger, 2015, Lemma 2.3) because it implies that, up to the constant $U(z)$, the agent's payoff when its type is x is pinned down by the allocation rule q .²⁴ That is, all incentive-compatible mechanisms with the same allocation rule give the agent the same payoff up to an additive constant. The payoff equivalence theorem provides the important insight that the allocation rule is the key component. For example, if one is interested in efficient allocation, an agent's payoff is, up to an additive constant, pinned down. In addition, the payoff equivalence theorem has the important implication for IO applications that perfect price discrimination is ruled out. If two different types of the same agent get the same allocation, they must get the same payoff. Hence, holding the allocation fixed, the price cannot vary.

Using the payoff equivalence theorem, it follows from (3) with $z = \underline{v}$ that for a buyer, $U'(v) = q(v) \geq 0$, and so the buyer with type $\underline{v}$ is the worst-off type of buyer. Using (3) with $z = \bar{c}$, we have $U(c) = U(\bar{c}) + \int_{\bar{c}}^c q(t)dt = U(\bar{c}) - \int_c^{\bar{c}} q(t)dt$, so for a seller, who has $q(c) \leq 0$, $U'(c) = q(c) \leq 0$, implying that the seller with type $\bar{c}$ is the worst-off type of seller. It is then necessary and sufficient for individual rationality that for a buyer $U(\underline{v}) \geq 0$ and for a seller $U(\bar{c}) \geq 0$.

Equating the expressions for $U(x)$ from (1) and (3) yields the following implication of the payoff equivalence theorem:

$$m(x) = q(x)x - \int_z^x q(t)dt - U(z). \quad (4)$$

This is sometimes referred to as the *revenue equivalence theorem* (Börger, 2015, Lemma 2.4) because it states that the (expected) payment of an agent of type v is pinned down, up to a constant, by the allocation rule. Riley and Samuelson (1981) noticed an instance of this principle in their analysis of revenue equivalence among "standard auctions" with identically distributed types, where a standard auction is defined as an auction that always allocates the good to the highest bidder and gives a payoff of zero to the lowest possible type, that is, is characterized by $U(\underline{x}) = 0$.

While the analysis above shows that incentive compatibility implies that q is nondecreasing, one can also show that if q is nondecreasing, then the mechanism with allocation rule

²⁴An early instance of this theorem, under the assumption that the ex post allocation rule is efficient, was established by Holmström (1979).

q and payment rule given by (4) is incentive compatible. To see this, note that given the payment rule in (4), the payoff of an agent with type x that reports $\hat{x}$ is

$$U(x, \hat{x}) \equiv (x - \hat{x})q(\hat{x}) + \int_z^{\hat{x}} q(t)dt + U(z),$$

which means that the incremental payoff of an agent with type x that reports truthfully versus reporting $\hat{x}$ is

$$\begin{aligned} U(x) - U(x, \hat{x}) &= \int_z^x q(t)dt - (x - \hat{x})q(\hat{x}) - \int_z^{\hat{x}} q(t)dt \\ &= \int_{\hat{x}}^x q(t)dt - (x - \hat{x})q(\hat{x}) \\ &= \int_{\hat{x}}^x (q(t) - q(\hat{x}))dt, \end{aligned}$$

which is nonnegative if q is nondecreasing, establishing incentive compatibility. Thus, the mechanism $\langle q, m \rangle$ is incentive compatible if and only if q is nondecreasing and m satisfies (4) (Börger, 2015, Proposition 2.2).

If we now make use of information on the distribution of the agent's type, we can derive an expression for the agent's expected payment. For this purpose, assume that the agent's type is distributed according to distribution H with support $[\underline{x}, \bar{x}]$ and continuous, positive density h . Define the agent's *virtual value* function Φ_H and *virtual cost* function Γ_H by

$$\Phi_H(x) \equiv x - \frac{1 - H(x)}{h(x)} \quad \text{and} \quad \Gamma_H(x) \equiv x + \frac{H(x)}{h(x)}.$$

For the case of a buyer, we have $H = F$, $h = f$, and $[\underline{x}, \bar{x}] = [\underline{v}, \bar{v}]$. In this case, the relevant virtual value function is $\Phi_F \equiv \Phi$, where we drop the distribution subscript. For the case of a seller, we have $H = G$, $h = g$, and $[\underline{x}, \bar{x}] = [\underline{c}, \bar{c}]$, and the relevant virtual cost function is $\Gamma_G \equiv \Gamma$. We assume that Φ and Γ are increasing. We discuss the role of this assumption below—it is analogous to the assumption of decreasing marginal revenue and increasing marginal cost.

Given $z \in [\underline{x}, \bar{x}]$, we can write the agent's expected payment as follows (for details on

how one obtains this expression, see Appendix A.1):

$$\begin{aligned}
\mathbb{E}_x[m(x)] &= \int_{\underline{x}}^{\bar{x}} m(t)h(t)dt \\
&= \int_{\underline{x}}^z \Gamma_H(t)q(t)h(t)dt + \int_z^{\bar{x}} \Phi_H(t)q(t)h(t)dt - U(z) \\
&= \mathbb{E}_x[\Gamma_H(x)q(x) \mid x \leq z] \Pr(x \leq z) + \mathbb{E}_x[\Phi_H(x)q(x) \mid x \geq z] \Pr(x \geq z) - U(z).
\end{aligned}$$

Thus, we have the result that in any incentive-compatible direct mechanism $\langle q, m \rangle$, the agent's expected payment to the mechanism (and, symmetrically, the expected revenue to the mechanism) is pinned down by the allocation rule, virtual type functions, and type distribution, up to an additive constant.

For example, if the agent is a buyer, then its worst-off type is $\underline{v}$. Letting $z = \underline{v}$ in the expression above, the buyer's expected payment to the mechanism is

$$\mathbb{E}_v[m(v)] = \mathbb{E}_v[\Phi(v)q(v)] - U(\underline{v}).$$

If the agent is a seller, then its worst-off type is $\bar{c}$. Letting $z = \bar{c}$ implies that the seller's expected payment *from* the mechanism is

$$-\mathbb{E}_c[m(c)] = \mathbb{E}_c[\Gamma(c)|q(c)|] + U(\bar{c}).$$

Of course, by combining these expressions for the expected payment to the mechanism by buyers and the expected payment from the mechanism to sellers, we can obtain expressions for the expected budget surplus of the mechanism itself, which, as this shows, depends critically on the allocation rule. This leads us to the following discussion of optimal and efficient mechanisms.

3.3 Optimal and efficient mechanisms

Given the mechanism-design setup just discussed, we can consider key questions of what is the optimal mechanism and whether efficient mechanisms exist that satisfy incentive compatibility and individual rationality and do not run a deficit.

We begin by considering optimal mechanisms. We revisit Myerson's (1981) derivation of the optimal auction of an object for which the auctioneer has commonly known cost $c < \bar{v}$ to a buyer with private value v defined as above to be drawn from distribution F on $[\underline{v}, \bar{v}]$. We seek the mechanism $\langle q, m \rangle$ that maximizes the auctioneer's expected payoff subject to incentive compatibility and individual rationality. We know from the analysis above that we

can write the auctioneer's expected payoff from any incentive compatible mechanism $\langle q, m \rangle$ as

$$\mathbb{E}_v [m(v) - cq(v)] = \mathbb{E}_v [(\Phi(v) - c)q(v)] - U(\underline{v}).$$

If Φ is increasing, which Myerson (1981) refers to as the *regular* case, then the auctioneer optimally sets $q(v) = 1$ if $\Phi(v) \geq c$ and $q(v) = 0$ otherwise, and sets $U(\underline{v}) = 0$, which is the lowest amount such that the buyer's individual-rationality constraint is satisfied. This amounts to a take-it-or-leave-it offer by the auctioneer of $\Phi^{-1}(c)$ (or an ascending-bid auction with a reserve of $\Phi^{-1}(c)$).²⁵ Analogously, assuming that Γ is increasing, the optimal mechanism for the purchase of one unit for which the auctioneer has commonly known value $v > \underline{c}$ from a seller with private cost c drawn from distribution G on $[\underline{c}, \bar{c}]$ is a take-it-or-leave-it offer of $\Gamma^{-1}(v)$.²⁶

As observed by Mussa and Rosen (1978), virtual value functions can be interpreted as marginal revenue functions and, analogously, virtual cost functions can be interpreted as marginal cost functions. To see this, drawing on Bulow and Roberts (1989), note that the demand at price p by a buyer whose value is drawn from distribution F is $D(p) = 1 - F(p)$, and so the revenue associated with quantity Q is $R(Q) \equiv QF^{-1}(1 - Q)$. Noting that $R'(Q) = F^{-1}(1 - Q) - \frac{1 - F(F^{-1}(1 - Q))}{f(F^{-1}(1 - Q))} = \Phi(F^{-1}(1 - Q))|_{x=F^{-1}(1 - Q)} = \Phi(x)$, the assumption of decreasing marginal revenue is then equivalent to the assumption that the virtual value function Φ is increasing. Thus, Myersonian regularity, i.e., the assumption of an increasing virtual value function, corresponds to decreasing marginal revenue, or equivalently, concave revenue, and, analogously, the assumption of an increasing virtual cost function corresponds to increasing marginal cost.²⁷ Further, because a monopoly seller with a constant or increasing marginal cost function and concave revenue function optimally sets a uniform price, uniform pricing follows as a conclusion, not an assumption, from our mechanism-design setup with regularity.²⁸

To consider the existence of efficient mechanisms, let us now move to a bilateral trade setting with one buyer with private value v and one seller with private cost c . In that case, the direct mechanism $\langle Q, M_B, M_S \rangle$ maps the type vector (v, c) to a quantity traded,

²⁵For $c < \Phi(\underline{v})$, define $\Phi^{-1}(c) \equiv \underline{v}$, implying that the seller's take-it-or-leave-it offer is $\underline{v}$.

²⁶For $v > \Gamma(\bar{c})$, define $\Gamma^{-1}(v) \equiv \bar{c}$, implying that the buyer's take-it-or-leave-it offer is $\bar{c}$.

²⁷If the virtual type functions are not increasing, then more care is required to ensure that incentive compatibility holds, i.e., that the allocation rule is nondecreasing. This is where Myerson's (1981) "ironing" techniques come into play. While the expected payment $\mathbb{E}_v [m(v)]$ is characterized by the "unironed" virtual types, regardless of whether the virtual type functions are increasing, the allocation rule for the optimal mechanism must be defined in terms of "ironed" virtual types in the absence of regularity. See Section 4.4 for further discussion.

²⁸In contrast, if the monopoly has increasing marginal costs and the revenue function that it faces is not concave, then setting non-market-clearing prices may be optimal; see the discussion in Section 4.4.

$Q(v, c) \in \{0, 1\}$, a payment from the buyer, $M_B(v, c)$, and a payment to the seller $M_S(v, c)$. Of course, the mechanism is required to be incentive compatible and individually rational. As mentioned above, incentive compatibility and individual rationality are defined in terms of the agent's interim expected allocation and payment. For the buyer, those are $q_B(v) \equiv \mathbb{E}_c[Q(v, c)]$ and $m_B(v) \equiv \mathbb{E}_c[M_B(v, c)]$, and for the seller, they are $q_S(c) \equiv \mathbb{E}_v[Q(v, c)]$ and $m_S(c) \equiv \mathbb{E}_v[M_S(v, c)]$. We say that a mechanism satisfies the *no-deficit* condition if the mechanism does not run a deficit in expectation, that is, if $\mathbb{E}_{v,c}[m_B(v) - m_S(c)] \geq 0$. Under the view that a mechanism represents a market outcome, it makes sense to restrict attention to mechanisms that satisfy the no-deficit condition so that no resources have to be injected into the system, and a risk neutral intermediary would be willing to facilitate the mechanism.

Consider the question whether efficient trade can be achieved in the bilateral trade setup. That is, does there exist an incentive-compatible, individually-rational, no-deficit mechanism that has the efficient allocation rule, $Q^E(v, c) = 1$ if $v \geq c$, and $Q^E(v, c) = 0$ otherwise? Myerson and Satterthwaite (1983) provide an answer to this question: it depends on whether the supports of the distributions of v and c overlap. Economically, overlapping supports mean that trade is sometimes but not always ex post efficient. Myerson and Satterthwaite (1983) show that if $\underline{v} < \bar{c}$, i.e., the type spaces overlap, then there is no such mechanism. This is their famous *impossibility* result.

We can now see how the power of incentive compatibility comes into play to prove the Myerson-Satterthwaite impossibility result. We first prove the impossibility result for the special case in which the buyer's and seller's support are identical and, without further loss of generality, equal $[0, 1]$, and then we generalize. Define $q_B^E(v) \equiv \mathbb{E}_c[Q^E(v, c)] = G(v)$ and $q_S^E(c) \equiv \mathbb{E}_v[Q^E(v, c)] = 1 - F(c)$. It follows that the expected payment from the buyer is $\mathbb{E}_v[\Phi(v)G(v)]$ and the expected payment to the seller is $\mathbb{E}_c[\Gamma(c)(1 - F(c))]$ up to constants which must be nonnegative by individual rationality. If any such mechanism satisfies the no deficit constraint, then it is the ones with the constants equal to zero, so we proceed under that assumption. Hence, expected revenue to the mechanism is

$$\begin{aligned} \mathbb{E}_v[\Phi(v)G(v)] - \mathbb{E}_c[\Gamma(c)(1 - F(c))] &= \int_0^1 \Phi(v)G(v)f(v)dv - \int_0^1 \Gamma(c)(1 - F(c))g(c)dc \\ &= - \int_0^1 (1 - F(x))G(x)dx \\ &< 0, \end{aligned}$$

where the second equality follows from substituting the expression for Φ and Γ , noting that $G(x)f(x) - (1 - F(x))g(x) = -\frac{d}{dx}[(1 - F(x))G(x)]$, and integrating by parts. The inequality is a consequence of the integrand being positive for all $x \in (0, 1)$ and nonnegative at the

bounds. Notice that the mechanism's deficit $\int_0^1 (1 - F(x))G(x)dx$ is equal to expected social surplus, which is $\int_0^1 \int_0^v (v - c)g(c)f(v)dc dv$. For example, for F and G uniform, expected social surplus and the deficit are $1/6$.

