

Incomplete Information Bargaining with Applications to Mergers, Investment, and Vertical Integration[†]

By SIMON LOERTSCHER AND LESLIE M. MARX*

We provide an incomplete information bargaining framework that captures the effects of differential bargaining power in markets with multiple buyers and multiple suppliers. The market is modeled as a mechanism that maximizes the expected weighted welfare of the firms, subject to the constraints of incentive compatibility, individual rationality, and no deficit. We show that, in this model, there is no basis for the presumption that vertical integration increases equally weighted social surplus, while it is possible that horizontal mergers that appropriately change bargaining weights increase social surplus. Moreover, efficient bargaining implies that in equilibrium noncontractible investments are efficient. (JEL C78, D82, D83, G34, K21, L22, L40)

Bargaining has come to the forefront in industrial organization and antitrust. It plays a prominent role in recent cases, including in health care, telecommunications, mass media, and patents. Common practice in modeling bargaining is to assume that the firms have complete information about each other's values and costs and to adhere to axiomatic approaches based on Nash bargaining or the Shapley value according to which bargaining outcomes are efficient. Apart from bargaining losing "much of its interest" when information is complete (Fudenberg and Tirole 1991), the complete information approach with efficient bargaining has the downside that shifts of bargaining power, perhaps due to a merger, or more generally changes in market structure, only affect the distribution of surplus and not its size since

*Loertscher: Department of Economics, University of Melbourne (email: simonl@unimelb.edu.au); Marx: Duke University (email: marx@duke.edu). Liran Einav was the coeditor for this article. We are grateful to three anonymous reviewers and two coeditors of this journal whose comments and suggestions have helped us to improve the paper substantially. We also thank Matt Backus, Allan Collard-Wexler, Nitin Dua, Sean Ennis, Soheil Ghili, Brad Larsen, Robin Lee, Joao Montez, Volker Nocke, Rob Porter, Patrick Rey, George Rozanski, Glen Weyl, Matthijs Wildenbeest, Steve Williams, and seminar audiences at Bocconi University, Boston University, Columbia University, CREST/École Polytechnique, Harvard University, MIT, New York University, University of East Anglia, University of Melbourne, Australasian Microeconomic Theory Seminars, Australasian Meeting of the Econometric Society, Conference on Mechanism & Institution Design, Interactive Online IO Seminar, CEPR Virtual IO Seminar Series, 3rd BEET Workshop, 33rd Summer Conference on Industrial Organization: Advances in Competition Policy, 13th Competition Law, Economics and Policy Conference in Pretoria, MaCCI Workshop on Mergers and Antitrust, 19th IOCC, and 13th Northwestern Antitrust Conference for valuable discussions and comments. Edwin Chan, Toan Le and Bing Liu provided excellent research assistance. Financial support from the Samuel and June Hordern Endowment, European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement 340903, University of Melbourne Faculty of Business and Economics Eminent Research Scholar Grant, and Australian Research Council Discovery Project Grant DP200103574 is also gratefully acknowledged.

[†]Go to <https://doi.org/10.1257/aer.20201092> to visit the article page for additional materials and author disclosure statements.

bargaining is, by assumption, efficient. Of course, the popularity of the complete information bargaining approach is in no small part due to the perceived challenges associated with the alternative of incomplete information bargaining, such as a lack of tractability of extensive-form representations and the dependence of bargaining outcomes on higher-order beliefs and assumptions of common knowledge of type distributions.

In this paper, we develop an incomplete information bargaining framework that sidesteps the lack of tractability of extensive-form games by taking an “as-if” approach in which allocations and transfers are on the Pareto frontier achievable through mediated mechanisms. Specifically, we stipulate that there is a market mechanism that, for given bargaining weights, maximizes the weighted sum of the firms’ surplus, subject to the constraints that the mechanism is incentive compatible and individually rational and does not run a deficit. For the case of one buyer and one supplier with equal bargaining weights, our model specializes to the bilateral trade problem of Myerson and Satterthwaite (1983).

We apply this framework to analyze several long-standing questions in antitrust. We show that with incomplete information bargaining, there is no basis for a presumption that vertical integration increases social surplus.¹ The intuition is simple and related to the fact that whether incomplete information bargaining is efficient is endogenous. In a nutshell, vertical integration can create a Myerson-Satterthwaite problem by rendering hitherto efficient bargaining inefficient. More generally, because changes in bargaining weights and market structures have direct effects on the social surplus resulting from incomplete information bargaining, the framework opens scope for a countervailing power defense of, say, horizontal mergers that appropriately shift bargaining powers, or more generally the analysis of policies that equalize bargaining powers. Although the concept that power on one side of a market could neutralize power on the other side has been controversial since its inception,² it has popular appeal and has influenced antitrust policies and regulation.³ Our paper provides a framework that permits the evaluation of arguments based on equalization of bargaining power.

The incomplete information bargaining framework also has the feature that when firms make noncontractible and nonobservable investments that improve their own type distributions, efficient incomplete information bargaining implies efficient

¹The notion that vertical integration improves outcomes remains influential in antitrust. A case in point is the 2020 update of the US Department of Justice (DOJ) and Federal Trade Commission’s (FTC’s) *Vertical Merger Guidelines*, which after recognizing that “vertical mergers are not invariably innocuous” state that “vertical mergers often benefit consumers through the elimination of double marginalization, which tends to lessen the risks of competitive harm” and that “vertical mergers combine complementary economic functions and eliminate contracting frictions, and therefore have the capacity to create a range of potentially cognizable efficiencies that benefit competition and consumers” (US DOJ and FTC 2020, pp. 2, 11).

²Galbraith (1954, p. 1) saw “the neutralization of one position of power by another” as a mitigant of economic power of “substantial, and perhaps central, importance,” while Stigler (1954, p. 13) lamented the lack of any explanation for “why bilateral oligopoly should in general eliminate, and not merely redistribute, monopoly gains.” The controversy arises in no small part because formalizing notions of countervailing power has proven challenging and because “it is difficult to model bilateral monopoly or oligopoly, and there exists no single canonical model” (Snyder 2008, p. 1188).

³For example, OECD (2011, pp. 50–51) and OECD (2007, pp. 58–59) raise the possible role of collective negotiation and group boycotts for counterbalancing the market power of providers of payment card services. Potential benefits from allowing physician network joint ventures are recognized by the US DOJ and FTC’s 1996 *Statement of Antitrust Enforcement Policy in Health Care*.

equilibrium investments. Thereby, the model sheds new light on ongoing debates in industrial organization and antitrust in the wake of the Dow-DuPont merger decision on the interaction between market structure and investments. It also epitomizes the contrast to complete information models, which with incomplete contracting obtain inefficient investments because of hold-up. With incomplete information, incentive compatibility protects the firms from hold-up, and if bargaining is efficient, it perfectly aligns individuals' investment incentives with the planner's objective.

Our framework uses the Myersonian mechanism design approach (Myerson 1981) to elicit firms' private information and determine prices and builds on the bilateral trade model of Myerson and Satterthwaite (1983), augmented by bargaining weights and multiple buyers and suppliers. Thereby, it combines elements of Myerson and Satterthwaite (1983), Williams (1987), and Gresik and Satterthwaite (1989). Specifically, our model allows for multiple buyers and multiple suppliers without imposing restrictions on the supports of the buyers' values and the suppliers' costs other than assuming that all buyers' value distributions have the same support and all suppliers' cost distributions have the same support. We generalize Williams' (1987) approach of maximizing an objective that assigns differential weights in a bilateral trade problem by allowing for multiple agents. Put differently, our paper reinterprets Myerson and Satterthwaite (1983) as a bilateral monopoly problem, extends it to allow for bargaining weights and multiple agents on both sides of the market, and shows that it is tractable and has all the required features. In particular, inherent to the independent private values setting is the key economic trade-off between rent extraction and social surplus. We defer further discussion of the literature to Section VII.

While our paper does, of course, not resolve the deep problems related to agents' higher-order beliefs and common knowledge assumptions in economics, it seems fair to deflect criticism of incomplete information bargaining models based on these concerns by noting that assuming common knowledge of distributions is weaker than the assumption of complete information models that there is common knowledge of values and costs.

The remainder of the paper is structured as follows. Section I introduces the setup. In Section II, we provide a model of incomplete information bargaining. Section III derives results pertaining to horizontal mergers, and Section IV derives results for vertical integration. Section V analyzes investment incentives. We discuss extensions in Section VI and related literature in Section VII. Section VIII concludes the paper. Formal mechanism design results, longer proofs, and additional results and extensions are relegated to the online Appendix.

I. Setup

We consider a pre-merger market with n^S suppliers and n^B buyers, denoting the sets of suppliers and buyers, respectively, by $\mathcal{N}^S \equiv \{1, \dots, n^S\}$ and $\mathcal{N}^B \equiv \{1, \dots, n^B\}$. Each supplier j has the capacity to produce k_j^S units of a good at a constant marginal cost, and each buyer i has constant marginal value for up to k_i^B units of the good, where k_j^S and k_i^B are positive integers. Total demand is $K^B \equiv \sum_{i \in \mathcal{N}^B} k_i^B$, and total supply is $K^S \equiv \sum_{j \in \mathcal{N}^S} k_j^S$, and we define $K \equiv \min\{K^B, K^S\}$.

Supplier j draws its constant marginal cost c_j independently from distribution G_j with support $[\underline{c}, \bar{c}]$ and density g_j that is positive on the interior of the support. Buyer i draws its constant marginal value v_i independently from distribution F_i with support $[\underline{v}, \bar{v}]$ and density f_i that is positive on the interior of the support. The problem is trivial if $\bar{v} \leq \underline{c}$ because then it is never ex post efficient to have any trade. Therefore, we assume that $\bar{v} > \underline{c}$. We assume that $G_1, \dots, G_n^S, F_1, \dots, F_n^B, k_1^S, \dots, k_n^S$, and k_1^B, \dots, k_n^B are common knowledge, while the realized costs and values are the private information of the individual suppliers and buyers. To save on notation, we ignore ties among the agents' costs and values. While we adhere to a setup with constant marginal costs and values, with additional structure, one can allow for decreasing marginal values and increasing marginal costs.

The suppliers and buyers have quasi-linear preferences. The payoff of supplier j with type c_j when producing $q \in \{0, \dots, k_j^S\}$ units of the good and receiving the monetary transfer m is $m - c_j q$. The payoff of buyer i with type v_i when receiving $q \in \{0, \dots, k_i^B\}$ units of the good and making the monetary payment m is $v_i q - m$. For every agent, we normalize the value of the outside option of not trading to 0.

Because both the buyers' values and the suppliers' costs are random variables whose realizations are the agents' private information, the setup is symmetric with respect to the privacy of information, with the important consequence that ex post efficiency need not be possible.⁴ Indeed, our setup encompasses the classic Myerson-Satterthwaite (1983) setting, where, as they show, for $n^S = n^B = 1$, ex post efficient trade is possible if and only if $\bar{c} \leq \underline{v}$ (see online Appendix A). We refer to the case with $\bar{c} \leq \underline{v}$ as the case of *nonoverlapping supports* and to the case with $\bar{c} > \underline{v}$ as the case of *overlapping supports*. Thus, with one supplier and one buyer, incomplete information prevents ex post efficient trade in the case of overlapping supports, but not in the case of nonoverlapping supports.

We denote supplier j 's *virtual cost* and buyer i 's *virtual value* function by

$$\Gamma_j(c) \equiv c + \frac{G_j(c)}{g_j(c)} \quad \text{and} \quad \Phi_i(v) \equiv v - \frac{1 - F_i(v)}{f_i(v)},$$

which we assume to be increasing (this can be relaxed through the use of “ironing”), and for $a \in [0, 1]$, we define the a -weighted virtual cost functions and the a -weighted virtual value functions by

$$\Gamma_j^a(c) \equiv c + (1 - a) \frac{G_j(c)}{g_j(c)} \quad \text{and} \quad \Phi_i^a(v) \equiv v - (1 - a) \frac{1 - F_i(v)}{f_i(v)}.$$

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. The monotonicity of $\Gamma_j(c)$ and $\Phi_i(v)$ implies that $\Gamma_j^a(c)$ and $\Phi_i^a(v)$ and are also monotone.⁵

⁴To avoid informed-principal problems, we model the mechanism-design problem as one in which a third party without private information—such as a broker or “the market”—organizes the exchange. Although our setup has properties that are sufficient for the informed-principal problem to have no material consequences (see Mylovanov and Tröger 2014), it seems wise to circumvent the associated technicalities. Of course, by giving all the bargaining power to one agent, we still obtain the optimal mechanism for that agent, just as one would if one assumed that the agent with all the bargaining power organized the exchange.

⁵If $g_j(\underline{c}) = 0$, then define $\Gamma_j(\underline{c}) \equiv \lim_{c \downarrow \underline{c}} \Gamma_j(c)$. If $g_j(\bar{c}) = 0$, then $\Gamma_j(\bar{c}) = \infty$. Likewise, if $f_i(\bar{v}) = 0$, then define $\Phi_i(\bar{v}) \equiv \lim_{v \uparrow \bar{v}} \Phi_i(v)$. If $f_i(\underline{v}) = 0$, then $\Phi_i(\underline{v}) = -\infty$. The notation Γ_j^a and Φ_i^a departs from the more

II. Incomplete Information Bargaining

At the heart of any economic model of exchange with transfers are assumptions that govern the price-formation process. For example, oligopoly models specify a mapping from firms' actions to prices, and models based on Nash bargaining specify a mapping from preferences to trades and transfer payments. Our model stays within this tradition and adds to it by introducing an incomplete information bargaining model that allows for heterogeneous bargaining weights. It has neither the shortcoming of standard oligopoly models that buyers are price takers nor the problem of Nash bargaining that outcomes are efficient by assumption. As noted by Ausubel, Cramton, and Deneckere (2002, p. 1934), the results of Myerson and Satterthwaite (1983) imply that the search for efficiency is "fruitless." For the purpose of exposition, it is useful to think of incomplete information bargaining as what the market does and to contrast it with what society, represented by a planner, would choose, with the planner facing the same constraints as the market—*incentive compatibility, individual rationality and no deficit*—while giving equal weight to all agents.

A. Market Mechanism

We model incomplete information bargaining as a direct mechanism $\langle \mathbf{Q}, \mathbf{M} \rangle$ operated by the *market*, where the allocation rule, $\mathbf{Q} = (\mathbf{Q}^S, \mathbf{Q}^B)$ with $Q_j^S: [\underline{v}, \bar{v}]^{n^B} \times [\underline{c}, \bar{c}]^{n^S} \rightarrow \{0, \dots, k_j^S\}$ and $Q_i^B: [\underline{v}, \bar{v}]^{n^B} \times [\underline{c}, \bar{c}]^{n^S} \rightarrow \{0, \dots, k_i^B\}$, maps the agents' types to the quantities provided by the suppliers and the quantities received by the buyers, and the payment rule, $\mathbf{M} = (\mathbf{M}^S, \mathbf{M}^B)$ with $\mathbf{M}^S: [\underline{v}, \bar{v}]^{n^B} \times [\underline{c}, \bar{c}]^{n^S} \rightarrow \mathbb{R}^{n^S}$ and $\mathbf{M}^B: [\underline{v}, \bar{v}]^{n^B} \times [\underline{c}, \bar{c}]^{n^S} \rightarrow \mathbb{R}^{n^B}$, maps types to the payments to the suppliers and the payments from the buyers.⁶ Feasibility requires that for all type realizations, $\sum_{j \in \mathcal{N}^S} Q_j^S(\mathbf{v}, \mathbf{c}) \geq \sum_{i \in \mathcal{N}^B} Q_i^B(\mathbf{v}, \mathbf{c})$.