To generalize this argument, and to see how the deficit result depends on $\underline{v} < \bar{c}$, consider the dominant-strategy implementation of the efficient allocation rule whose payment rule is given by $\hat{M}_B(v, c) \equiv \max\{c, \underline{v}\}$ if $v \geq c$ and $\hat{M}_B(v, c) \equiv 0$ otherwise, and $\hat{M}_S(v, c) \equiv \min\{v, \bar{c}\}$ if $v \geq c$ and $\hat{M}_S(v, c) \equiv 0$ otherwise.²⁹ The mechanism $\langle Q^E, \hat{M}_B, \hat{M}_S \rangle$ is incentive compatible by the usual second-price auction logic, and the mechanism satisfies individual rationality since $U(\underline{v}) = 0 = U(\bar{c})$. If $\underline{v} < \bar{c}$, then

$$\hat{M}_B(v, c) - \hat{M}_S(v, c) \leq 0,$$

with strict inequality with probability given that the distributions of v and c are continuous.³⁰ Because the dominant-strategy implementation runs a deficit, it follows that *any* incentive-compatible, individually-rational bilateral trade mechanism with the efficient allocation rule runs a deficit. (For a more formal demonstration, see Appendix A.2.) In other words, efficient bilateral trade is not possible without running a deficit.

Intuitively, the deficit for $\underline{v} < \bar{c}$ arises because the buyer and seller are complementary inputs into social surplus. By the usual Vickrey-Clarke-Groves logic (Vickrey, 1961; Clarke, 1971; Groves, 1973), each agent obtains as its payoff its social marginal product, which because of complementarity is larger with the other agent present. Hence, the sum of marginal products exceeds the total product, which is social surplus. This is maybe easiest to see with identical support, that is, if $\underline{v} = \underline{c}$ and $\bar{v} = \bar{c}$. Then each agent's marginal product is simply social surplus. Consequently, the mechanism has to pay two times social surplus and generates it only once, hence incurring a deficit equal to social surplus, which explains the observation made above that with identical supports the expected deficit is equal to expected social surplus.³¹ Notice also that if $\underline{v} \geq \bar{c}$, then trade is always efficient and $\hat{M}_B(v, c) - \hat{M}_S(v, c) = \underline{v} - \bar{c} \geq 0$. That is, ex post efficient trade is possible.³²

²⁹This argument follows closely Krishna (2010, Section 5.3.3).

³⁰To see this, notice first that if $v \leq \bar{c}$ and $c \geq \underline{v}$, we have $\hat{M}_B(v, c) - \hat{M}_S(v, c) = c - v$, which is negative given that trade is efficient. If $v > \bar{c}$ and $c \geq \underline{v}$, then $\hat{M}_B(v, c) - \hat{M}_S(v, c) = c - \bar{c}$, which is negative (unless $c = \bar{c}$) and if $v > \bar{c}$ and $c < \underline{v}$, then $\hat{M}_B(v, c) - \hat{M}_S(v, c) = \underline{v} - \bar{c}$, which is negative by assumption. Finally, if $c < \underline{v}$ and $v \leq \bar{c}$, then $\hat{M}_B(v, c) - \hat{M}_S(v, c) = \underline{v} - v$, which is negative unless $v = \underline{v}$.

³¹Generalizing Loertscher et al. (2015), Delacrétaz et al. (2019) explore this intuition systematically and show that it extends to a large range of environments with multi-unit traders. The complementarity between agents is also at the source of the deficit in public goods problems, of which, as noted by Mailath and Postlewaite (1990), the bilateral trade problem is a special case.

³²Another observation is worth making. Under the dominant strategy implementation of the efficient allocation rule, every agent is the residual claimant to the social surplus it creates. Because the implementation does not depend on the distribution from which the agent draws its type, keeping its support fixed,

The Myerson-Satterthwaite impossibility result raises the question what is the maximum expected social surplus that can be achieved in a bilateral trade setting when $\underline{v} < \bar{c}$. Chatterjee and Samuelson (1983) investigate this and analyze second-best mechanisms that maximize expected social surplus subject to incentive compatibility, individual rationality, and no deficit. As they emphasize, the efficiency of the allocation must be sacrificed in order to satisfy the no-deficit condition. This is illustrated with their k -double auction mechanism, in which the buyer submits bid p_B and the seller submits bid p_S and trade occurs if and only if $p_B \geq p_S$, in which case the buyer pays $kp_B + (1 - k)p_S$ to the seller. The mechanism satisfies the no-deficit condition by construction. For the case of agents whose types are drawn from the same uniform distribution and who play according to their linear Bayes Nash equilibrium bidding strategies, the k -double auction with $k = 1/2$ is a second-best mechanism. Any movement in k away from $1/2$ not only shifts surplus between the two agents, but also reduces the total expected surplus, highlighting the centrality of the surplus-payoff tradeoff: Agents' bargaining power affects only the division of surplus in bilateral bargaining, but also the total surplus created through bilateral trade.

3.4 Uniqueness of the properties of the independent private values setting

As we have seen, the mechanism-design setup with independent private types and risk neutral agents allows for the powerful payoff equivalence and revenue equivalence theorems. One might wonder whether even greater applicability could be obtained by venturing outside of these bounds. However, as we now discuss, the literature identifies significant challenges when one relaxes any of the core assumptions.

We have assumed a setup with single-dimensional, independent, private types drawn from continuous distributions, which implies that for a given objective, the mechanism that maximizes this objective, subject to incentive compatibility, individual rationality, and no deficit, is well defined and pinned down (up to a constant in the payments) by the allocation rule, which is unique. As discussed above, it also has the feature that, quite generally, there is a tradeoff between allocating efficiently and extracting payments, which is of particular interest for Industrial Organization. If one relaxes the assumption of private types, then agents become subject to hold up in a way that they are not when types are private, and

it thus provides the agent with incentives to invest to improve this distribution that align with those of the social planner (see, e.g., Milgrom, 2004). As explored in more detail in Loertscher and Marx (2021), efficient bargaining under incomplete information thus implies efficient investments. In other words, under efficient incomplete-information bargaining there is no hold-up of the form that is central to the inefficient investments results with complete information bargaining (see, e.g., Grossman and Hart, 1986; Hart and Moore, 1990).

the privacy of types means that first-degree price discrimination is not possible.

Dropping the assumption of risk neutrality, Maskin and Riley (1984) and Matthews (1984) show that optimal mechanisms depend on the nature of risk aversion, are not easily characterized, and may require payments to and/or from losers. Without independence, as foreshadowed by Myerson (1981), Crémer and McLean (1985, 1988) show that there is no tradeoff between profit and social surplus. Without private values, additional and, therefore, in some sense arbitrary, restrictions may be required to maintain tractability and/or the tradeoff between profit and social surplus (Mezzetti, 2004, 2007). Notwithstanding recent progress, with multi-dimensional private information and multiple agents, the optimal mechanism is not known (see, e.g., Daskalakis et al., 2017). With discrete types, there is no payoff equivalence theorem. In other words, the mechanism is not pinned down by the allocation rule.

Thus, our focus on the setup with single-dimensional, independent, private types drawn from continuous distributions and risk neutral agents is motivated not only by the power and applicability of that setup, but also by the conclusion that these are essentially the *only* assumptions that permit a tractable approach and that maintain the core tradeoff between profit and social surplus.

4 Progress to date

In this section, we review progress in analyzing agents’ power in incomplete-information IO models. We start with the definition and analysis of buyer power in Section 4.1. Then in Section 4.2, we examine the implications of bargaining power for efficiency. In Section 4.3, we illustrate these ideas and others in the context of a k -double auction. In Section 4.4, we revisit standard monopoly pricing through the lens of mechanism design. Finally, in Section 4.5, we discuss and illustrate estimation in the context of incomplete-information bargaining.

4.1 Buyer power

Notions of buyer power and countervailing power figure prominently in antitrust debates—indeed, there’s a long history, with Galbraith referring to the mitigation or regulation of economic power as the oldest of economic problems. Both buyer power and countervailing power are discussed in the context of merger review. In the United States, agencies’ merger guidelines state that “The Agencies consider the possibility that powerful buyers may constrain the ability of the merging parties to raise prices” (*US Guidelines*), and that it is possible that “buyer power would act as a countervailing factor to an increase in market

power resulting from the merger” (*EC Guidelines*). Further, the notion has traction in a litigation context: “Although the strong-buyer defense may be valid in a variety of circumstances, I believe that the courts have sometimes embraced it as if it had talismanic power that cured all doubts about a merger” (Steptoe, 1993). As an example, the buyer power of the large oil companies Shell, BP, and ExxonMobil in the purchase of oilfield services was recognized by competition authorities in their evaluation of the proposed merger of Halliburton and Baker Hughes. Let us consider what these seemingly power buyers do when they purchase inputs. The large oil companies hold auctions plus follow-on negotiations for the oilfield services that they purchase. Computer manufacturers such as Dell and HP hold online auctions plus face-to-face negotiations with suppliers of computer components. And governments often hold competitive procurements. This suggests that a natural model of buyer power involves a buyer that uses a procurement mechanism, where buyer power is the ability to use an optimal procurement.

Contributions in the theoretical literature related to buyer power include Galbraith (1952); Stigler (1954); Snyder (1996, 2008); Nocke and Thanassoulis (2014); Caprice and Rey (2015); Loertscher and Marx (2019a); Decarolis and Rovigatti (2020). Here we focus on the definition of buyer power provided in Loertscher and Marx (2019a), which, similar to Bulow and Klemperer (1996), view buyer power as a buyer’s ability to use the optimal mechanism rather than an efficient mechanism. Under this definition, one can analyze buyer power in a mechanism-design setup with one-sided private information where, pre merger, m sellers have costs drawn from distribution G and 1 buyer has a commonly known value v . We model a merger of sellers as creating a merged entity whose cost is drawn from the distribution of the minimum of two independent draws from G , which is the distribution $\hat{G}(c) \equiv 1 - (1 - G(c))^2$.³³

Having no buyer power means that the buyer holds an efficient procurement, thereby ensuring, as suggested by the quote from Ausubel et al. (2002), that no gains from trade remain unrealized. Specifically, one can think of the buyer with no power as using a second-price auction with a reserve r equal to the minimum of the buyer’s value and the upper support of the sellers’ cost distribution, $r = \min\{v, \bar{c}\}$. This says that a buyer with power can commit not to purchase even if the lowest cost is less than v , but a buyer without power cannot.

As shown in Loertscher and Marx (2019a) and as discussed below, mergers always harm the buyer, regardless of whether the buyer is powerful.³⁴ Essentially, a merger eliminates a

³³As discussed in Loertscher and Marx (2019a), most results extend to sellers who are asymmetric pre merger.

³⁴While here we focus on the so-called unilateral effects of a merger, an incomplete-information approach to analyzing additional merger harms due to coordinated effects can be found in Loertscher and Marx

bidder, and that is bad for the buyer. Without buyer power, mergers are always profitable and neutral for rivals because, regardless of the merger, the low-cost seller always trades with the buyer whenever their cost is below the buyer's value. The only change caused by the merger is that the buyer has to pay more if the bid that was eliminated by the merger was the second-lowest cost. In contrast, with buyer power, mergers are not always profitable and, with pre-merger symmetry, benefit rivals and entrants. This occurs because, with buyer power, the buyer discriminates against the merged entity because of its stronger productive power, which tips the scale towards the rivals and potential entrants.

Digging in a bit more on what the optimal procurement looks like in this setup, in the pre-merger market with pre-merger symmetry, one can think of the optimal procurement as selecting the bidder with the lowest cost and then allowing the buyer to make a take-it-or-leave-it offer to that bidder. In the post-merger market, there is no symmetry because the merged entity now draws its cost the distribution $\hat{G}$. In that case, the optimal procurement selects the bidder with the lowest virtual cost—this has the effect of handicapping the merged entity, so that it has to submit a distinctly lower bid to win—and then allows the buyer to make a take-it-or-leave-it offer to that bidder, where the buyer's optimal take-it-or-leave-it offer is lower for the merged entity than for the other sellers.

As above, denote the virtual cost function for the pre-merger sellers by Γ , and denote the virtual cost function for the merged entity by $\hat{\Gamma}$. Then the optimal take-it-or-leave-it offers are the prices p and $\hat{p}$ that equate the buyer's value with the corresponding virtual cost. Notice that for all $c \in (\underline{c}, \bar{c})$,

$$\Gamma(c) = c + \frac{G(c)}{g(c)} < c + \frac{2 - G(c)}{2 - 2G(c)} \frac{G(c)}{g(c)} = c + \frac{\hat{G}(c)}{\hat{g}(c)} = \hat{\Gamma}(c),$$

which means that the take-it-or-leave-it offer to the merged entity, $\hat{p} = \hat{\Gamma}^{-1}(v)$, is less than the take-it-or-leave-it offer to the other sellers, $p = \Gamma^{-1}(v)$. That is, the buyer negotiates more aggressively with the merged entity than with the pre-merger sellers.

This illustrates a distinction between productive power and bargaining power. A buyer without power is harmed by a merger—buyer surplus decreases because the buyer receives one fewer bid. But, interestingly, a buyer with power is also harmed by the merger. To see this, note that a buyer with power could replicate the post-merger mechanism in the pre-merger market, but because it chooses not to, by revealed preference, the buyer is better off in the pre-merger market than in the post-merger market, and so is harmed by the merger.³⁵ But as long as the merging sellers are symmetric, buyer power does act as a

(2020a).

³⁵To be more precise, as there is one more bidder pre merger, the mechanisms necessarily differ. However,

countervailing force in the sense that a buyer with power is harmed less by a merger than a buyer without power. Intuitively, a buyer with power is able to use discrimination in the post-merger market, which has asymmetric sellers, which reduces the harm experienced by a buyer with power relative to that experienced by a buyer without power, but at the same time introduces distortions in the allocation rule that reduce expected social surplus.

A limitation of analysis in Loertscher and Marx (2019a) is that sellers have no power and whether the buyer is powerful or not is, largely, treated as exogeneously given.³⁶ Consequently, there can be no possibility of a mergers of sellers causing the bargaining power of those sellers to increase, possibly countervailing the buyer’s power. And, the buyer’s power is defined to be all or nothing and not to change with a merger. This is the motivation for recent work that relaxes these assumptions, including Loertscher and Marx (2019b, 2021), which we discuss next. We first review the key results from the prior literature on which this work builds.