The mechanism is required to satisfy *incentive compatibility, individual rationality, and no deficit*. A direct mechanism is *Bayesian incentive compatible* if truthfully reporting to the mechanism constitutes a Bayes Nash equilibrium. It is *interim individually rational* if each agent, for every possible type, is weakly better off in expectation by participating in the mechanism than by walking away and taking the outside option with value zero, where the expectations are taken with respect to the type distributions of the other agents, assuming truthful reporting.⁷ A direct mechanism incurs *no deficit* if the sum of the expected payments from the buyers is greater than or equal to the sum of the expected payments to the suppliers. For formal definitions, see online Appendix A.

standard notation in that the coefficient on the hazard rate term is $1 - a$ rather than a , but because we will be introducing bargaining weights, this modification is useful.

⁶Any model of trade maps agents' types into quantities and payments, regardless of whether the model has complete or incomplete information. However, for complete information models, the dependence on agents' types is often degenerate insofar as each agent has only one (known) type.

⁷In our independent private values setting, any Bayesian incentive compatible and interim individually rational mechanism can be implemented with dominant strategies and ex post individual rationality. By construction, it yields the same interim and hence ex ante expected payoffs and revenue. Thus, while we formally state our assumptions in online Appendix A in terms of Bayesian incentive compatibility and interim individual rationality, one could also use the ex post versions of those constraints. However, under what conditions the no-deficit constraint can be allowed to hold ex post remains an open question.

Fixing a mechanism $\langle \mathbf{Q}, \mathbf{M} \rangle$, supplier j 's and buyer i 's ex post surpluses as a function of the type realizations are

$$U_{j;\mathbf{Q},\mathbf{M}}^S(\mathbf{v}, \mathbf{c}) \equiv M_j^S(\mathbf{v}, \mathbf{c}) - c_j Q_j^S(\mathbf{v}, \mathbf{c}),$$

and

$$U_{i;\mathbf{Q},\mathbf{M}}^B(\mathbf{v}, \mathbf{c}) \equiv v_i Q_i^B(\mathbf{v}, \mathbf{c}) - M_i^B(\mathbf{v}, \mathbf{c}).$$

The budget surplus generated by the mechanism is

$$R_{\mathbf{M}}(\mathbf{v}, \mathbf{c}) \equiv \sum_{i \in \mathcal{N}^B} M_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} M_j^S(\mathbf{v}, \mathbf{c}),$$

and the *welfare* or *social surplus* generated by the mechanism is

$$W_{\mathbf{Q}}(\mathbf{v}, \mathbf{c}) \equiv \sum_{i \in \mathcal{N}^B} v_i Q_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} c_j Q_j^S(\mathbf{v}, \mathbf{c}).$$

To capture bargaining power, we endow the agents with bargaining weights $\mathbf{w} = (\mathbf{w}^S, \mathbf{w}^B)$, where $w_j^S \in [0, 1]$ is supplier j 's bargaining weight and $w_i^B \in [0, 1]$ is buyer i 's bargaining weight. We assume that at least one agent's bargaining weight is positive. We define *weighted welfare* with bargaining weights \mathbf{w} to be

$$(1) \quad W_{\mathbf{Q},\mathbf{M}}^{\mathbf{w}}(\mathbf{v}, \mathbf{c}) \equiv \sum_{i \in \mathcal{N}^B} w_i^B U_{i;\mathbf{Q},\mathbf{M}}^B(\mathbf{v}, \mathbf{c}) + \sum_{j \in \mathcal{N}^S} w_j^S U_{j;\mathbf{Q},\mathbf{M}}^S(\mathbf{v}, \mathbf{c}).$$

Thus, the social surplus $W_{\mathbf{Q}}(\mathbf{v}, \mathbf{c})$ created by the mechanism is equal to weighted welfare with all weights equal to 1, that is, $W_{\mathbf{Q},\mathbf{M}}^1(\mathbf{v}, \mathbf{c})$, plus the budget surplus of the mechanism.

We assume that the market maximizes $E_{\mathbf{v},\mathbf{c}}[W_{\mathbf{Q},\mathbf{M}}^{\mathbf{w}}(\mathbf{v}, \mathbf{c})]$, subject to incentive compatibility, individual rationality, and the *no-deficit* constraint, which requires a non-negative expected budget surplus:

$$(2) \quad E_{\mathbf{v},\mathbf{c}}[R_{\mathbf{M}}(\mathbf{v}, \mathbf{c})] \geq 0.$$

We let \mathcal{M} denote the set of incentive compatible, individually rational, no-deficit mechanisms. The *payoff equivalence theorem* (see, e.g., Myerson 1981, Krishna 2010, Börgers 2015) implies that, given $\langle \mathbf{Q}, \mathbf{M} \rangle \in \mathcal{M}$, the expected payoff of an agent is pinned down by the allocation rule and incentive compatibility up to a constant that is equal to the interim expected payoff of the *worst-off type* for that agent, which by incentive compatibility is \bar{c} for a supplier and \underline{v} for a buyer (see online Appendix A). Thus, we have

$$(3) \quad E_{\mathbf{v},\mathbf{c}}[M_j^S(\mathbf{v}, \mathbf{c})] = E_{\mathbf{v},\mathbf{c}}[\Gamma_j(c_j) Q_j^S(\mathbf{v}, \mathbf{c})] + \hat{u}_j^S(\bar{c})$$

and

$$(4) \quad E_{\mathbf{v},\mathbf{c}}[M_i^B(\mathbf{v}, \mathbf{c})] = E_{\mathbf{v},\mathbf{c}}[\Phi_i(v_i) Q_i^B(\mathbf{v}, \mathbf{c})] - \hat{u}_i^B(\underline{v}),$$

where $\hat{u}_j^S(c) \equiv E_{\mathbf{v}, \mathbf{c}, -j}[U_{j; \mathbf{Q}, \mathbf{M}}^S(\mathbf{v}, \mathbf{c})]$ and $\hat{u}_i^B(v) \equiv E_{\mathbf{v}, -i, \mathbf{c}}[U_{i; \mathbf{Q}, \mathbf{M}}^B(\mathbf{v}, \mathbf{c})]$.

It is possible that multiple agents have the maximum bargaining weight and that incentive compatible implementation of the allocation rule for the market mechanism induces a strict budget surplus in expectation while still satisfying all individual rationality constraints. To accommodate this possibility, a complete specification of the outcome of incomplete information bargaining needs to include a sharing rule that pins down the distribution of the budget surplus among these agents. With that in mind, we assume that there are tiebreaking shares $(\eta^S, \eta^B) \in [0, 1]^{n^S+n^B}$ satisfying $\eta_i^x = 0$ if $w_i^x < \max \mathbf{w}$ and $\sum_{j \in \mathcal{N}^S} \eta_j^S + \sum_{i \in \mathcal{N}^B} \eta_i^B = 1$. The market then selects the mechanism in \mathcal{M} that maximizes expected weighted welfare and that distributes the budget surplus absent fixed payments among the agents according to their tiebreaking shares.

We define an *incomplete information bargaining* mechanism with bargaining weights \mathbf{w} to be a mechanism that, among all mechanisms in \mathcal{M} , maximizes expected weighted welfare, $E_{\mathbf{v}, \mathbf{c}}[W_{\mathbf{Q}, \mathbf{M}}^{\mathbf{w}}(\mathbf{v}, \mathbf{c})]$. Notice that, because we evaluate outcomes according to expected welfare $E_{\mathbf{v}, \mathbf{c}}[W_{\mathbf{Q}}(\mathbf{v}, \mathbf{c})]$, the bargaining weights \mathbf{w} are indeed only bargaining weights; that is, they do not affect how outcomes are evaluated, although they do affect the distribution of social surplus and, as we will see, sometimes the size of social surplus.