4.2 Bargaining power and social surplus

Given Myerson and Satterthwaite’s impossibility theorem, the efficient allocation rule is “off the table,” which means that any incentive compatible, individually rational mechanism that avoids a deficit will be inefficient. Given this, Williams (1987) notes that there can be many allocation rules such that the expected payoffs to the buyer and seller are not Pareto dominated by the expected payoffs under any other allocation rule. And, there will be an allocation rule or rules that is best for the buyer and others that are best for the seller. Thus, Williams (1987) brings the notion of bargaining power into the incomplete-information bargaining setup. Williams (1987) derives the Pareto frontier of payoff vectors that are possible in incentive-compatible, individually-rational, no-deficit, bilateral-trade mechanisms, which we refer to as the *Williams frontier*. Figure 2 illustrates the Williams frontier for the case when the agents’ types are drawn from the uniform distribution on $[0, 1]$. As the figure shows, the farther from equal division of surplus one moves, the greater is the loss of social surplus. For general distributions, Williams shows that the outcomes of the k -double auction with $k = 1$ (“buyer’s bid mechanism”) and with $k = 0$ (“seller’s price mechanism”) are on the Pareto frontier. These points correspond, respectively, to the buyer-optimal and the seller-optimal mechanism.

Keeping bargaining weights the same across all agents, Gresik and Satterthwaite (1989)

one can implement the post-merger optimal allocation rule pre merger by having the two merging suppliers compete for the right to produce whenever this allocation induces production by the merged entity; see Loertscher and Marx (2019a) for more details.

³⁶The exception pertains to a discussion showing that a merger makes a buyer more inclined to *acquire* buyer power.

show that as the number of traders increases, the inefficiency is reduced. They characterize the rate at which a bound on the expected inefficiency of the mechanism’s allocations relative to ex post efficient allocations diminishes as the number of traders increases. Of course, this has implications for how we evaluate the effects of mergers, which reduce the number of traders, in addition to potentially affecting agents’ type distributions and bargaining power. In what follows, we illustrate the combination of these effects in the context of a k -double auction.

4.3 Illustration based on the k -double auction

We now use a well-known incomplete-information model, the k -double auction of Chatterjee and Samuelson (1983), to illustrate key attributes of the incomplete-information framework. As we show, the model embodies the key payoff-versus-social-surplus tradeoff, allows one to parameterize separately firms’ productive power and their bargaining power, and offers scope for analyzing countervailing power. Further, it allows an illustration of the link between investment incentives and the efficiency of the bargaining process.

As mentioned above, the k -double auction model has one buyer with a value drawn from distribution F and one seller with a cost drawn from distribution G . Given $k \in [0, 1]$, the buyer and seller in a k -double auction simultaneously submit bids p_B and p_S , and trade occurs at the price $kp_B + (1 - k)p_S$ if and only if $p_B \geq p_S$. By construction, the k -double auction never incurs a deficit. For $k \in \{0, 1\}$, the k -double auction reduces to take-it-or-leave-it offers—by the buyer if $k = 1$ and by the seller if $k = 0$.

Payoff versus social surplus tradeoff

If the agents’ types are uniformly distributed on $[0, 1]$, then in the linear Bayes Nash equilibrium of the k -double auction, a buyer of type v bids $p_B(v) \equiv (1 - k)k/(2(1 + k)) + v/(1 + k)$, and a seller with cost c bids $p_S(c) \equiv (1 - k)/2 + c/(2 - k)$. Thus, the linear Bayes Nash equilibrium of the k -double auction results in trade if and only if $v \geq c \frac{1+k}{2-k} + \frac{1-k}{2}$. The associated expected payoffs of the buyer and seller are illustrated in Figure 2.

As illustrated in Figure 2, bargaining in the incomplete-information framework entails a payoff-surplus tradeoff. As k increases, the buyer’s (seller’s) expected payoff increases (decreases) but the sum of expected payoffs is maximized at $k = 1/2$. Indeed, as first noted by Myerson and Satterthwaite (1983), for $k = 1/2$ and uniformly distributed types, the k -double auction yields the second-best outcome. Williams (1987) noted that this insight generalizes insofar as, for uniformly distributed types and any $k \in [0, 1]$, the k -double auction implements the outcomes of incomplete-information bargaining for some bargaining weights.




Figure 2: Expected payoffs of the buyer, u_B , and the seller, u_S , in the k -double auction for $k \in [0, 1]$. Assumes that the buyer's value and the seller's cost are uniformly distributed on $[0, 1]$.

Bargaining power versus productive power and scope for countervailing power

It is straightforward to extend the uniform-uniform k -double auction model by parameterizing the distributions so that for some $a \in [0, 1/2]$, the buyer draws its value from the uniform distribution on $[a, 1]$ and the seller draws its cost from the uniform distribution on $[0, 1 - a]$. Then, as a increases from 0 to $1/2$, the productive power of both the buyer and the seller increases, and, separately, the k of the k -double auction parameterizes the agents' relative bargaining powers.

In this setup,³⁷ there is a linear Bayes Nash equilibrium in which a buyer of type v bids

$$p_B(v) \equiv \begin{cases} \frac{v}{1+k} + \frac{k(1-k)}{2(1+k)} & \text{if } v \leq \frac{1+k}{2-k}(1-a) + \frac{1-k}{2}, \\ \frac{1-a}{2-k} + \frac{1-k}{2} & \text{otherwise,} \end{cases}$$

and a seller with type c bids

$$p_S(c) \equiv \begin{cases} \frac{a}{1+k} + \frac{k(1-k)}{2(1+k)} & \text{if } c < \frac{2-k}{1+k}a - \frac{(1-k)(2-k)}{2(1+k)}, \\ \frac{c}{2-k} + \frac{1-k}{2} & \text{otherwise.} \end{cases}$$

Expected payoffs of the buyer and seller are illustrated in Figure 3.

As suggested by these results, the k of the k -double auction can be viewed as a *bargaining power* parameter, with $k = 1/2$ corresponding to equal bargaining power between the buyer

³⁷Example 2 of Chatterjee and Samuelson (1983) encompasses this setup (NB typos in the bid functions as stated in the 1983 paper).




Figure 3: Expected payoffs of the buyer, u_B , and the seller, u_S , in the k -double auction for $k \in [0, 1]$. Assumes that the buyer's value is uniformly distributed on $[a, 1]$ and that the seller's cost is uniformly distributed on $[0, 1 - a]$. The dashed lines show the second-best frontiers for different values of a .

and the seller. In contrast, the *productive power* of the agents is captured by their distributions. We illustrate this in Figure 3, which shows how productive power and bargaining power separately affect the agents' payoffs and the efficiency of bargaining.³⁸

As this shows, the incomplete-information framework allows one to capture the notion of countervailing power—an equalization of bargaining power has the potential to outweigh productive power effects. In a market with extreme bargaining power, a change in agents' productive power, for example due to a merger, that reduces social surplus might be counterbalanced by an equalization of bargaining power.

Investment incentives

To consider investment incentives, we adapt the prior k -double auction setup so that the buyer's value is drawn from the uniform distribution on $[a_B, 1]$, denoted F_{a_B} , and the seller's cost is drawn from the uniform distribution on $[0, 1 - a_S]$, denoted G_{a_S} . Suppose a_B is determined by the buyer's investment, with $a_B \in [0, 1]$ having cost $I_B(a_B)$, and that a_S is determined by the seller's investment, with $a_S \in [0, 1]$ having cost $I_S(a_S)$. Assume that

³⁸Using the Nash-in-Nash model, Crawford and Yurukoglu (2012) find that differences in bargaining power explain pricing differentials in multichannel television markets, finding that small cable operators and satellite providers have slightly less bargaining power than large cable operators. They rely on bargaining power differences to explain price differences, without allowing differences in the productive power of, for example, small and large cable operators and satellite providers. Incomplete-information models have the potential to differentiate between the two sources of price differentials.

investment occurs prior to type realizations and that the buyer and seller do not observe the investment of the other agent. For our purposes here, we assume that all optima are defined by first-order conditions.

Suppose that there is a symmetric Nash equilibrium investment a^* ; that is, in equilibrium, $a_B = a_S = a^*$. Denote by p_B^* and p_S^* the bid functions defined above but with $a = a^*$. Let $u_B^*(a_B, a_S)$ be the expected payoff of a buyer, not including investment costs, when its type is drawn from F_{a_B} , the seller's type is drawn from G_{a_S} , the seller bids according to p_S^* , and the buyer bids optimally given that it believes that the seller bids according to p_S^* and has type drawn from G_{a^*} . Analogously, let $u_S(a_B, a_S)$ be the expected payoff of a seller, not including investment costs, when its type is drawn from G_{a_S} , the buyer's type is drawn from F_{a_B} , the buyer bids according to p_B^* , and the seller bids optimally given that it believes that the buyer bids according to p_B^* and has type drawn from F_{a^*} .

By the assumptions that investments $a_B = a_S = a^*$ are a Nash equilibrium and that the rival's investment is not observed, each agent's investment maximizes its expected payoff given that the other agent bids according to its equilibrium bid function and has a type drawn from its equilibrium distribution, i.e.,

$$\left. \frac{\partial u_B(a_B, a^*)}{\partial a_B} \right|_{a_B=a^*} = I'_B(a^*) \quad \text{and} \quad \left. \frac{\partial u_S(a^*, a_S)}{\partial a_S} \right|_{a_S=a^*} = I'_S(a^*).$$

In contrast, investments $a_B = a_S = a^*$ are optimal for the planner if and only if

$$\left. \frac{\partial u_B(a_B, a^*)}{\partial a_B} \right|_{a_B=a^*} + \left. \frac{\partial u_S(a_B, a^*)}{\partial a_B} \right|_{a_B=a^*} = I'_B(a^*)$$

and

$$\left. \frac{\partial u_B(a^*, a_S)}{\partial a_S} \right|_{a_S=a^*} + \left. \frac{\partial u_S(a^*, a_S)}{\partial a_S} \right|_{a_S=a^*} = I'_S(a^*).$$

It follows that Nash equilibrium investments are optimal for the planner if and only if the allocation and payment are not affected by the investments. If the solution to the planner's problem has $a_B = a_S = a^* < 1/2$, so that the buyer's and seller's type distributions overlap, then the agents' Nash equilibrium investments are not socially optimal. However, if $a^* \geq 1/2$ and if trade occurs at the fixed price of, say, $1/2$, then each agent is the residual claimant to its investment, and Nash equilibrium investments are socially optimal. Loertscher and Marx (2021) show in a more general incomplete-information bargaining setup that first-best investments are a Nash equilibrium outcome if (and under additional conditions only if) incomplete information bargaining is efficient.

Illustration that mergers harm the buyer

Let us now focus on the case of $k = 1$ and $a = 0$, that is, assume that the buyer has all the bargaining power and types are drawn from the uniform distribution on $[0, 1]$. When the buyer faces 2 or more sellers, it holds an optimal procurement. This is consistent with the k -double auction in that when $k = 1$ and there is only one seller, the equilibrium of the k -double auction is equivalent to the buyer making a take-it-or-leave-it offer to the seller, which is the optimal procurement in that case.

As above, we assume that when the two sellers merge, the merged entity's cost is drawn from the distribution of the minimum of two independent draws from G . Supposing that the sellers draw their costs from the distribution $G(c) = 1 - \sqrt{1 - c}$ on $[0, 1]$, then the merged entity's cost distribution, $\hat{G}(c) \equiv 1 - (1 - G(c))^2 = c$, is the uniform distribution on $[0, 1]$.

Before the merger, the optimal reserve of a buyer with value $v \leq 1$ is p_B satisfying $\Gamma(p_B) = v$.³⁹ Thus, $p_B(v) = \frac{1}{9} (4 + 3v - 2\sqrt{4 - 3v})$. However, after the merger, the buyer's optimal reserve is $\hat{p}_B(v) = v/2$, which is lower for all $v \in (0, 1]$. This illustrates the point discussed above that the buyer negotiates more aggressively when facing the merged entity than when facing the two competing sellers.

Nevertheless, the buyer is harmed by the merger of sellers. Before the merger, a buyer with value 1 has an optimal reserve of $5/9$ and an expected payoff of

$$\mathbb{E}_{c_1, c_2}[(1 - \min\{\max\{c_1, c_2\}, \Gamma^{-1}(1)\}) \cdot \mathbf{1}_{\min\{c_1, c_2\} < \Gamma^{-1}(1)}] = 0.41.$$

After the merger, its optimal reserve is $1/2$, and its expected payoff is $1/4$.

Consider a buyer with a value of greater than or equal to $\max_{c \in [\underline{c}, \bar{c}]} \Gamma(c)$. Prior to the merger, the buyer benefits from competition between the sellers, paying the maximum of the two sellers' costs, which is less than $\bar{c}$ with probability one. But after the merger, the transaction price is just $\bar{c}$, implying that the buyer is harmed by the merger.

Countervailing power

In pre-merger markets in which sellers sell their product to powerful buyers, a merger among sellers has the potential to have countervailing effects (Galbraith, 1952). In particular, consolidation may increase expected social surplus if the enhanced market power of the combined seller serves to counterbalance the power of the buyers. This thinking is laid out in the Australian government's 1999 (now superseded) guidelines, which state: "If pre-merger prices are distorted from competitive levels by market power on the opposite side

³⁹For $v \in [0, 1]$, $\Gamma^{-1}(v)$ is well defined.

of the market, a merger may actually move prices closer to competitive levels and increase market efficiency. For example, a merger of buyers in a market may create countervailing power which can push prices down closer to competitive levels” (ACCC, 1999, para. 5.131).

As an illustration, reconsider the setup above with 1 buyer and 2 pre-merger sellers, where the buyer’s value is uniformly distributed on $[0, 1]$ and the pre-merger sellers’ costs are drawn from $G(c) = 1 - \sqrt{1 - c}$. If the buyer has all the bargaining power in the pre-merger market, then pre-merger expected social surplus is

$$\mathbb{E}_{v, c_1, c_2} [(v - \min\{c_1, c_2\}) \cdot \mathbf{1}_{\min\{c_1, c_2\} \leq p_B(v)}] = \frac{953}{7290} \approx 0.13,$$

where, as above, $p_B(v) = \frac{1}{9}(4 + 3v - 2\sqrt{4 - 3v})$.

In contrast, if one assumes that after the merger between the two sellers, k changes, then social surplus can be larger after merger, as shown in Figure 3 (for the case of $a = 0$). In particular, given a post-merger value of k denoted by $\hat{k}$, the post-merger expected social surplus is $\frac{1}{16}(2 - \hat{k})(1 + \hat{k})$, which is concave in $\hat{k}$ and achieves its maximum value of $\frac{9}{64} \approx 0.14$ when $\hat{k} = 1/2$. Expected social surplus increases as a result of the merger combined with the reduction in k to $\hat{k}$ for all $\hat{k} \in (0.102, 0.898)$.⁴⁰ In particular, the bargaining power of the buyer must be reduced from the pre-merger level of $k = 1$, but the buyer must retain some bargaining power in the post-merger market, so that we do not simply trade an all-powerful buyer for an all-powerful seller. This gives rise to the countervailing power result that if, as a result of the merger, the bargaining power of the buyer and seller are sufficiently equalized, then the merger-plus-reduction-in-buyer-power results in an increase in expected social surplus (and is profitable for sellers and harms the buyer). While here we focus on the k -double auction, Loertscher and Marx (2021) analyze the scope for social-surplus-increasing countervailing power in a general incomplete-information bargaining setup.

4.4 Monopoly pricing revisited

Optimal rationing and randomization

It is commonly perceived as puzzling that sellers would set below market clearing prices, randomly ration consumers, and then either try to prevent the resale that ensues from the inefficient initial allocation or at least complain about the emergence of resale.⁴¹ Would the

⁴⁰With one buyer and one seller drawing their types from the uniform distribution, $k = 1$ and $k = 0$ are equivalent for social surplus. It may seem natural to expect that for any $k \in (0, 1)$, social surplus increases. However, this is not the case because pre merger there are two sellers, and the competition between them mitigates the detrimental effects of the buyer’s bargaining power.

⁴¹See, for example, Becker (1991).

seller not be better off setting a market clearing price, whereby it would preempt resale and *make* more revenue, thereby killing two birds with one stone? Following Loertscher and Muir (2020), we now briefly revisit monopoly pricing problems and show that rationing is optimal for a monopoly only if the revenue function is not concave and, using their insights, connect this to opaque pricing of vertically differentiated products. Heuristically, this exercise also illustrates why market clearing pricing is optimal when revenue is concave.

Consider a monopoly who, for now, has K units of a homogeneous good for sale, facing a continuum of consumers with single-unit demand described by the downward sloping inverse demand function $P(Q)$ that, under market clearing pricing, gives rise to the revenue $R(Q) \equiv P(Q)Q$. As an alternative to selling K units at the market clearing price $P(K)$, suppose the seller uses a two-price mechanism with prices $p_1 > p_2$. These prices are such that in equilibrium all $Q_1 < K$ consumers with values no less than $P(Q_1)$ buy at price p_1 and are served for sure, while $Q_2 - Q_1$ with $Q_2 > K$ consumers participate in a lottery where they are served with probability $\alpha \equiv \frac{K-Q_1}{Q_2-Q_1} < 1$. (While for notational ease we keep the dependence of α on K , Q_1 , and Q_2 implicit, we emphasize this dependence as it matters below.) Making the participation for the marginal consumers bind yields $p_2 = P(Q_2)$. The incentive compatibility constraint for the consumers with value $P(Q_1)$ is $P(Q_1) - p_1 = \alpha(P(Q_1) - p_2)$, which, using $p_2 = P(Q_2)$, is equivalent to $p_1 = (1 - \alpha)P(Q_1) + \alpha P(Q_2)$. In other words, p_1 is a convex combination of $P(Q_1)$ and $P(Q_2)$. Multiplying p_1 by Q_1 and p_2 by the quantity sold at that price, which is $K - Q_1$, yields the revenue

$$(1 - \alpha)R(Q_1) + \alpha R(Q_2), \tag{5}$$

which is larger than the revenue $R(K)$ from market-clearing pricing if and only if R is not concave at K . Maximizing $(1 - \alpha)R(Q_1) + \alpha R(Q_2)$ over Q_1 and Q_2 yields the convex hull (or concavification) of R , denoted $\bar{R}(K)$, that is,

$$\bar{R}(K) \equiv \max_{Q_1, Q_2} (1 - \alpha)R(Q_1) + \alpha R(Q_2),$$

which is the smallest concave function such that $\bar{R}(K) \geq R(K)$ for all K in the domain of R . See Figure 4 for an illustration.

Denoting by F the distribution of consumers' values, we have $P(Q) = F^{-1}(1 - Q)$, and taking the derivative of $\bar{R}(K)$ yields what Myerson (1981) termed the *ironed virtual valuation*, which we denote $\bar{\Phi}(v)$, that is, $\bar{R}'(K) \equiv \bar{\Phi}(F^{-1}(1 - K))|_{v=F^{-1}(1-K)} = \bar{\Phi}(v)$. Using Myerson's arguments, one can show that rationing is not only better than uniform pricing,




Figure 4: Illustration of nonconcave revenue and the construction of ironed marginal revenue. In panel (a), the red line concavifies the revenue function. In panel (b), the shared blue area is equal to the shaded red area, and the dashed black line indicates the ironed marginal revenue function.

but is the optimal mechanism for the monopoly (see, e.g., Bulow and Roberts, 1989).⁴² Intuitively, rationing when revenue is not concave—or equivalently, marginal revenue is not monotone—is better than market clearing prices because it allows the monopoly to serve the highest marginal revenue consumers (who are not necessarily the consumers with the largest values) subject to not violating the incentive compatibility constraint that $q(v)$ must be nondecreasing. The best the monopoly can do, subject to incentive compatibility, is to serve consumers with lower values and higher marginal revenue with the same probability as it serves higher value consumers with lower marginal revenue. Thus, uniform (or random) rationing is optimal.

Note that consumers pay the price p_2 if only if they are successful in the lottery. As Loertscher and Muir (2020) observe, the monopoly can alternatively achieve the revenue $\bar{R}(K)$ by creating a vertically differentiated, inferior good that consists of $K - Q_1$ high quality units (with quality 1) and $Q_2 - K$ lower quality units (of quality 0), whose expected quality is thus α . The market clearing price for this *opaquely priced* good is $\tilde{p}_2 = \alpha P(Q_2)$, and incentive compatibility for the consumers with value $P(Q_1)$ implies that the price $\tilde{p}_1$ for the high-quality good is $\tilde{p}_1 = (1 - \alpha)P(Q_1) + \tilde{p}_2$ because the consumers can only be charged the incremental quality improvement $1 - \alpha$. Multiplying $\tilde{p}_1$ by Q_1 and $\tilde{p}_2$ by $Q_2 - Q_1$ yields the expression in (5), whose maximum is $\bar{R}(K)$. Of course, the conflated good consisting of

⁴²Without relating to mechanism-design arguments or to Myerson (1981), Wilson (1988) and Ferguson (1994) establish, respectively, the optimality of two-price mechanisms and their relation to rationing. Optimal monopoly pricing with nonconcave revenue was first analyzed by Hotelling (1931) and Mussa and Rosen (1978).

$Q_2 - K$ units of quality 0 is, in a sense, degenerately conflated. Loertscher and Muir (2020) show that the argument extends to any quality $\theta \in [0, 1)$ and to an arbitrary number of vertically differentiated goods that are available in given capacities.

Loertscher and Muir (2020) also analyze the effects of resale that ensue from randomization and show that, while it harms the seller, it does not offset all the benefits from randomization. In other words, they show that once one accounts for incomplete information (and allows for nonconcave revenue), there is no puzzle. The seller dislikes resale much in the same way that an incumbent in a game with an entrant dislikes entry, yet accommodates it because it is more profitable than deterring entry. They also show that randomization—in the form of rationing and conflation—has positive effects on the quantity that the monopoly sells in equilibrium, and that if this effect is strong enough, then randomization increases consumer surplus. By a similar logic, allowing for highly efficient resale can have an adverse quantity effect, and if this effect is strong enough, then consumers may be harmed by too efficient resale. This suggests the possibility of a consumer surplus rationale for the *Better Online Tickets Sales (BOTS)* Act of 2016 in the United States, which reduced the efficiency of resale without banning it. The reason for why there can be large quantity effects of randomization is that randomization renders the optimized revenue function of the seller concave. Hence, the optimal quantity sold varies smoothly in the parameters of the problem. In contrast, when the optimized revenue is not concave (either because randomization is prohibited or made less effective, which is the typical effect of resale), the revenue-maximizing quantity can vary discontinuously.

Of course, the practical relevance of these insights hinges on whether or not revenue is concave. Empirically, when concavity of revenue is not imposed, the hypothesis that revenue is concave (or equivalently, that virtual values are monotone) is regularly rejected; see, for example, Larsen et al. (2020, Section 5), Larsen and Zhang (2018, Appendix D), and Celis et al. (2014).⁴³ Loertscher and Muir (2020) provide theoretical reasons—such as the integration of markets—as to why revenue may not be concave. Maybe more fundamentally, even though it is standard to assume that revenue is concave, this does not make it so, and there are simply no theoretical reasons for revenue to be concave.

Digital monopolies

As noted for example by Loertscher and Marx (2020b), larger markets are better, all else equal, because they can execute the same trades as smaller, standalone markets, and some-

⁴³Celis et al. (2014) analyze advertising bidding data from Microsoft Advertising Exchange and find that, to a large extent because, empirically, virtual values are not monotone, one can do better in terms of revenue and consumer surplus with mechanisms other than a second-price auction with an optimal reserve.

times execute more or more valuable trades. This feature of increasing returns to scale in market making is inherent in the independent private values model, in which the market maker’s payoff function (be it social surplus, profit, a weighted average of the two, or social surplus under the second-best mechanism) increases more than linearly in market size. More precisely, denote by $W(\mathcal{N})$ the expected value of the market maker’s objective when there is a set $\mathcal{N}$ of buyers and sellers, each drawing their values and costs independently from some distributions F and G with overlapping supports. Then for two disjoint and finite sets $\mathcal{N}$ and $\mathcal{N}'$, we have

$$W(\mathcal{N}) + W(\mathcal{N}') < W(\mathcal{N} \cup \mathcal{N}'),$$

because, reflecting Williamson’s (1985) notion of selective intervention, the market maker can always replicate what it would do in the stand-alone markets and sometimes do better.⁴⁴ This basic insight is of relevance for the ongoing antitrust debates concerning big tech mentioned in Section 2 because it provides a simple and coherent rationale for why, in principle, gigantic market makers can emerge without engaging in anticompetitive conduct. Indeed, when firms are primarily market makers, or otherwise subject to increasing returns to scale (for example, because of the value of aggregating data), efficiency may dictate that there is a single active firm. Borrowing from the notion of natural monopolies, Loertscher and Marx (2020b) refer to firms in such environments as *digital monopolies*.

The analogy is also useful because it suggests that seemingly “old-school”, brick-and-mortar approaches that proved useful in dealing with natural monopolies may also be productive in the digital age. In particular, when increasing returns to scale loom large, breaking up large firms may deprive society of these benefits of increased market thickness, and price regulation may be an effective instrument to distribute these gains more evenly across society. It is not clear, a priori, why it should not be possible to regulate, say, the fee Apple charges for selling products via its App Store or to restrict the algorithms that Amazon and others use to price (and target) the goods that they sell.

To illustrate the problems that may arise in the starkest and simplest terms, consider a platform in the presence of two users that generates a value of $V > 0$ if both users join (and, say, forego their privacy) and has zero value otherwise. V could be the social value of the platform or the profit it generates to the platform. The cost of joining the platform for each user i with $i \in \{1, 2\}$ is c_i , which is net of the private benefit that the platform generates for the user. So, if V is the social value of the platform, it is efficient for the platform to operate

⁴⁴The insight that larger markets perform better has been prominently explored by Gresik and Satterthwaite (1989), who show that the performance of the second-best mechanism relative to first-best improves as market size increases, and in the subsequent double-auctions literature. However, the insight is not restricted to this class of mechanisms or to a specific objective.

if and only if $\sum_{i=1}^2 c_i < V$.

To see that a competitive (or Walrasian) price p that induces the efficient allocation may fail to exist, notice that for the platform to be willing to acquire both inputs/attract both users (assuming for simplicity that it appropriates the entire social value V if it operates), it must be that $V - 2p > 0$, assuming that the outside option of the platform is 0. For user i to be willing to sell at price p , it must be that $p - c_i > 0$, where 0 is the value of the outside option of the user. Taken together, this means that a market clearing price that supports the efficient allocation exists if and only if $V - 2\max\{c_1, c_2\} > 0$, while efficiency dictates that the platform should operate whenever $V - (c_1 + c_2) > 0$. Thus, whenever

$$2\max\{c_1, c_2\} > V > c_1 + c_2,$$

the platform should operate, but there is no competitive price that supports that outcome.