An immediate implication of this approach is that with equal bargaining weights, incomplete information bargaining delivers the *second-best* allocation rule, which maximizes expected welfare subject to incentive compatibility, individual rationality, and no deficit. Depending on the specifics, the second-best allocation rule may differ from the *first-best* allocation rule that maximizes welfare ex post without accounting for the no-deficit constraint. The first-best allocation rule is monotone and hence permits incentive compatible implementation, and without the no-deficit constraint, individual rationality is trivial to satisfy. In what follows, it will be useful to denote the first-best allocation for a given realization of types by $\mathbf{Q}^{FB}(\mathbf{v}, \mathbf{c})$. Then, for a given realization of types, first-best welfare is

$$(5) \quad W^{FB}(\mathbf{v}, \mathbf{c}) \equiv \sum_{i \in \mathcal{N}^B} v_i Q_i^{FB, B}(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} c_j Q_j^{FB, S}(\mathbf{v}, \mathbf{c}).$$

Another implication of our approach is that if a group of agents on one side of the market have all the bargaining power, e.g., each agent in the group has a bargaining weight of one while all other agents have a bargaining weight of zero, then the incomplete information bargaining outcome is the perfectly collusive outcome for the agents with all the bargaining weight. Although collusive outcomes are not necessarily inconsistent with having large numbers of agents,⁸ under the view that increasing competition on one side of the market reduces the bargaining power of those agents, one could, for example, fix a set of buyers with positive bargaining weight and assume that when there are n^S suppliers, each supplier has bargaining weight $1/n^S$, in which case we get the “usual” result that supplier power goes to zero as the number of suppliers increases.

⁸Hatfield et al. (2020) show that collusion in syndicated markets may become easier as market concentration falls, and that market entry may facilitate collusion because firms can sustain collusion by refusing to syndicate with any firm that undercuts the collusive price.

B. Allocation Rule for Incomplete Information Bargaining

Letting ρ be the Lagrange multiplier on the no-deficit constraint, the Lagrangian associated with maximizing expected weighted welfare (1) subject to the no-deficit constraint (2) is $E_{\mathbf{v}, \mathbf{c}} [W_{\mathbf{Q}, \mathbf{M}}^{\mathbf{w}}(\mathbf{v}, \mathbf{c}) + \rho R_{\mathbf{M}}(\mathbf{v}, \mathbf{c})]$.⁹ Using (3) and (4), it can be rewritten as

$$(6) \quad E_{\mathbf{v}, \mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} w_i^B \overbrace{(v_i - \Phi_i(v_i)) Q_i^B(\mathbf{v}, \mathbf{c})}^{\text{buyer } i\text{'s surplus}} + \sum_{j \in \mathcal{N}^S} w_j^S \overbrace{(\Gamma_j(c_j) - c_j) Q_j^S(\mathbf{v}, \mathbf{c})}^{\text{supplier } j\text{'s surplus}} + \rho \left(\underbrace{\sum_{i \in \mathcal{N}^B} \Phi_i(v_i) Q_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j(c_j) Q_j^S(\mathbf{v}, \mathbf{c})}_{\text{budget surplus}} \right) \right]$$

plus the term $\sum_{i \in \mathcal{N}^B} (w_i^B - \rho) \hat{u}_i^B(\mathbf{y}) + \sum_{j \in \mathcal{N}^S} (w_j^S - \rho) \hat{u}_j^S(\bar{c})$, which can be treated parametrically.¹⁰ The buyer, supplier, and budget surpluses identified in (6) are the parts of the respective surpluses that vary with the allocation rule.

Given the Lagrange multiplier ρ , the allocation rule that maximizes (6) can be defined pointwise. For the case of one supplier with cost c and one buyer with value v , it is straightforward to show that the optimum has $Q_1^S(v, c) = Q_1^B(v, c) = \min\{k_1^S, k_1^B\}$ if $\Gamma_1^{w_1^S/\rho}(c) \leq \Phi_1^{w_1^B/\rho}(v)$, and $Q_1^S(v, c) = Q_1^B(v, c) = 0$ otherwise. For the general case, this basic rule extends as one might expect, but requires some additional notation.

Let $\Gamma_j^a(c)$ denote the constant vector $(\Gamma_j^a(c), \dots, \Gamma_j^a(c))$ with k_j^S elements and denote by $\mathbf{\Gamma}^a(\mathbf{c}) \equiv (\mathbf{\Gamma}_j^a(c_j))_{j \in \mathcal{N}^S}$ the merged list of these weighted virtual costs. Analogously, let $\Phi_i^a(v)$ denote the constant vector $(\Phi_i^a(v), \dots, \Phi_i^a(v))$ with k_i^B elements and denote by $\mathbf{\Phi}^a(\mathbf{v}) \equiv (\mathbf{\Phi}_i^a(v_i))_{i \in \mathcal{N}^B}$ the merged list of these weighted virtual values. For a given type vector (\mathbf{v}, \mathbf{c}) , bargaining weight vector \mathbf{w} , and Lagrange multiplier ρ , the objective in (6) is maximized when the quantity traded $q^*(\rho)$ is the largest element of $\{0, 1, \dots, K\}$ such that the $q^*(\rho)$ lowest elements of $\mathbf{\Gamma}^{\mathbf{w}/\rho}(\mathbf{c})$ are less than or equal to the $q^*(\rho)$ greatest elements of $\mathbf{\Phi}^{\mathbf{w}/\rho}(\mathbf{v})$. We select, arbitrarily but without loss of generality, the *largest* quantity consistent with the virtual values associated with traded units being greater than or equal to the virtual costs associated with traded units.

Defining $\Gamma^*(\rho)$ to be the $q^*(\rho)$ -th lowest element of $\mathbf{\Gamma}^{\mathbf{w}/\rho}(\mathbf{c})$ and $\Phi^*(\rho)$ to be the $q^*(\rho)$ -th highest element of $\mathbf{\Phi}^{\mathbf{w}/\rho}(\mathbf{v})$, it follows that $\Gamma^*(\rho) \leq \Phi^*(\rho)$ and that $\Gamma^*(\rho)$ and $\Phi^*(\rho)$ are thresholds that separate, on each side of the market, the agents

⁹While we do not pursue it here, our approach generalizes to the requirement that the mechanism needs to generate an expected budget surplus of $\kappa \in \mathbb{R}$, which is not more than the maximum expected budget surplus that any incentive-compatible, individually rational mechanism can generate.

¹⁰Denoting by $\hat{q}_i^z(z)$ and $\hat{m}_i^z(z)$ the interim expected quantity and payment of agent i when its type is z for $x \in \{B, S\}$, which are formally defined in online Appendix A, we have $\hat{u}_i^B(\mathbf{y}) = \hat{q}_i^B(\mathbf{y})\mathbf{y} - \hat{m}_i^B(\mathbf{y})$ and $\hat{u}_j^S(\bar{c}) = \hat{m}_j^S(\bar{c}) - \hat{q}_j^S(\bar{c})\bar{c}$. Consequently, no matter what the pointwise maximizer implies for $\hat{q}_i^B(\mathbf{y})$ and $\hat{q}_j^S(\bar{c})$, one can achieve any value for $\hat{u}_i^B(\mathbf{y})$ and $\hat{u}_j^S(\bar{c})$ by appropriately varying $\hat{m}_i^B(\mathbf{y})$ and $\hat{m}_j^S(\bar{c})$, respectively.

that trade from those that do not. We denote the set of *inframarginal* suppliers and buyers, respectively, by

$$\mathcal{N}_I^S(\rho) \equiv \{j \in \mathcal{N}^S \mid \Gamma_j^{w_j^S/\rho}(c_j) < \Gamma^*(\rho)\}$$

and

$$\mathcal{N}_I^B(\rho) \equiv \{i \in \mathcal{N}^B \mid \Phi_i^{w_i^B/\rho}(v_i) > \Phi^*(\rho)\}.$$

Observe that q^* , Γ^* , Φ^* , \mathcal{N}_I^S , and \mathcal{N}_I^B also depend on \mathbf{v} , \mathbf{c} , and \mathbf{w} , but to ease notation we do not make this dependence explicit.