Not surprisingly, once one has to account for the agents' private information about their costs, the obstacles do not necessarily become smaller. To see this, suppose the two agents are privately informed about their costs c_i of joining the platform and consider first the dominant-strategy implementation of the efficient allocation rule. As above, let the value of their outside option be 0. To make the problem interesting, let $\bar{c}$ be the largest possible cost for each agent and assume that $V < 2\bar{c}$, so that it is not always efficient to operate the platform. Let $p_i(\mathbf{c})$ be the price offered to i when the costs are $\mathbf{c} = (c_1, c_2)$. If $\sum_{i=1}^2 c_i \geq V$, then it is efficient not to operate the platform, and so the smallest price that can be offered to the agents, without violating their individual rationality constraints, is 0. If $\sum_{i=1}^2 c_i < V$, then it is efficient to operate the platform, and each agent i is offered a price equal to the highest cost it could report without changing the allocation, that is, $p_i(\mathbf{c}) = \min\{V - c_j, \bar{c}\}$, where $j \neq i$. It is straightforward to establish, that whenever the platform operates under efficiency, the payments to the agents exceed the value of the platform, that is, $\sum_{i=1}^2 p_i(\mathbf{c}) > V$.⁴⁵ Thus, this efficient mechanism that satisfies dominant strategies and agents' (ex post) individual rationality constraints always runs a deficit when it is efficient to operate the platform.⁴⁶ If we assume that the agents' costs are independently distributed on the interval $[\underline{c}, \bar{c}]$ according to some distributions G_1 and G_2 with positive densities, then it follows that *any* efficient,

⁴⁵To see this, assume first that $V - c_i < \bar{c}$ for $i \in \{1, 2\}$. In this case, $\sum_{i=1}^2 p_i(\mathbf{c}) > V$ is equivalent to $2V - (c_1 + c_2) > V$, which is equivalent to $V > c_1 + c_2$, which is the case because it is efficient to operate the platform. Next, assume that $V - c_i \geq \bar{c}$ for $i \in \{1, 2\}$, in which case, $\sum_{i=1}^2 p_i(\mathbf{c}) > V$ is equivalent to $2\bar{c} > V$, which is true by assumption. Last, assume that $V - c_i < \bar{c} \leq V - c_j$ for $i \neq j$, in which case $\sum_{i=1}^2 p_i(\mathbf{c}) > V$ is equivalent to $\bar{c} > c_i$, which is the case.

⁴⁶Individual rationality constraints are said to be satisfied *ex post* if agents are better off participating even after the payments have been and the allocation is determined. It contrasts with interim individual rationality, which only stipulates that agents are better off participating at the time they report their types.

incentive compatible and individually rational mechanism runs a deficit.⁴⁷ Admittedly, the assumption that the agents are perfect complements (and that there are only two agents) is strong.⁴⁸ However, it is useful as it makes the point that the competitive benchmark cannot be taken for granted for pricing when there are strong complementarities among agents, which is, for example, the case in data-driven industries where one consumer’s data are primarily valuable in conjunction with data from other consumers.

Of course, the main activity of many platforms and big tech companies is not necessarily that they match literally buyers and sellers and gain from the trades they execute. This is certainly the case for companies such as Google, Facebook, and Twitter, which create value for their users by facilitating certain flows of information. Accordingly, a major policy concern in regards to companies like these is how these platforms steer information flows, to what extent they create “echo chambers,” facilitate the spread of conspiracy theories, what (if anything) are the limits to freedom of speech, and the like. Our discussion sheds little light on these important aspects of the problem other than perhaps suggesting that the old-school approach of subjecting social media to qualitatively similar rules and restrictions to those that were imposed when radio and TV broadcasting emerged in the first half of the 20th century may be preferable to dismantling social media networks, in particular if the returns to scale that are the driving force behind them are so strong that as soon as one network is dismantled another one pops up.

4.5 Empirical work

Estimation in the context of incomplete-information bargaining

We now turn to empirical estimation of incomplete-information bargaining and illustrate the recent work by Larsen and Zhang (2018) mentioned in Section 2. They make use of the power of the revelation principle and the implications of incentive compatibility to develop estimation methods tailored to a general incomplete-information bargaining setups.

To illustrate the estimation approach of Larsen and Zhang (2018), consider the case of a bilateral bargaining. Suppose that you have repeated observations on bargaining outcomes between a buyer and a seller (not necessarily the same buyer and seller over time, but assumed to draw their types from the same distribution over time). You do not observe the exact mechanism through which bargaining occurs, but you do observe reports by agents (think of these as the agents’ initial offers), and you observe bargaining outcomes. Specifically, for

⁴⁷This is, of course, reflective of the public goods problem and, more generally, the issues arising from complementarities mentioned in and around footnote 31.

⁴⁸One can show that the deficit problem remains with an arbitrary number of agents n , provided $n\bar{c} > V$ and provided the agents are perfect complements.

each bargain j , you observe actions a_j^B and a_j^S , an indicator x_j that is 1 if there is trade and 0 otherwise, and transfer t_j from the buyer to the seller.

For example, if the underlying mechanism is a k -double auction with $k = 3/4$ and the underlying type distributions are the uniform distribution on $[0, 1]$, then one might observe the following dataset:

j	a_j^B	a_j^S	x_j	t_j
1	0.30	0.20	1	0.28
2	0.43	0.49	0	0
3	0.59	0.52	1	0.57
4	0.28	0.57	0	0
5	0.34	0.16	1	0.30

Note that one should not expect to observe buyer bids less than $\frac{1-k}{2} = \frac{1}{8}$ or seller bids greater than $\frac{2-k}{2} = \frac{5}{8}$ because such bids have zero probability of resulting in trade in equilibrium.

Letting $p_B(v)$ be the (unobserved) mapping from buyer types to actions and $p_S(c)$ be the (unobserved) mapping from seller types to actions, one can use this data to estimate the inverses p_B^{-1} and p_S^{-1} using the methodology of Larsen and Zhang (2018). Specifically, for $a \in \mathbb{R}$, let

$$\hat{\beta}^B(a) \equiv (\hat{\beta}_0^B(a), \hat{\beta}_1^B(a), \hat{\beta}_2^B(a)) \equiv \arg \min_{\beta} \sum_j \left(\left(t_j - \sum_{\ell=0}^2 \beta_{\ell} (a_j^B - a)^{\ell} \right)^2 \kappa_{h^t}(a_j^B - a) \right)$$

and

$$\hat{\gamma}^B(a) \equiv (\hat{\gamma}_0^B(a), \hat{\gamma}_1^B(a), \hat{\gamma}_2^B(a)) \equiv \arg \min_{\gamma} \sum_j \left(\left(x_j - \sum_{\ell=0}^2 \gamma_{\ell} (a_j^B - a)^{\ell} \right)^2 \kappa_{h^x}(a_j^B - a) \right),$$

where κ_h is a kernel functions with bandwidth h (see Larsen and Zhang (2018) for a discussion of appropriate kernel and bandwidth choices). Then $\hat{\beta}^B(a)$ and $\hat{\gamma}^B(a)$ are defined for any a not included as a buyer action in the dataset, and, as shown by Larsen and Zhang (2018), for any such a , a consistent estimator of $p_B^{-1}(a)$ is given by

$$\hat{p}_B^{-1}(a) \equiv \frac{\hat{\beta}_1^B(a)}{\hat{\gamma}_1^B(a)}.$$

One can then extend the definition of $\hat{p}_B^{-1}(a)$ to all buyer actions using linear interpolation. Defining things analogously for sellers, and using $-t_j$ instead of t_j and $-x_j$ instead of x_j ,

$$\hat{p}_S^{-1}(a) \equiv \frac{\hat{\beta}_1^S(a)}{\hat{\gamma}_1^S(a)}$$

provides a consistent estimate of $p_S^{-1}(a)$ for any a not included as a seller action in the dataset. Similarly, this can be extended to all seller actions using linear interpolation.

Given these estimates of the inverse strategy mappings, we can estimate the underlying types and construct the empirical distributions of those types using $\hat{v}_j \equiv \hat{p}_B^{-1}(a_j^B)$ and $\hat{c}_j \equiv \hat{p}_S^{-1}(a_j^S)$. We illustrate the results of such estimation in Figure 5, with the details contained in Appendix A.3.




Figure 5: Empirical distributions of estimated type compared with the theoretical distributions. Estimates are based on the observation of 4000 randomly drawn outcomes of a k -double auction with $k = 3/4$ and actual types drawn from the uniform distribution on $[0, 1]$, assuming agents use their linear Bayes Nash equilibrium strategies. In this case, buyers with values less than $1/8$ and sellers with costs greater than $5/8$ do not participate, and so the relevant theoretical distributions are the uniform distributions truncated at those types, as shown.

Because this approach allows one to estimate the underlying types, one can then use those types to estimate the agents' expected payoffs by using the average payoffs in the data. That is, given J total observations, we have

$$\hat{u}_B \equiv \frac{1}{t} \sum_{j=1}^J (\hat{v}_j x_j - t_j) \quad \text{and} \quad \hat{u}_S \equiv \frac{1}{t} \sum_{j=1}^J (t_j - \hat{c}_j x_j),$$

where it is important to note that the theoretical expected payoffs would also “average in” some zeros because we only observe agents with types sufficiently “efficient” that they trade

with positive probability. In the simulation underlying Figure 5, we get, $\hat{u}_B = 0.1430$, as compared with the relevant theoretical benchmark of $u_B = 0.1458$, which is the expected payoff of a buyer in a k -double auction with $k = 3/4$ and uniformly distributed values on $[\frac{1}{8}, 1]$ and uniformly distributed costs on $[0, \frac{5}{8}]$. For the sellers, we get $\hat{u}_S = 0.0995$ as compared with the theoretical value of $u_S = 0.1041$.

One can further use the methods of Loertscher and Marx (2020) to estimate the virtual types. Specifically, letting $\hat{v}_{(j)}$ denote the j -th highest estimated value and, letting $\hat{c}_{[j]}$ denote the j -th lowest estimated cost, uniformly consistent estimators for $\frac{1-F(v_{(j)})}{f(v_{(j)})}$, where F is the theoretical value distribution, and $\frac{G(c_{[j]})}{g(c_{[j]})}$, where G is the theoretical cost distribution, respectively, are given by

$$\hat{h}_j^B \equiv j\sigma_j^B \quad \text{and} \quad \hat{h}_j^S \equiv j\sigma_j^S,$$

where σ_j^B and σ_j^S are the average spacing between nearby types:

$$\sigma_j^B \equiv \frac{\hat{v}_{(j-\min\{r,j-1\})} - \hat{v}_{(j+\min\{r,J-j\})}}{\min\{r,j-1\} + \min\{r,J-j\}} \quad \text{and} \quad \sigma_j^S \equiv \frac{\hat{c}_{[j+\min\{r,J-j\}]} - \hat{c}_{[j-\min\{r,j-1\}]}}{\min\{r,J-j\} + \min\{r,j-1\}}$$

where r is the number of spacings on each side of a type that are included in the spacing estimate, defined to be $J^{4/5}$, rounded to the nearest integer.⁴⁹

One can then use $(\hat{v}_{(j)}, \hat{h}_j^B)_{j=1}^J$ to define the estimate $\hat{h}^B(v)$ of $\frac{1-F(v)}{f(v)}$ for all v in the relevant range using linear interpolation.⁵⁰ Similarly, one can use $(\hat{c}_{[j]}, \hat{h}_j^S)_{j=1}^J$ to define the estimate $\hat{h}^S(c)$ of $\frac{G(c)}{g(c)}$ for all c in the relevant range using linear interpolation.⁵¹ We illustrate in Figure 6.

Given these estimates, for any bargaining weights $(w_B, w_S) \in [0, 1]^2$, with at least one weight being positive, and any $\rho \geq 1$, we can define the allocation rule $Q^{\mathbf{w}, \rho}$ as follows:

$$Q^{\mathbf{w}, \rho}(v, c) \equiv \begin{cases} 1 & \text{if } v - (1 - \frac{w_B}{\rho})\hat{h}^B(v) \geq c + (1 - \frac{w_S}{\rho})\hat{h}^S(c), \\ 0 & \text{otherwise.} \end{cases}$$

Given $\mathbf{w}$, define the empirical budget surplus by

$$\hat{\pi}^{\mathbf{w}, \rho} \equiv \frac{1}{J} \sum_{j=1}^J \left[\left(\hat{v}_j - \hat{h}^B(\hat{v}_j) \right) - \left(\hat{c}_j + \hat{h}^S(\hat{c}_j) \right) \right] Q^{\mathbf{w}, \rho}(\hat{v}_j, \hat{c}_j)$$

⁴⁹See Loertscher and Marx (2020) for a discussion of the choice of the number of spacings to include in the estimator.

⁵⁰Interpolation is only used for the purposes of creating figures—in our estimates we only evaluate $\hat{h}^B$ and $\hat{h}^S$ at estimated types.

⁵¹One might be able to improve estimates by constraining $\hat{h}^B$ and $\hat{h}^S$ to be positive. We have not done this for the purposes of our illustration.




Figure 6: Estimated hazard rates based on the observation of 4000 randomly drawn outcomes of a k -double auction with $k = 3/4$ and actual types drawn from the uniform distribution on $[0, 1]$, assuming agents use their linear Bayes Nash equilibrium strategies. In this case, buyers with values less than $1/8$ and sellers with costs greater than $5/8$ do not participate, and so the relevant theoretical distributions are the uniform distributions truncated at those types.

and define $\rho^{\mathbf{w}}$ to be the smallest $\rho \geq 1$ such that $\hat{\pi}^{\mathbf{w}, \rho} \geq 0$.

Then, letting $\eta_B^{\mathbf{w}} \equiv 1$ if $w_B > w_S$, $\eta_B^{\mathbf{w}} \equiv 0$ if $w_B < w_S$, and $\eta_B^{\mathbf{w}} \equiv 1/2$ if $w_B = w_S$, and analogously for $\eta_S^{\mathbf{w}}$ (i.e., $\eta_S^{\mathbf{w}} \equiv 1 - \eta_B^{\mathbf{w}}$), estimated payoffs in the incomplete-information bargaining mechanism with bargaining weights $\mathbf{w}$ are

$$\hat{u}_B(\mathbf{w}) \equiv \eta_B^{\mathbf{w}} \hat{\pi}^{\mathbf{w}, \rho^{\mathbf{w}}} + \frac{1}{J} \sum_{j=1}^J \hat{h}^B(\hat{v}_j) Q^{\mathbf{w}, \rho^{\mathbf{w}}}(\hat{v}_j, \hat{c}_j)$$

and

$$\hat{u}_S(\mathbf{w}) \equiv \eta_S^{\mathbf{w}} \hat{\pi}^{\mathbf{w}, \rho^{\mathbf{w}}} + \frac{1}{J} \sum_{j=1}^J \hat{h}^S(\hat{c}_j) Q^{\mathbf{w}, \rho^{\mathbf{w}}}(\hat{v}_j, \hat{c}_j).$$

As before, these estimates do not account for payoffs of zero associated with types that have no probability of trading, which are missing from our dataset.