With this in hand, we are in a position to describe the allocation rule that maximizes (6) pointwise for a given ρ . That allocation rule induces each supplier $j \in \mathcal{N}_I^S(\rho)$ to produce k_j^S and each buyer $i \in \mathcal{N}_I^B(\rho)$ to obtain k_i^B units. Ignoring ties among the weighted virtual types of different agents at these threshold values, which occur with probability zero, the “residual” quantity $q^*(\rho) - \sum_{j \in \mathcal{N}_I^S(\rho)} k_j^S$ is procured from the supplier whose weighted virtual cost is equal to $\Gamma^*(\rho)$, and the quantity $q^*(\rho) - \sum_{i \in \mathcal{N}_I^B(\rho)} k_i^B$ is allocated to the buyer whose weighted virtual value is equal to $\Phi^*(\rho)$. We refer to this allocation rule as the *pointwise maximizer* given ρ .

It only remains to specify the solution value ρ^w for the Lagrange multiplier ρ . Following the same arguments that were first developed in the working paper version of Gresik and Satterthwaite (1989) and that were first used in published form in Myerson and Satterthwaite (1983), ρ^w is the smallest feasible value for ρ such that the no-deficit constraint is satisfied by the pointwise maximizer given ρ . Because any budget surplus can be reallocated to the agents through fixed payments, and because ρ^w is the shadow price of the no-deficit constraint, we have $\rho^w \geq \max \mathbf{w}$. In addition, because a positive expected budget surplus is always possible given our assumption that $\bar{v} > \underline{c}$, the shadow price is finite. Thus, we can define ρ^w as the smallest value of ρ greater than or equal to $\max \mathbf{w}$ such that the “budget surplus” term in (6) is nonnegative for the pointwise maximizer given ρ , and we have the following result:

LEMMA 1: *The allocation rule for incomplete information bargaining with bargaining weights \mathbf{w} , \mathbf{Q}^w , is defined by*

$$Q_j^{w,S}(\mathbf{v}, \mathbf{c}) \equiv \begin{cases} k_j^S & \text{if } \Gamma_j^{w_j^S/\rho^w}(c_j) < \Gamma^*(\rho^w), \\ q^*(\rho^w) - \sum_{\ell \in \mathcal{N}_I^S(\rho^w)} k_\ell^S & \text{if } \Gamma_j^{w_j^S/\rho^w}(c_j) = \Gamma^*(\rho^w), \end{cases}$$

and $Q_j^{w,S}(\mathbf{v}, \mathbf{c}) \equiv 0$ otherwise, and

$$Q_i^{w,B}(\mathbf{v}, \mathbf{c}) \equiv \begin{cases} k_i^B & \text{if } \Phi_i^{w_i^B/\rho^w}(v_i) > \Phi^*(\rho^w), \\ q^*(\rho^w) - \sum_{\ell \in \mathcal{N}_I^B(\rho^w)} k_\ell^B & \text{if } \Phi_i^{w_i^B/\rho^w}(v_i) = \Phi^*(\rho^w), \end{cases}$$

and $Q_i^{w,B}(\mathbf{v}, \mathbf{c}) \equiv 0$ otherwise.

PROOF:

See online Appendix B.

Because weighted virtual values are weakly less than the associated values, and weighted virtual costs are weakly greater than the associated costs, an immediate implication of Lemma 1 is that the total quantity traded in incomplete information bargaining never exceeds the (largest) quantity traded under the first best.

C. *Payoffs under Incomplete Information Bargaining*

We are left to augment the allocation rule of Lemma 1 with a consistent payment rule. As described above, based on the payoff equivalence theorem, all that remains to be done is to define the fixed payments to the agents’ worst-off types. Individual rationality is satisfied if and only if the fixed payment to each supplier is nonnegative and the fixed payment from each buyer is nonpositive. The optimization of the weighted objective requires that no money be left on the table. So, we first define the “money on the table” before fixed payments are made, i.e., the budget surplus under the mechanism of Lemma 1, not including the fixed payments, given by

$$(7) \quad \pi^w \equiv E_{\mathbf{v}, \mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} \Phi_i(v_i) Q_i^{w,B}(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j(c_j) Q_j^{w,S}(\mathbf{v}, \mathbf{c}) \right].$$

Because all expected budget surplus is distributed to the agents, it follows that

$$\pi^w = \sum_{j \in \mathcal{N}^S} \hat{u}_j^S(\bar{c}) + \sum_{i \in \mathcal{N}^B} \hat{u}_i^B(\underline{v}).$$

Of course, if $\rho^w > \max \mathbf{w}$, then the no-deficit constraint binds, implying that $\pi^w = 0$ and that the question of how to allocate the budget surplus is moot. In contrast, if $\rho^w = \max \mathbf{w}$, then $\pi^w \geq 0$. In this case, weighted welfare is maximized when π^w is allocated among the suppliers and buyers with bargaining weights equal to $\max \mathbf{w}$, which is accomplished by having interim expected payoffs to the agents’ worst-off types of

$$(8) \quad \hat{u}_j^S(\bar{c}; \mathbf{w}, \boldsymbol{\eta}) = \eta_j^S \pi^w \text{ and } \hat{u}_i^B(\underline{v}; \mathbf{w}, \boldsymbol{\eta}) = \eta_i^B \pi^w,$$

where, as defined above, $\eta_j^S = 0$ and $\eta_i^B = 0$ for any supplier j and buyer i that does not have the maximum bargaining weight.

The outcome of incomplete information bargaining with bargaining weights \mathbf{w} and tiebreaking shares $\boldsymbol{\eta}$ is then given by the expected buyer and supplier payoffs implied by the allocation rule \mathbf{Q}^w given in Lemma 1 and interim expected payoffs to agents’ worst-off types given by (8). Thus, we have:

PROPOSITION 1: *Incomplete information bargaining with bargaining weights \mathbf{w} and shares $\boldsymbol{\eta}$ generates expected supplier payoffs for $j \in \mathcal{N}^S$ of*

$$u_j^S(\mathbf{w}, \boldsymbol{\eta}) \equiv \eta_j^S \pi^w + E_{\mathbf{v}, \mathbf{c}} \left[(\Gamma_j(c_j) - c_j) Q_j^{w,S}(\mathbf{v}, \mathbf{c}) \right],$$

and expected buyer payoffs for $i \in \mathcal{N}^B$ of

$$u_i^B(\mathbf{w}, \boldsymbol{\eta}) \equiv \eta_i^B \pi^w + E_{\mathbf{v}, \mathbf{c}} \left[(v_i - \Phi_i(v_i)) Q_i^{w,B}(\mathbf{v}, \mathbf{c}) \right].$$

The outcomes from incomplete information bargaining given in Proposition 1 coincide with the set of Pareto undominated payoffs associated with mechanisms in \mathcal{M} . To see this, first note that because no money is left on the table, any expected payoffs from incomplete information bargaining are Pareto undominated among payoffs resulting from mechanisms in \mathcal{M} . Conversely, given a vector of expected payoffs $\tilde{\mathbf{u}}$ that is the outcome of $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle \in \mathcal{M}$ and that is Pareto undominated in the set of expected payoff vectors that obtain from mechanisms in \mathcal{M} , the weights \mathbf{w} and shares $\boldsymbol{\eta}$ that induce $\tilde{\mathbf{u}}$ follow from the dual characterization of maximal elements (see, for example, Boyd and Vandenberghe 2004, and online Appendix B for a derivation).