One can then find the $\mathbf{w}$ that minimize the distance between $(\hat{u}_S(\mathbf{w}), \hat{u}_B(\mathbf{w}))$ and $(\hat{u}_S, \hat{u}_B)$, giving us an estimate $\hat{\mathbf{w}}$ of the agents' bargaining weights. We illustrate this in Figure 7(a). As shown there, we find $(\hat{w}_B, \hat{w}_S) = (1, 0.5)$, consistent with the data being simulated based on a k -double auction with $k = 3/4$. To explore the precision of the estimation procedure and the effects of sample size, we calculate bootstrap confidence intervals for the agents' expected payoffs for sample sizes of 1,000 and 10,000, as shown in Figure 7(b). The figure compares the bootstrap 95% confidence intervals to the theoretical values for the agents' expected payoffs.




Figure 7: Estimated Williams frontier and bootstrap confidence intervals for agents’ expected payoffs compared with theoretical values. Panel (a): The red point labeled “data” is $(\hat{u}_S, \hat{u}_B)$. The small black points along the frontier are $(\hat{u}_S(\mathbf{w}), \hat{u}_B(\mathbf{w}))$ for a grid of values for $\mathbf{w}$ with one bargaining weight equal to 1 and the other an element of $\{0, 0.1, \dots, 1\}$. The green point labeled with bargaining weights is the grid point that is closest to the red point. Estimated types and hazard rates are based on the observation of 4000 randomly drawn outcomes (Larsen and Zhang (2018) also perform estimates for a sample size of 4000) of a k -double auction with $k = 3/4$ and actual types drawn from the uniform distribution on $[0, 1]$, assuming agents use their linear Bayes Nash equilibrium strategies. Panel (b): 95% confidence intervals based on estimates of buyer and seller expected surplus for 100 independently drawn simulated datasets of 1000 observations each and, separately, for 100 independently drawn simulated datasets of 10,000 observations each for a k -double auction with $k = 3/4$ and actual types drawn from the uniform distribution on $[0, 1]$, assuming agents use their linear Bayes Nash equilibrium strategies; u_B and u_S are the associated theoretical values.

Thus, as we have just illustrated, data on bids and outcomes are sufficient to allow the identification of agents’ type distributions and bargaining weights associated with incomplete-information bargaining. Although the k -double auction setup is special and results in that setup depend on specific distributions, this illustration shows that the incomplete-information framework provides a tractable, intuitive approach to: the payoff-surplus trade-off, the separation of bargaining and productive power, countervailing power, and merger harms.

Bargaining breakdown on the equilibrium path

As mentioned above, standard complete-information bargaining models focus on the division of a fixed surplus, not the size of the surplus. In those models, when there are gains from trade, trade always occurs. In contrast, in incomplete-information bargaining with over-

lapping type distributions, negotiations break down on the equilibrium path with positive probability. This occurs for two reasons. First, it may be that the buyer’s value is below the seller’s cost, but because of private information, the two parties do not know this before they sit down at the negotiating table. Second, by the Myerson-Satterthwaite impossibility result, even if the buyer’s value exceeds the seller’s cost, the constraints imposed by incentive compatibility, individual rationality, and no deficit may prevent ex post efficient trade from taking place.

The possibility of bargaining breakdown that arises in incomplete-information bargaining captures features of reality that complete-information models cannot accommodate. Our common experience suggests that bargaining breakdown is common—we observe strikes and blackouts and government shutdowns. More systematically, Backus et al. (2020) examine millions of eBay transactions involving bilateral bargaining and document the prevalence of bargaining breakdown. Data suggest that approximately half of negotiations that are initiated end without a transaction. This, of course, is not a stylized fact that can be accommodated by complete information bargaining models such as Nash bargaining. However, in an incomplete-information context, as we have said, bargaining breakdown arises naturally. Indeed, the observation of bargaining breakdown in data can be used to calibrate distributions in an incomplete-information setup.

For example, under the assumption that incomplete-information bargaining maximizes expected social surplus subject to incentive compatibility, individual rationality, and no deficit, and that the buyer and seller are treated symmetrically, one can use the observed frequencies of negotiation breakdowns to back out the parameters of the distributions from which the buyer and the seller draw their types. To illustrate, assume that the buyer’s value v is drawn from the distribution $F(v) = 1 - (1 - v)^{1/\kappa}$ and the seller’s cost c is drawn from the distribution $G(c) = c^{1/\kappa}$, whose supports are $[0, 1]$, where $\kappa \in (0, \infty)$ has the interpretation of a “capacity.” Figure 8 plots the probability that negotiations break down as a function of κ under the assumptions stated. For example, if, as in the data set of Backus et al. (2020), 55 percent of all negotiations break down, eyeballing the figure indicates that κ must be around 1.5.⁵² Rather than treating negotiation breakdowns as measurement error as in the case with complete information, with incomplete information the frequency of those breakdowns is valuable information that can be used for estimation.

The suitability of our approach for applied work is reinforced by the ability to calibrate straightforward parameterizations of the agents’ type distributions based on data such as the

⁵²More precisely, for the case considered, a breakdown probability of 55 percent corresponds to $\kappa = 1.60904$.




Figure 8: Probability of the breakdown of negotiations in incomplete-information bargaining as a function of κ assuming $F(v) = 1 - (1 - v)^{1/\kappa}$ and $G(c) = c^{1/\kappa}$ on $[0, 1]$.

observed probability with which negotiations break down.⁵³ Given the calibrated distributions, the bargaining solution is easily computed. For example, with one buyer and one seller, each drawing their types from the uniform distribution on $[0, 1]$,⁵⁴ incomplete-information bargaining modeled as a k -double auction with $k = 1/2$ gives each agent an expected surplus of $9/128 \approx 0.07$, which is slightly less than the expected surplus under Nash bargaining of $1/12 \approx 0.08$. One can also accommodate general distributions, asymmetries, and one-to-many bargaining (Loertscher and Marx, 2021).

Of course, it may be that in many instances empirical analysts do not have the luxury of observing bargaining breakdown. But this does not necessarily mean that what is not observed does not matter (just like it would be hard to argue that nuclear weapons played no role during the cold war on the grounds that nuclear war did not occur). Methods to account for unobserved bargaining breakdown might advance progress here.

5 Open issues and the road forward

Of course, there are many open issues, and a lot of work remains to be done. In this section, we highlight a few of these issues and of the possible avenues for future work.

⁵³For one-to-many bargaining, the market shares of the competing agents can also be used for calibration.

⁵⁴A buyer and seller each drawing its type from the uniform distribution on $[0, 1]$ corresponds to a breakdown probability of approximately 72%, which is greater than the 50% breakdown probability in an efficient bargaining setup with the same type distributions.

5.1 Determinants of bargaining power

As discussed above, the mechanism-design based approach to incomplete-information bargaining allows one to incorporate differential bargaining power for agents and, given data, to estimate the relative bargaining power of agents. Questions remains, however, regarding what determines this bargaining power in the first place. This is particularly pertinent given the recent upsurge in interest in bargaining. Avenues for future work includes the analysis of how bargaining power varies with market structure and how it depends on firm size, as well as other things.

A possible avenue to pursue this more systematically is touched upon in Loertscher and Marx (2019a), who discuss the incentives of the buyer to acquire, at some cost, buyer power and show that these incentives are stronger post merger. This suggests that one could more generally ask whether merging parties have stronger incentives to increase their bargaining weight before or after the merger.

5.2 Mergers that create multi-unit traders

As discussed above, the incomplete-information IO framework loses its power if agents' types are not single dimensional. One can easily work within those bounds for a fixed market structure. In particular, the setup accommodates multiple buyers with multi-unit demands, if a buyer's marginal value for each unit is a linear function of the agent's single-dimensional type; and there can be multiple sellers with multi-unit capacities, if a seller's marginal cost for each unit is a linear function of the agent's single-dimensional type.

Modeling mergers is then straightforward in some cases. For example, if the total quantity demanded by the buyers in the market is D , and if sellers i and j each have constant marginal cost (equal to their types) up to a capacity of D and, then a merger of sellers i and j can be modeled as creating a merged entity that has capacity D and a constant marginal cost equal to the minimum of two independent cost draws, one from seller i 's type distribution and one from seller j 's type distribution. This scenario applies in the case of one buyer with single-unit demand purchasing from some number of sellers analyzed above, or more generally when multiple agents on one side of the market trade with a single agent with single-unit demand or capacity on the other side of the market.

Modeling challenges arise when the merged entity might trade more units than either of the pre-merger agents can produce or consume. For example, suppose that there is one buyer with two-unit demand and there are two sellers, each with one-unit capacity. Let G_1 and G_2 be the cost distributions of the two pre-merger sellers. The merged entity might be expected to supply two units for some type realizations, but how should one model the merged entity's

costs for those two units? A way to transform G_1 and G_2 into a one-dimension distribution is required to maintain the features of the Myersonian mechanism design approach. A number of possibilities suggest themselves. Noting that the joint supply by the two pre-merger sellers at price p is

$$S(p) = G_1(p) + G_2(p) = 2 \frac{G_1(p) + G_2(p)}{2},$$

one could model the merged entity as having two-unit supply at a constant marginal cost drawn from the distribution $\tilde{G}(c) \equiv \frac{G_1(c) + G_2(c)}{2}$. Or, one could view the merged entity as having a type $\hat{c}$ drawn from the distribution of the minimum of c_1 and c_2 , given by $\hat{G}(c) \equiv 1 - (1 - G_1(c))(1 - G_2(c))$, and then viewing that type $\hat{c}$ as the merged entity's marginal cost for its first unit and letting $a\hat{c}$ be the merged entity's marginal cost for its second unit for some commonly known $a \geq 1$. For example, one could choose a to be the ratio of the expected value of $\max\{c_1, c_2\}$ to the expected value of $\min\{c_1, c_2\}$.

As we now discuss, additional, related challenges present themselves for the case of vertical mergers.

5.3 Vertical integration

As in the case of horizontal mergers discussed above, one can easily accommodate vertical integration within the incomplete-information setup for some market structures, assuming that after integration, the integrated firm can efficiently solve its internal agency problem, which is a standard assumption.⁵⁵ For example, suppose that the pre-integration market has one buyer with single-unit demand and value drawn from distribution F and multiple sellers with costs drawn from distribution G . The integration of the buyer with seller 1 creates an entity that has demand for one unit, which can either be sourced internally or externally. The integrated firm never acts as a seller because there are no external buyers. Thus, the integrated firm's type that is relevant for trade with the external sellers is the integrated firm's willingness to pay for an external unit. Viewing the downstream component of the integrated firm as having a willingness to pay of v and the upstream component as having a cost to supply of c_1 , the willingness to pay of the integrated firm for an external unit is then $\min\{v, c_1\}$. So one can naturally view the integrated firm as drawing a single-dimensional willingness to pay for an external unit from the distribution of the minimum of v and c_1 . Things are similarly straightforward if the pre-integration market has only one seller with single-unit supply. In that case the willingness of the integrated firm to sell in the external market is the maximum of the types of the integrated buyer and integrated seller.

⁵⁵This assumption can be rationalized, for example, on the grounds that integration slackens the individual rationality constraints within the integrated firm.

Interestingly, as shown in Loertscher and Marx (2021), this immediately gives us the result that vertical integration can, under general conditions, cause an efficient market to no longer be efficient. To see this, consider the case of nonoverlapping supports with $\bar{c} \leq \underline{v}$ and suppose that there are two suppliers and one buyer in the pre-integration market. That is, efficient trade is possible prior to integration—for example, a second-price auction would suffice—but following integration, the vertically integrated firm’s willingness to pay for an external unit is $\min\{v_1, c_1\} = c_1$, which has overlapping support with the cost of the external supplier, so the Myerson-Satterthwaite impossibility result applies. Efficient trade is possible prior to integration, but following integration, there is no incentive-compatibility, individually-rational, no-deficit mechanism that implements the efficient allocation.

Turning to the case of a pre-integration market with multiple buyers and multiple suppliers, challenges arise because the integrated firm might act as either a buyer or a seller in the external market, depending on type realizations. Vertical integration creates a market in which at least one agent, the vertically integrated firm, cannot be categorized *ex ante* as either a buyer or a seller. Following Loertscher and Marx (2020c), we refer to such a market as an *asset market*.⁵⁶ While many models presume that markets have two sides—agents are either buyers or sellers with *ex ante* determined trading positions—in asset markets some agents’ trading positions—buy, sell, do not trade—are determined endogeneously; see, e.g., Lu and Robert (2001); Chen and Li (2018); Johnson (2019); Loertscher and Marx (2020c). In practice, the gig-economy has resulted in formerly two-sided markets becoming asset markets, with examples such as eBay, Uber, and AirBnB.

Interest in accommodating asset markets adds to the appeal of the incomplete-information approach because it offers scope for productively modeling asset markets. In contrast, for standard oligopoly models with price-taking consumers on one side of the market, it is not clear how one would even shift bargaining power between the buyer side and seller side of the market, much less how one would adapt such models to handle *ex ante* nonidentified traders.

Extending the incomplete-information framework to accommodate flexibly asset markets would allow analysis of mergers between intermediaries who take positions and, depending on circumstances, either buy or sell. This would pave the way for the analysis of settings that have “three sides” in the sense that there are buyers, sellers, and integrated firm that can be either buyers or sellers.

⁵⁶Lu and Robert (2001), and Chen and Li (2018) following them, refer to these markets as markets with “*ex ante* unidentified traders.” More recently, Johnson (2019) uses the term “one-sided” to refer to these problems.

6 Conclusions

We show that the incomplete-information model is tractable for capturing central issues in Industrial Organization and Antitrust, such as the tradeoff between social surplus and profit and the effects of buyer power, bargaining power, and countervailing power. Moreover, the model predicts bargaining break down on the equilibrium path; implies that perfect price discrimination is not possible because, holding fixed the quantity a buyer is allocated, the price cannot vary with the buyer's willingness to pay; provides conditions on the primitives that imply that uniform pricing by a monopoly is optimal; and gives rise, without any additional assumptions, to increasing return to scale in market making. This chapter provides an overview of the relevant literature, explains the key mechanics, and lays out a roadmap for the work ahead.

A Appendix

A.1 Expected payment details

Given $z \in [\underline{v}, \bar{v}]$, we have:

$$\begin{aligned}
& \mathbb{E}_v [m(v)] \\
&= \int_{\underline{v}}^{\bar{v}} m(v) f(v) dv \\
&= \int_{\underline{v}}^{\bar{v}} \left(q(v)v - \int_z^v q(x) dx \right) f(v) dv - U(z) \\
&= \int_{\underline{v}}^z \left(q(v)v + \int_v^z q(x) dx \right) f(v) dv + \int_z^{\bar{v}} \left(q(v)v - \int_z^v q(x) dx \right) f(v) dv - U(z) \\
&= \int_{\underline{v}}^z q(v)v f(v) dv + \int_{\underline{v}}^z \int_{\underline{v}}^x q(x) f(v) dv dx + \int_z^{\bar{v}} q(v)v f(v) dv - \int_z^{\bar{v}} \int_x^{\bar{v}} q(x) f(v) dv dx - U(z) \\
&= \int_{\underline{v}}^z q(v)v f(v) dv + \int_{\underline{v}}^z q(x) F(x) dx + \int_z^{\bar{v}} q(v)v f(v) dv - \int_z^{\bar{v}} q(x)(1 - F(x)) dx - U(z) \\
&= \int_{\underline{v}}^z \left(x + \frac{F(x)}{f(x)} \right) q(x) f(x) dx + \int_z^{\bar{v}} \left(x - \frac{1 - F(x)}{f(x)} \right) q(x) f(x) dx - U(z) \\
&= \int_{\underline{v}}^z \Gamma(x) q(x) f(x) dx + \int_z^{\bar{v}} \Phi(x) q(x) f(x) dx - U(z) \\
&= \mathbb{E}_v [\Gamma(v)q(v) \mid v \leq z] \Pr(v \leq z) + \mathbb{E}_v [\Phi(v)q(v) \mid v \geq z] \Pr(v \geq z) - U(z),
\end{aligned}$$

where the first equality uses the definition of the expectation, the second uses (4), the third through sixth rearrange, including switching the order of integration, the seventh uses the definition of the virtual cost and virtual value functions Γ and Φ , and the last equality uses the definition of the expectation.

A.2 Details of the demonstration of the Myerson-Satterthwaite impossibility result

Suppose that $\underline{v} < \bar{c}$ and define $\hat{M}_B(v, c) \equiv \max\{c, \underline{v}\}$ if $v \geq c$ and $\hat{M}_B(v, c) \equiv 0$ otherwise, and $\hat{M}_S(v, c) \equiv \min\{v, \bar{c}\}$ if $v \geq c$ and $\hat{M}_S(v, c) \equiv 0$ otherwise. The mechanism $\langle Q^E, \hat{M}_B, \hat{M}_S \rangle$ is incentive compatible by the usual second-price auction logic: an agent's report affects only whether it trades, not its payment if it does trade, and trade is only an agent's interest when it is efficient, so agents cannot profit from misreporting their types. In addition, the mechanism satisfies individual rationality with equality: a buyer with value $\underline{v}$ is the worst-off

type of buyer and has an expected payoff of zero because it trades with probability zero, and similarly for a seller with cost $\bar{c}$.

We can now apply the revenue equivalence theorem. Suppose that there exists a payment rule $(\tilde{M}_B, \tilde{M}_S)$ such that mechanism $\langle Q^E, \tilde{M}_B, \tilde{M}_S \rangle$, with interim expected payoff function for the buyer of $\tilde{U}_B$ and for the seller of $\tilde{U}_S$, is incentive compatible and individually rational. Then we have

$$\begin{aligned} \mathbb{E}_{v,c} [\tilde{m}_B(v) - \tilde{m}_S(c)] &= \mathbb{E}_{v,c} [(\Phi(v) - \Gamma(c)) Q^E(v, c)] - \tilde{U}_B(\underline{v}) - \tilde{U}_S(\bar{c}) \\ &\leq \mathbb{E}_{v,c} [(\Phi(v) - \Gamma(c)) Q^E(v, c)] \\ &= \mathbb{E}_{v,c} [\hat{m}_B(v) - \hat{m}_S(c)] \\ &< 0, \end{aligned}$$

where the first equality uses the revenue equivalence theorem applied to $\langle Q^E, \tilde{M}_B, \tilde{M}_S \rangle$, the first inequality uses the implication of individual rationality that $\tilde{U}_B(\underline{v}) \geq 0$ and $\tilde{U}_S(\bar{c}) \geq 0$, the second equality uses the revenue equivalence theorem applied to $\langle Q^E, \hat{M}_B, \hat{M}_S \rangle$ and the fact that the interim expected payoffs of the worst-off types are zero in that mechanism, and the final inequality follows because under mechanism $\langle Q^E, \hat{M}_B, \hat{M}_S \rangle$, the payment by the buyer, $\max\{c, \underline{v}\}$ is less than or equal to the payment to the seller, $\min\{v, \bar{c}\}$, and strictly so for an open interval of types when $\underline{v} < \bar{c}$. Thus, when the set of possible buyer values overlaps with the set of possible seller costs, $\underline{v} < \bar{c}$, *any* incentive-compatible, individually-rational bilateral trade mechanism with the efficient allocation rule runs a deficit. In other words, efficient bilateral trade is not possible without running a deficit.

A.3 Details of the estimation illustration

We begin by considering the theoretical application of the k -double auction, assuming that agents' types are uniformly distributed on $[0, 1]$ and that agents use their linear Bayes Nash equilibrium strategies, where a buyer of type v bids $p_B(v) \equiv (1 - k)k/(2(1 + k)) + v/(1 + k)$, and a seller with cost c bids $p_S(c) \equiv (1 - k)/2 + c/(2 - k)$.

Let (P_B, T_B) be the menu faced by the buyer, where $P_B(b)$ is the probability that the buyer trades given a bid b and $T_B(b)$ is the expected payment by the buyer to the mechanism given a bid b . Then we have

$$P_B(b) = p_S^{-1}(b)$$

and

$$T_B(b) = \int_0^{p_S^{-1}(b)} (kb + (1 - k)p_S(c)) dc.$$

As required by incentive compatibility, the menu defines a convex frontier as shown in Figure 9(a).




Figure 9: Menu of expected allocations and transfers as a function of bids for a k -double auction with $k = 3/4$ and types drawn from the uniform distribution on $[0, 1]$, assuming agents use their linear Bayes Nash equilibrium strategies.

Analogously, let (P_S, T_S) be the menu faced by the buyer, where $-P_S(b)$ is the probability that the seller trades given a bid b and $-T_S(b)$ is the expected payment to the seller from the mechanism given a bid b . Then we have

$$P_S(b) = -(1 - p_B^{-1}(b))$$

and

$$T_S(b) = - \int_{p_B^{-1}(b)}^1 (kp_B(v) + (1 - k)b)dv.$$

Again, as required by incentive compatibility, the menu defines a convex frontier as shown in Figure 9(b).

Simulated data

We generate a large number of pairs of uniform $[0, 1]$ values and costs, discarding those where the value is less than $\frac{1}{8}$ or seller the cost is greater than $\frac{5}{8}$, so that we ultimately have 4000 type vectors. From these, we calculate the bids of the buyer and seller according to p_B and p_S with $k = \frac{3}{4}$ and calculate the associated allocation and transfers.

Bandwidth estimation

We use the data to calculate the bandwidth for use in the kernel estimation procedure.⁵⁷ Following this procedure, we find bandwidths between 0.05 and 0.06.

Estimation of inverse strategy for the buyer

We now estimate the first derivative of T_B using the local polynomial regression estimator for $T_B(a)$ at a given bid a with degree 2, bandwidth h_T , and the Gaussian kernel. We use the value of $\hat{\beta}_1(a)$ derived from

$$(\hat{\beta}_0^B(a), \hat{\beta}_1^B(a), \hat{\beta}_2^B(a)) \equiv \arg \min_{\beta} \sum_{j=1}^{4000} \left(\left(t_j - \sum_{\ell=0}^2 \beta_{\ell} (a_j^B - a)^{\ell} \right)^2 \kappa_{h_T}(a_j^B - a) \right).$$

Similarly, we estimate the first derivative of P_B using the value of $\hat{\gamma}_1(a)$ derived from

$$(\hat{\gamma}_0^B(a), \hat{\gamma}_1^B(a), \hat{\gamma}_2^B(a)) \equiv \arg \min_{\gamma} \sum_{j=1}^{4000} \left(\left(x_j - \sum_{\ell=0}^2 \gamma_{\ell} (a_j^B - a)^{\ell} \right)^2 \kappa_{h_P}(a_j^B - a) \right).$$

Then we estimate $\hat{p}_B^{-1}(b)$ using $\frac{\hat{\beta}_1^B(a)}{\hat{\gamma}_1^B(a)}$, and similarly for $\hat{p}_S^{-1}(b)$.

Further, $\hat{\beta}_1^B(a)$ and $\hat{\gamma}_1^B(a)$ are estimates of the first derivatives of T_B and P_B , respectively, so we can use those to reconstruct T_B and P_B and calculate the (P_B, T_B) frontier, which we confirm is convex.

⁵⁷To do this, regress the vector of 4000 transfers on a quintic polynomial of the observed buyer bids and calculate the variance of the residuals σ_T^2 . Let $\hat{\hat{T}}'''$ denote the third derivative of the fitted equation. Define bandwidth

$$h_T \equiv C \left(\frac{\sigma_T^2}{\sum_{j=1}^{4000} \left(\hat{\hat{T}}'''(a_j^B) \right)^2} \right)^{1/7},$$

where $C = 0.884$ is a constant for the Gaussian kernel. Analogously, define the bandwidth for the allocation rule by

$$h_P \equiv C \left(\frac{\sigma_P^2}{\sum_{j=1}^{4000} \left(\hat{\hat{P}}'''(a_j^B) \right)^2} \right)^{1/7}.$$




Figure 10: Estimated and theoretical inverse bid functions for a k -double auction with $k = 3/4$ and types drawn from the uniform distribution on $[0, 1]$, assuming agents use their linear Bayes Nash equilibrium strategies.

References

- AUSTRALIAN COMPETITION & CONSUMER COMMISSION (1999): “Merger Guidelines,” <https://www.accc.gov.au/>.
- (2019): “Digital Platforms Inquiry: Final Report,” <https://www.accc.gov.au/focus-areas/inquiries-ongoing/digital-platforms-inquiry>.
- AUSUBEL, L. M., P. CRAMTON, AND R. J. DENECKERE (2002): “Bargaining with Incomplete Information,” in *Handbook of Game Theory*, ed. by R. Aumann and S. Hart, Elsevier Science B.V., vol. 3, 1897–1945.
- BACKUS, M., T. BLAKE, B. LARSEN, AND S. TADELIS (2020): “Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions,” *Quarterly Journal of Economics*, 135, 1319–1361.
- BACKUS, M., T. BLAKE, AND S. TADELIS (2019): “On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining,” *Journal of Political Economy*, 127, 1599–1628.
- BAYE, M. R. AND J. D. WRIGHT (2011): “Is Antitrust Too Complicated for Generalist Judges? The Impact of Economic Complexity and Judicial Training on Appeals,” *Journal of Law and Economics*, 54, 1–24.
- BECKER, G. S. (1991): “A Note on Restaurant Pricing and Other Examples of Social Influences on Price,” *Journal of Political Economy*, 99, 1109–1116.
- BESSEN, J. E. AND M. J. MEURER (2006): “Patent Litigation with Endogenous Disputes,” *American Economic Review*, 96, 77–81.
- BÖRGERS, T. (2015): *An Introduction to the Theory of Mechanism Design*, Oxford University Press.
- BULOW, J. AND P. KLEMPERER (1996): “Auctions versus Negotiations,” *American Economic Review*, 86, 180–194.
- BULOW, J. AND J. ROBERTS (1989): “The Simple Economics of Optimal Auctions,” *Journal of Political Economy*, 97, 1060–90.
- BYRNE, D. P., L. A. MARTIN, AND J. S. NAH (2019): “Price Discrimination, Search, and Negotiation in an Oligopoly: A Field Experiment in Retail Electricity,” Working Paper, University of Melbourne.
- CAPPS, C. AND D. DRANOVE (2004): “Hospital Consolidation and Negotiated PPO Prices,” *Health Affairs*, 23, 87–109.
- CAPPS, C., D. DRANOVE, AND M. SATTERTHWAITE (2003): “Competition and Market Power in Option Demand Markets,” *RAND Journal of Economics*, 34, 737–763.

- CAPRICE, S. AND P. REY (2015): “Buyer Power from Joint Listing Decision,” *Economic Journal*, 125, 1677–1704.
- CAVES, R. E., H. CROOKELL, AND J. P. KILLING (1983): “The Imperfect Market for Technology Licenses,” *Oxford Bulletin of Economics and Statistics*, 249–267.
- CELIS, L. E., G. LEWIS, M. MOBIUS, AND H. NAZERZADEH (2014): “Buy-It-Now or Take-a-Chance: Price Discrimination through Randomized Auctions,” *Management Science*, 60, 2927 – 2948.
- CÉMER, J., Y.-A. DE MONTJOYE, AND H. SCHWEITZER (2019): “Competition policy for the digital era,” Final Report, European Commission.
- CHATTERJEE, K. AND W. SAMUELSON (1983): “Bargaining under Incomplete Information,” *Operations Research*, 31, 835–851.
- CHEN, Y.-C. AND J. LI (2018): “Revisiting the Foundations of Dominant-Strategy Mechanisms,” *Journal of Economic Theory*, 178, 294–317.
- CHOI, J. P. AND H. GERLACH (2017): “A Theory of Patent Portfolios,” *American Economic Journal: Microeconomics*, 9, 315–351.
- CLARKE, E. (1971): “Multipart Pricing of Public Goods,” *Public Choice*, 11, 17–33.
- COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. S. LEE (2019): “Nash-in-Nash: A Microfoundation for Applied Work,” *Journal of Political Economy*, 127, 163–195.
- CRAIG, S., M. GRENNAN, AND A. SWANSON (2020): “Mergers and Marginal Costs: New Evidence on Hospital Buyer Power,” Working Paper, University of Pennsylvania.
- CRAWFORD, G., R. LEE, M. WHINSTON, AND A. YURUKOGLU (2018a): “The Welfare Effects of Vertical Integration in Multichannel Television Markets,” *Econometrica*, 86.
- (2018b): “The Welfare Effects of Vertical Integration in Multichannel Television Markets,” *Econometrica*, 86, 891–954.
- CRAWFORD, G. AND A. YURUKOGLU (2012): “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, 102, 643–685.
- CRAWFORD, G. S. (2014): “Cable Regulation in the Internet Era,” in *Economic Regulation and Its Reform: What Have We Learned?*, ed. by N. L. Rose, University of Chicago Press, 137–193.
- CRÉMER, J. AND R. MCLEAN (1985): “Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent,” *Econometrica*, 53, 345–361.
- (1988): “Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1257.