PROPOSITION 2: *If expected payoff vector $\tilde{\mathbf{u}}$ associated with $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle \in \mathcal{M}$ is Pareto undominated among expected payoff vectors for mechanisms in \mathcal{M} , then there exist bargaining weights \mathbf{w} and shares $\boldsymbol{\eta}$ such that $\mathbf{Q}^{\mathbf{w}} = \tilde{\mathbf{Q}}$ and $(\mathbf{u}^S(\mathbf{w}, \boldsymbol{\eta}), \mathbf{u}^B(\mathbf{w}, \boldsymbol{\eta})) = \tilde{\mathbf{u}}$. Conversely, given bargaining weights \mathbf{w} and shares $\boldsymbol{\eta}$, $(\mathbf{u}^S(\mathbf{w}, \boldsymbol{\eta}), \mathbf{u}^B(\mathbf{w}, \boldsymbol{\eta}))$ is Pareto undominated among expected payoff vectors for mechanisms in \mathcal{M} .*

PROOF:

See online Appendix B.

As we show in online Appendix D, incomplete information bargaining includes the k -double auction of Chatterjee and Samuelson (1983) as a special case (when $n^S = n^B = 1$ and agents draw their types from the same uniform distribution). In incomplete information bargaining, just as in the k -double auction, equalization of bargaining power increases expected social surplus, which is what we turn to next.

D. Social-Surplus-Increasing Equalization of Bargaining Weights

Despite the result of Proposition 2 that incomplete information bargaining is Pareto efficient, its outcome may differ from what the planner would choose. This creates potential for social-surplus-increasing *equalization of bargaining power*—by which we mean changing some asymmetric vector of bargaining weights $\tilde{\mathbf{w}}$ to $\mathbf{w} = (w, \dots, w)$ —and the possibility that the negative consequences of, say, a merger on social surplus might be reversed by an associated equalization of bargaining power.

In particular, denoting by \mathbf{Q}^* and $\mathbf{Q}^{\mathbf{w}}$ the allocation rule chosen by the planner and the market, and by $W^* \equiv E_{\mathbf{v}, \mathbf{c}}[W_{\mathbf{Q}^*}(\mathbf{v}, \mathbf{c})]$ and $W^{\mathbf{w}} \equiv E_{\mathbf{v}, \mathbf{c}}[W_{\mathbf{Q}^{\mathbf{w}}}(\mathbf{v}, \mathbf{c})]$ the value of the planner's objective under \mathbf{Q}^* and $\mathbf{Q}^{\mathbf{w}}$, respectively, we have $W^{\mathbf{w}} \leq W^*$ because the allocation rule $\mathbf{Q}^{\mathbf{w}}$ is available when the planner chooses \mathbf{Q}^* . Notice also that $\mathbf{Q}^* = \mathbf{Q}^{(w, \dots, w)}$ for any $w \in (0, 1]$.¹¹ Hence, for any $w \in (0, 1]$, we have $W^* = W^{(w, \dots, w)}$.

¹¹To see this, note that $W_{\mathbf{Q}^{\mathbf{w}}}^{(w, \dots, w)}(\mathbf{v}, \mathbf{c}) = w(W_{\mathbf{Q}^*}(\mathbf{v}, \mathbf{c}) - R_{\mathbf{M}}(\mathbf{v}, \mathbf{c}))$, which is maximized, subject to no deficit, at \mathbf{Q}^* . With symmetric bargaining weights, the weight w has a multiplicative effect on the solution value of the Lagrange multiplier on the no-deficit constraint, but ultimately no effect on the allocation rule $\mathbf{Q}^{\mathbf{w}}$, which depends on w divided by the Lagrange multiplier.

Given a market with weights \mathbf{w} , we say that the *planner prefers an equalization of bargaining weights* if $W^{\mathbf{w}} < W^*$, or equivalently, $\mathbf{Q}^{\mathbf{w}}(\mathbf{v}, \mathbf{c}) \neq \mathbf{Q}^*(\mathbf{v}, \mathbf{c})$ for all (\mathbf{v}, \mathbf{c}) in an open subset of $[\underline{v}, \bar{v}]^{n^B} \times [\underline{c}, \bar{c}]^{n^S}$. As stated in the next proposition, specific conditions are required for the planner *not* to prefer an equalization of bargaining weights. Of course, the question of equalization of bargaining weights is moot if these weights are already all the same. But even when the weights differ, there may be no benefit to the planner if the market has full trade, that is,

$$(9) \quad \left(\text{K-th lowest of } \left\{ \Gamma_j^{w_j^S / \rho^{\mathbf{w}}}(\bar{c}) \right\}_{j \in \mathcal{N}^S} \right) \leq \left(\text{K-th highest of } \left\{ \Phi_i^{w_i^B / \rho^{\mathbf{w}}}(\underline{v}) \right\}_{i \in \mathcal{N}^B} \right),$$

which implies that $\rho^{\mathbf{w}} = \max \mathbf{w}$, and if there is sufficient symmetry among the agents that it is always the highest-value buyers and lowest-cost suppliers that trade. Specifically, the suppliers must have equal bargaining weights and the buyers must have equal bargaining weights, and if one side of the market has a lower bargaining weight, i.e., does not have weight equal to $\max \mathbf{w}$, then agents on that side must have symmetric distributions so that the ordering of virtual types matches the ordering of actual types.¹²

PROPOSITION 3: *In a market with asymmetric bargaining weights \mathbf{w} , the planner strictly prefers an equalization of bargaining weights unless all of the following conditions are satisfied:*

- (i) *The full-trade condition (9) holds;*
- (ii) *For all $j \in \mathcal{N}^S$, $w_j^S = w^S$, and for all $i \in \mathcal{N}^B$, $w_i^B = w^B$;*
- (iii) *If $w^S < w^B$, then for all $j \in \mathcal{N}^S$, $G_j = G$;*
- (iv) *If $w^B < w^S$, then for all $i \in \mathcal{N}^B$, $F_i = F$.*

PROOF:

See online Appendix B.

Proposition 3 provides conditions on bargaining weights and primitives such that the planner does not benefit from an equalization of bargaining weights. (As we show in Section VI, the planner may prefer unequal bargaining weights that favor buyers in the presence of downstream consumers.) However, unless $n^B = n^S = 1$, $\bar{c} \leq \Phi_1(\underline{v})$, and $\underline{v} \geq \Gamma_1(\bar{c})$,¹³ there always exist asymmetric bargaining weights \mathbf{w} such that the planner benefits from an equalization of bargaining weights, i.e., $W^{\mathbf{w}} < W^*$.

¹²That ex ante symmetry among agents implies that there is no inefficiency in production when the production decision is based on (equally weighted) virtual types rather than actual types hinges on the assumption that the virtual type functions are increasing. Without that assumption, the weighted virtual type functions would have to be replaced by their “ironed” counterparts (see Myerson 1981), and the resulting allocation rules would induce inefficiency with positive probability because of randomness due to tiebreaking.

¹³These distributional assumptions are restrictive in the sense that they are not satisfied if the supports of the buyer’s and supplier’s type distributions overlap because $\Phi_1(v) < v$ for any $v < \bar{v}$ and $\Gamma_1(c) > c$ for any $c > \underline{c}$. Further, the conditions fail in many cases even when there is no overlap—for example, if the supplier draws its cost from the uniform distribution on $[0, 1]$ and the buyer draws its value from the uniform distribution on $[\underline{v}, \underline{v} + 1]$, then

This shows that, quite generally, equalization of bargaining power increases social surplus. Some of the benefits that the planner obtains from more equal bargaining weights stem from an equalization of bargaining weights among agents on the same side of the market, which eliminates socially wasteful discrimination among the agents based on differently weighted virtual types. While this effect is integral to the incomplete information bargaining model that we study here, equalization of bargaining power on one side of the market is arguably *not* what competition authorities and practitioners, or for that matter, John Galbraith, have in mind when speaking of countervailing power, which refers to an equalization of bargaining power across the two sides of the market.