- CUTLER, D. M. AND F. SCOTT MORTON (2013): “Hospitals, Market Share, and Consolidation,” *JAMA*, 310, 1964–1970.
- DAFNY, L. (2009): “Estimation and Identification of Merger Effects: An Application to Hospital Mergers,” *Journal of Law and Economics*, 52, 523–550.
- DASKALAKIS, C., A. DECKELBAUM, AND C. TZAMOS (2017): “Strong Duality for a Multiple-Good Monopolist,” *Econometrica*, 85, 735–767.
- DECAROLIS, F. AND G. ROVIGATTI (2020): “From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising,” Working Paper, Bocconi University.
- DELACRÉTAZ, D., S. LOERTSCHER, L. MARX, AND T. WILKENING (2019): “Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade,” *Journal of Economic Theory*, 416–454.
- EPSTEIN, R. J. AND D. L. RUBINFELD (2004): “Effects of Mergers Involving Differentiated Products,” Technical Report, COMP/B1/2003/07.
- EUROPEAN COMMISSION (2004): “Guidelines on the Assessment of Horizontal Mergers Under the Council Regulation on the Control of Concentrations Between Undertakings,” *Official Journal of the European Union* 2004/C 31/03.
- FERGUSON, D. G. (1994): “Shortages, Segmentation, and Self-Selection,” *The Canadian Journal of Economics*, 27, 183–197.
- FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, MIT Press, Cambridge, Massachusetts.
- FURMAN, J., A. FLETCHER, D. COYLE, AND D. M. P. MARSDEN (2019): “Unlocking Digital Competition,” Report of the Digital Competition Expert Panel.
- GAL-OR, E. AND A. DUKES (2003): “Minimum Differentiation in Commercial Media Markets,” *Journal of Economics & Management and Strategy*, 12, 291–325.
- (2006): “On the Profitability of Media Mergers,” *Journal of Business*, 79, 489–525.
- GALBRAITH, J. K. (1952): *American Capitalism: The Concept of Countervailing Power*, Houghton Mifflin, Boston.
- GANGLMAIR, B., L. M. FROEB, AND G. J. WERDEN (2012): “Patent Hold-Up and Antitrust: How a Well-Intentioned Rule Could Retard Innovation,” *Journal of Industrial Economics*, 60, 249–273.
- GAYNOR, M. AND R. J. TOWN (2012): “Competition in Health Care Markets,” in *Handbook of Health Economics*, ed. by T. G. McGuire, M. V. Pauly, and P. Barros, Elsevier, vol. 2, 449–637.
- GERSHKOV, A., J. K. GOEREE, A. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): “On

- the Equivalence of Bayesian and Dominant Strategy Implementation,” *Econometrica*, 81, 197–220.
- GOWRISANKARAN, G., A. NEVO, AND R. TOWN (2015): “Mergers When Prices Are Negotiated: Evidence from the Hospital Industry,” *American Economic Review*, 105, 172–203.
- GREEN, E. J. AND R. H. PORTER (1984): “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, 52, 87–100.
- GRENNAN, M. (2013): “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review*, 103, 145–177.
- (2014): “Bargaining Ability and Competitive Advantage: Empirical Evidence from Medical Devices,” *Management Science*, 60, 3011–3025.
- GRENNAN, M. AND A. SWANSON (2020): “Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining,” *Journal of Political Economy*, 128, 1234–1268.
- GRESIK, T. AND M. SATTERTHWAIT (1989): “The Rate At Which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms,” *Journal of Economic Theory*, 48, 304–332.
- GROSSMAN, S. J. AND O. D. HART (1986): “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy*, 94, 691–719.
- GROVES, T. (1973): “Incentives in Teams,” *Econometrica*, 41, 617–631.
- GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 68, 525–574.
- HART, O. AND J. MOORE (1990): “Property Rights and the Nature of the Firm,” *Journal of Political Economy*, 98, 1119–1158.
- HO, K. AND R. LEE (2017): “Insurer Competition in Health Care Markets,” *Econometrica*, 85, 379–417.
- (2019): “Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets,” *American Economic Review*, 109, 473–522.
- HOLMSTRÖM, B. (1979): “Groves’ Scheme on Restricted Domains,” *Econometrica*, 47, 1137–1144.
- HOLMSTRÖM, B. AND R. B. MYERSON (1983): “Efficient and Durable Decision Rules with Incomplete Information,” *Econometrica*, 51, 1799–1819.
- HORN, H. AND A. WOLINSKY (1988): “Bilateral Monopolies and Incentives for Merger,” *RAND Journal of Economics*, 19, 408–419.
- HOTELLING, H. (1931): “The Economics of Exhaustible Resources,” *Journal of Political*

- Economy*, 39, 137–175.
- IVALDI, M., B. JULLIEN, P. REY, P. SEABRIGHT, AND J. TIROLE (2007): “The Economics of Tacit Collusion: Implications for Merger Control,” in *The Political Economy of Antitrust*, ed. by V. Ghosal and J. Stennek, Elsevier, 217–239.
- JOHNSON, T. R. (2019): “One-Sided Markets,” Working Paper, University of Notre Dame.
- KANKANHALLI, G. AND A. KWAN (2019): “Bargaining Power in the Market for Intellectual Property: Evidence from Licensing Contract Terms,” Working Paper, Cornell University.
- KATZ, M. L. AND C. SHAPIRO (1985): “On the Licensing of Innovations,” *RAND Journal of Economics*, 16, 504–520.
- KLEMPERER, P. (2002): “What Really Matters in Auction Design,” *Journal of Economic Perspectives*, 16, 169–189.
- (2004): “Using and Abusing Auction Theory,” in *Auctions: Theory and Practice*, ed. by P. Klemperer, Princeton: Princeton University Press, 123–148.
- KRISHNA, V. (2002): *Auction Theory*, Elsevier Science, Academic Press.
- (2010): *Auction Theory*, Elsevier, second ed.
- KUHN, T. S. (1970): *The Structure of Scientific Revolutions*, University of Chicago Press.
- LARSEN, B. J. (2020): “The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions,” *Review of Economic Studies*.
- LARSEN, B. J., C. H. LU, AND A. L. ZHANG (2020): “Mediation in Bargaining: Evidence from Large-Scale Field Data on Business-to-Business Negotiations,” Working Paper, Stanford University.
- LARSEN, B. J. AND A. ZHANG (2018): “A Mechanism Design Approach to Identification and Estimation,” NBER Working Paper 24837.
- LEMUS, J. AND E. TEMNYALOV (2017): “Patent Privateering, Litigation, and R&D Incentives,” *RAND Journal of Economics*, 48, 1004–1026.
- LEON, R. (2018): “U.S. v. AT&T Inc.” Civil Case No. 17-2511 (RJL), Memorandum Opinion, June 12, 2018.
- LEWIS, M. S. AND K. E. PFLUM (2015): “Diagnosing Hospital System Bargaining Power in Managed Care Networks,” *American Economic Journal: Economic Policy*, 7, 243–274.
- LOERTSCHER, S., L. MARX, AND T. WILKENING (2015): “A Long Way Coming: Designing Centralized Markets with Privately Informed Buyers and Sellers,” *Journal of Economic Literature*, 53, 857–897.
- LOERTSCHER, S. AND L. M. MARX (2019a): “Merger Review for Markets with Buyer Power,” *Journal of Political Economy*, 127.

- (2019b): “Merger Review with Intermediate Buyer Power,” *International Journal of Industrial Organization*, 67, 1–16.
- (2020): “Asymptotically Optimal Prior-Free Clock Auctions,” *Journal of Economic Theory*, 187.
- (2020a): “Coordinated Effects in Merger Review,” Working Paper, University of Melbourne.
- (2020b): “Digital Monopolies: Privacy Protection or Price Regulation?” *International Journal of Industrial Organization*, 71.
- (2020c): “A Dominant-Strategy Asset Market Mechanism,” *Games and Economic Behavior*, 120, 1–15.
- (2021): “Incomplete Information Bargaining with Applications to Mergers, Investment, and Vertical Integration,” Working Paper, University of Melbourne.
- LOERTSCHER, S. AND E. V. MUIR (2020): “Monopoly Pricing, Optimal Randomization, and Resale,” Working Paper, University of Melbourne.
- LU, H. AND J. ROBERT (2001): “Optimal Trading Mechanisms with Ex Ante Unidentified Traders,” *Journal of Economic Theory*, 97, 50–80.
- MAILATH, G. J. AND A. POSTLEWAITE (1990): “Asymmetric Information Bargaining Problems with Many Agents,” *Review of Economic Studies*, 57, 351–367.
- MASKIN, E. AND J. RILEY (1984): “Optimal Auctions with Risk Averse Buyers,” *Econometrica*, 52, 1473–1518.
- MATTHEWS, S. (1984): “On the Implementability of Reduced Form Auctions,” *Econometrica*, 52, 1519–1522.
- MEZZETTI, C. (2004): “Mechanism Design with Interdependent Valuations: Efficiency,” *Econometrica*, 72, 1617–1626.
- (2007): “Mechanism Design with Interdependent Valuations: Surplus Extraction,” *Economic Theory*, 31, 473–488.
- MILGROM, P. (2004): *Putting Auction Theory to Work*, Cambridge University Press.
- MUSSA, M. AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301–317.
- MYERSON, R. AND M. SATTERTHWAITE (1983): “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory*, 29, 265–281.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- NEVO, A. (2014): “Mergers that Increase Bargaining Leverage,” Remarks

- at the Stanford Institute for Economic Policy Research, Jan. 22, 2014, <https://www.justice.gov/atr/file/517781/download>.
- NOCKE, V. AND J. THANASSOULIS (2014): “Vertical Relations under Credit Constraints,” *Journal of the European Economic Association*, 12, 337–367.
- PRAGER, E. AND M. SCHMITT (forth.): “Employer Consolidation and Wages: Evidence from Hospitals,” *American Economic Review*.
- RAMSEY, F. P. (1927): “A Contribution to the Theory of Taxation,” *The Economic Journal*, 37, 47–61.
- REY, P. AND T. VERGÉ (2019): “Secret Contracting in Multilateral Relations,” TSE Working Paper.
- RILEY, J. G. AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.
- ROYDEN, H. AND P. FITZPATRICK (2010): *Real Analysis*, Prentice Hall, fourth ed.
- SAMUELSON, W. (1985): “A Comment on the Coase Theorem,” in *Game-Theoretic Models of Bargaining*, ed. by A. E. Roth, New York: Cambridge University Press, 321–339.
- SHAPIRO, C. (2010): “Injunctions, Hold-Up, and Patent Royalties,” *American Law and Economics Review*, 12, 280–318.
- (2018): “Expert Report of Carl Shapiro,” U.S. v. AT&T Inc., DirecTV Group Holdings, LLC, and Time Warner Inc., Case No. 1:17-cv-02511 (RJL), <https://www.justice.gov/atr/case-document/file/1081336/download>.
- SIDAK, J. G. (2015): “Bargaining Power and Patent Damages,” *Stanford Technology Law Review*, 19, 1–31.
- SNYDER, C. M. (1996): “A Dynamic Theory of Countervailing Power,” *RAND Journal of Economics*, 27, 747–769.
- (2008): “Countervailing Power,” in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, London: Palgrave Macmillan, 1188–1189.
- SORENSEN, A. (2003): “Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut,” *Journal of Industrial Economics*, 51, 469–490.
- SPENGLER, J. J. (1950): “Vertical Integration and Antitrust Policy,” *Journal of Political Economy*, 58, 347–352.
- STEPTOE, M. L. (1993): “The Power-Buyer Defense in Merger Cases,” *Antitrust Law Journal*, 61, 493–504.
- STIGLER, G. J. (1954): “The Economist Plays with Blocs,” *American Economic Review (Papers and Proceedings)*, 44, 7–14.

- STIGLER CENTER (2019): “Stigler Committee on Digital Platforms,” Final Report.
- SWEETING, A., X. TAO, AND X. YAO (2018): “Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Mergers,” Working Paper, University of Maryland.
- U.S. DEPARTMENT OF JUSTICE AND THE FEDERAL TRADE COMMISSION (2010): “Horizontal Merger Guidelines,” <https://www.justice.gov/atr/horizontal-merger-guidelines-08192010>.
- (2020): “Vertical Merger Guidelines,” https://www.ftc.gov/system/files/documents/reports/us-department-justice-federal-trade-commission-vertical-merger-guidelines/vertical_merger_guidelines_6-30-20.pdf.
- VICKREY, W. (1961): “Counterspeculation, Auction, and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- WERDEN, G. J. (2008): “Unilateral Competitive Effects of Horizontal Mergers I: Basic Concepts and Models,” in *Issues in Competition Law and Policy*, ABA Section of Antitrust Law, vol. II, 1319–1341.
- WERDEN, G. J. AND M. L. FROEB (2008): “Unilateral Competitive Effects of Horizontal Mergers,” in *Handbook of Antitrust Economics*, ed. by P. Buccirossi, MIT Press, 43–104.
- WILLIAMS, S. R. (1987): “Efficient Performance in Two Agent Bargaining,” *Journal of Economic Theory*, 41, 154–172.
- WILLIAMSON, O. E. (1985): *The Economic Institutions of Capitalism*, New York: Free Press.
- WILSON, C. A. (1988): “On the Optimal Pricing Policy of a Monopolist,” *Journal of Political Economy*, 96, 1645–176.
- WILSON, R. (1987): “Game Theoretic Analysis of Trading Processes,” in *Advances in Economic Theory*, ed. by T. Bewley, Cambridge, U.K.: Cambridge University Press, 33–70.
- (1993): *Nonlinear Pricing*, Oxford University Press.