In light of this, we consider the frontier of total supplier and total buyer expected payoffs. Let $\eta_{\mathbf{w}}$ denote the tiebreaking shares that specify an equal division of any budget surplus among the agents with the maximum bargaining weight in \mathbf{w} ,¹⁴ and define the minimum and maximum total supplier payoffs by

$$\underline{u}_S \equiv \min_{\mathbf{w}} \sum_{j \in \mathcal{N}^S} u_j^S(\mathbf{w}, \eta_{\mathbf{w}}) \text{ and } \bar{u}_S = \max_{\mathbf{w}} \sum_{j \in \mathcal{N}^S} u_j^S(\mathbf{w}, \eta_{\mathbf{w}}).$$

For $u \in [\underline{u}_S, \bar{u}_S]$, let

$$\omega(u) \equiv \max_{\mathbf{w}} \sum_{i \in \mathcal{N}^B} u_i^B(\mathbf{w}, \eta_{\mathbf{w}}) \text{ subject to } \sum_{j \in \mathcal{N}^S} u_j^S(\mathbf{w}, \eta_{\mathbf{w}}) \geq u.$$

In honor of Williams (1987), who first analyzed problems of this kind in a bilateral trade setting, we call

$$\mathcal{F} \equiv \{(u, \omega(u)) \mid u \in [\underline{u}_S, \bar{u}_S]\}$$

the *Williams frontier*. This contrasts with the *first-best frontier*, which is

$$\mathcal{F}^{FB} \equiv \{(u_S, u_B) \in \mathbb{R}_+^2 \mid u_S + u_B = E_{\mathbf{v}, \mathbf{c}}[W^{FB}(\mathbf{v}, \mathbf{c})]\}.$$

The following lemma shows that the Williams frontier is defined by payoffs associated with bargaining weights that are symmetric on each side of the market. Defining, for $\Delta \in [0, 1]$, \mathbf{w}_Δ by $\mathbf{w}_\Delta^S \equiv (1 - \Delta, \dots, 1 - \Delta)$ and $\mathbf{w}_\Delta^B \equiv (\Delta, \dots, \Delta)$, and letting $\tilde{u}_j^S(\Delta) \equiv u_j^S(\mathbf{w}_\Delta, \eta_{\mathbf{w}_\Delta})$ and $\tilde{u}_i^B(\Delta) \equiv u_i^B(\mathbf{w}_\Delta, \eta_{\mathbf{w}_\Delta})$, we have:

LEMMA 2: $\mathcal{F} = \{(\sum_{j \in \mathcal{N}^S} \tilde{u}_j^S(\Delta), \sum_{i \in \mathcal{N}^B} \tilde{u}_i^B(\Delta)) \mid \Delta \in [0, 1]\}$.

PROOF:

See online Appendix B.

Using Lemma 2, we have the following characterization of the Williams frontier:

the conditions hold if and only if $\nu \geq 2$. If $\bar{c} > \Phi_1(\nu)$, then giving the supplier all the bargaining power reduces welfare below W^* , and if $\nu < \Gamma_1(\bar{c})$, then giving the buyer all the bargaining power reduces welfare below W^* .

¹⁴That is, for $x \in \{B, S\}$, if $w_i^x \neq \max \mathbf{w}$, then $\eta_{\mathbf{w}, i}^x \equiv 0$, and otherwise $\eta_{\mathbf{w}, i}^x(\mathbf{w}) \equiv 1/m$, where m is the number of elements of \mathbf{w} that are equal to $\max \mathbf{w}$.

PROPOSITION 4: *The Williams frontier is concave to the origin and the frontier is strictly concave if and only if its intersection with the first-best frontier contains at most one point.*

PROOF:

See online Appendix B.

As shown in Proposition 4, the Williams frontier is strictly concave if and only if it coincides with the first-best frontier at most at one point. This occurs, for example if the supports of the suppliers' and buyers' type distributions coincide, that is, if $\underline{v} = \underline{c}$ and $\bar{v} = \bar{c}$, because in that case the first best is not possible.¹⁵

If there is a gap between the supports, that is, $\underline{v} > \bar{c}$, then the Williams frontier follows the first-best frontier for a range of bargaining weights that are sufficiently symmetric. Along that segment, it is only weakly concave. This is illustrated in Figure 1. As shown in panel A, for the case of overlapping supports, the first best cannot be achieved and the Williams frontier is strictly concave. In contrast, as shown in panel B, with nonoverlapping supports, the first best is achieved for a range of Δ close to $1/2$ and in that range the frontier is linear. The reasons for the concavity of the Williams frontier are essentially the same as those invoked by Paul Samuelson to show that with constant returns to scale the production possibility frontier is concave: the convex combination between any two points on the frontier can be achieved by randomizing over the mechanisms associated with each of them. By reoptimizing, one may be able to do better. The linear segment of the frontier in Figure 1, panel B is a case where reoptimizing cannot improve outcomes because at both endpoints of that segment, the mechanism is already first best.

Building on Proposition 4 and letting $W(\Delta) \equiv \sum_{j \in \mathcal{N}^S} \tilde{u}_j^S(\Delta) + \sum_{i \in \mathcal{N}^B} \tilde{u}_i^B(\Delta)$, the concavity of the Williams frontier has the following implication:

COROLLARY 1: *Movement toward the equalization of buyer-side and supplier-side bargaining weights along the Williams frontier weakly increases social surplus, i.e., if $\Delta' < \Delta \leq 1/2$ or $1/2 \leq \Delta < \Delta'$, then $W(\Delta') \leq W(\Delta) \leq W(1/2)$.*

Corollary 1 has policy implications for settings in which competition authorities could allow actions that equalize bargaining power. For example, allowing merchants that purchase payment card services from powerful suppliers (e.g., Visa, Mastercard) to engage in group boycotts might improve social surplus by equalizing bargaining power (see footnote 3). Considering bargaining between powerful

¹⁵With single-unit traders and identical distributions on each side of the market, this follows from Williams (1999). Away from identical distributions and single-unit traders, the impossibility of the first best follows from the facts that with identical supports the lowest buyer and highest seller types never trade and that with multiunit traders the deficit of the Vickrey-Clarke-Groves (VCG) mechanism is bounded below by the Walrasian price gap times the quantity traded (Loertscher and Mezzetti 2019). By the payoff equivalence theorem, no efficient mechanism that satisfies incentive compatibility and individual rationality runs a smaller deficit in expectation than the VCG mechanism. As these arguments make clear, identical supports are only sufficient since $\underline{v} \leq \underline{c}$ and $\bar{v} \leq \bar{c}$ also implies that the least efficient types on each side never trade.

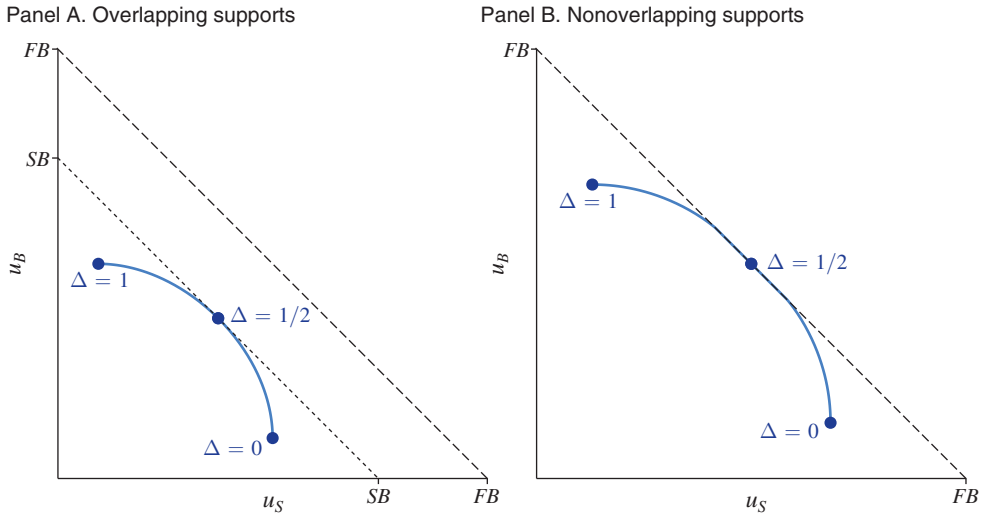


FIGURE 1. WILLIAMS FRONTIER \mathcal{F} FOR THE CASE OF ONE SINGLE-UNIT BUYER AND ONE SINGLE-UNIT SUPPLIER

Notes: FB denotes the first-best social surplus, and SB denotes the second-best social surplus. All types are uniformly distributed. Panel A assumes that $[\underline{c}, \bar{c}] = [\underline{v}, \bar{v}] = [0, 1]$. Panel B assumes that $[\underline{c}, \bar{c}] = [0, 1]$ and $[\underline{v}, \bar{v}] = [1.1, 2.1]$, in which case first-best and second-best total surplus are the same.

insurance companies and doctors for the supply of health services, allowing physician joint ventures might equalize bargaining power and improve social surplus.¹⁶

III. Horizontal Mergers

In this section, we analyze horizontal mergers, including both the effects of a merger on the merging parties' type distribution and the possibility of merger effects on the bargaining power of both the merging and nonmerging parties. We evaluate outcomes from an ex ante perspective, that is, before firms' types are realized.

A. Modeling Horizontal Mergers

To model mergers within our constant-returns-to-scale setup, we assume that the merged entity draws its constant marginal type from a distribution that combines the distributions of the merging firms. Further, we assume that the capacity of the merged entity combines the capacities of the merging firms. For a merger of suppliers i and j , we denote the merged entity's cost distribution by $G_{i,j}$ and its capacity by $k_{i,j}^S$, and for a merger of buyers i and j , we denote the merged entity's value

¹⁶For example, the US DOJ and FTC (1996) state: "The Agencies will not challenge, absent extraordinary circumstances, a nonexclusive physician network joint venture whose physician participants share substantial financial risk and constitute 30 percent or less of the physicians in each physician specialty with active hospital staff privileges who practice in the relevant geographic market" (p. 65). While the guidelines do not allow simple group negotiations over price, without some form of financial integration, under the theory that group negotiation would tend to raise prices, Corollary 1 finds benefits to group negotiation that equalizes bargaining weights even in the absence of financial integration.

distribution by $F_{i,j}$ and its capacity by $k_{i,j}^B$. We assume that the merged entity's virtual type function is increasing.

To model a merger, one needs to describe how a merger transforms the two pre-merger firms' distributions and capacities into the distribution and capacity of the merged entity. The natural mapping from pre-merger to postmerger firms is clear for a merger of firms whose capacities are sufficiently large that each could individually serve the entire other side of the market. For example, suppose that suppliers 1 and 2 merge, where $k_1^S = k_2^S = K^B$. In this case, we model the merged entity as having a constant marginal cost for K^B units that is drawn from the distribution of the minimum of a cost drawn from G_1 and a cost drawn from G_2 , i.e., $G_{1,2}(c) = 1 - (1 - G_1(c))(1 - G_2(c))$. This has the natural interpretation of a merged entity that has two facilities, each with constant marginal cost for K^B units, where the merged entity rationalizes its production by using only the facility with the lower marginal cost. In other words, we assume that there are no synergies associated with a merger beyond the ability to rationalize production or consumption between the component firms. Analogously, the merged entity created from the merger of buyers 1 and 2 with $k_1^B = k_2^B = K^S$ would draw its constant marginal value for K^S units from $F_{1,2}(v) = F_1(v)F_2(v)$.

B. Mergers That Do Not Affect the Bargaining Weights

We say that a merger *does not alter bargaining weights or shares* if (i) all nonmerging firms retain their pre-merger bargaining weights and shares in the postmerger market; and (ii) the two merging firms have the same bargaining weight in the pre-merger market, which is then inherited by the merged entity, that is, $w_i = w_j = w_{i,j}$ for a merger between i and j , and the share of the merged entity is equal to the sum of the shares of the merging firms, that is, $\eta_{i,j} = \eta_i + \eta_j$. We say that supplier j has *maximum capacity* if $k_j^S = K^B$ and that buyer i has *maximum capacity* if $k_i^B = K^S$.

A revealed-preference type of argument, which we sketch next, allows us to analyze the effects of mergers that do not alter bargaining weights or shares. Taking the case of merging maximum-capacity suppliers i and j , the postmerger incomplete information bargaining mechanism can be "applied" to the pre-merger market by taking $c = \min\{c_i, c_j\}$ to be the type of the merged entity and assigning the merged entity's quantity to the lower cost of the two merging suppliers. By the payoff equivalence theorem, we can, without loss, focus on payments that are the sum of a firm's "threshold types" for each unit traded, where the threshold type for a unit is the worst type that the firm could report and still trade that unit. Because nonmerging firms' threshold types depend only on the minimum of c_i and c_j , these threshold types are not affected by the merger. The expected payments to the merging suppliers strictly decrease, generating a budget surplus to be paid out to firms with the maximum bargaining weight, which increases weighted welfare if it transfers surplus to higher-weighted nonmerging firms. Any optimization of the pre-merger mechanism further benefits pre-merger expected weighted welfare. If the merging firms have all the bargaining weight, then the transfer has no effect and no further optimization of the mechanism is possible, implying that no firm is affected. The proof in online Appendix B applies this line of reasoning and gives us the following result:

PROPOSITION 5: *A horizontal merger of maximum-capacity firms with common bargaining weight w that does not alter bargaining weights or shares (i) weakly reduces expected weighted welfare (strictly if $w < \max \mathbf{w}$), and (ii) is neutral for social surplus and all firms if the merging firms have all the bargaining power.*

PROOF:

See online Appendix B.

The second part of Proposition 5 provides conditions under which a merger is neutral. Using the first part of the proposition, we have:

COROLLARY 2: *A horizontal merger of firms with maximum capacities that does not alter bargaining weights or shares (i) harms any nonmerging firm that has all the bargaining power, and (ii) weakly reduces expected social surplus when all firms have the same bargaining weight.*

The first part of the corollary follows from Proposition 5(i) by noting that when a nonmerging firm has all the bargaining power, then that firm's surplus is equal to weighted welfare. The second part follows similarly noting that when all firms have the same bargaining weight, then social surplus is equal to weighted welfare.

Proposition 5 generalizes the insights from Loertscher and Marx (2019) that a merger harms a powerful buyer to a setting in which incomplete information pertains to both sides of the market, there are multiple buyers and suppliers with multiunit demand and supply, and bargaining power is not restricted to be with the buyer. Proposition 5 implies that two maximum-capacity firms on the same side of the market that are the only firms with bargaining power have no incentive to merge, which is an effect that is depicted in Figure 2 below. In addition, as in Loertscher and Marx (2019), a merger need not be profitable for the merging suppliers—when the buyers have sufficient bargaining power, the resulting more aggressive behavior against the merged entity in response to its stronger type distribution can outweigh the benefits to the merging suppliers from the elimination of competition.

C. Bargaining Power Effects of Mergers

In addition to changing the type distribution of the merged entity compared to the merging firms, it is also conceivable that mergers alter firms' bargaining powers. Indeed, the idea that a merger somehow "levels the playing field" in terms of bargaining power is based on this very conception.¹⁷ It finds support in the empirical literature (Ho and Lee 2017, Bhattacharyya and Nain 2011, Decarolis and Rovigatti 2021) and features prominently in antitrust debates and cases. Nonetheless, a major obstacle to analyzing the effects of the equalization of bargaining power in existing modeling approaches is that these typically either assume efficient bargaining,

¹⁷As a case in point, the Australian government's 1999 (now superseded) guidelines stated: "If pre-merger prices are distorted from competitive levels by market power on the opposite side 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).

Panel A. Social-surplus-enhancing equalization of bargaining weights

Panel B. First-best-achieving equalization of bargaining weights

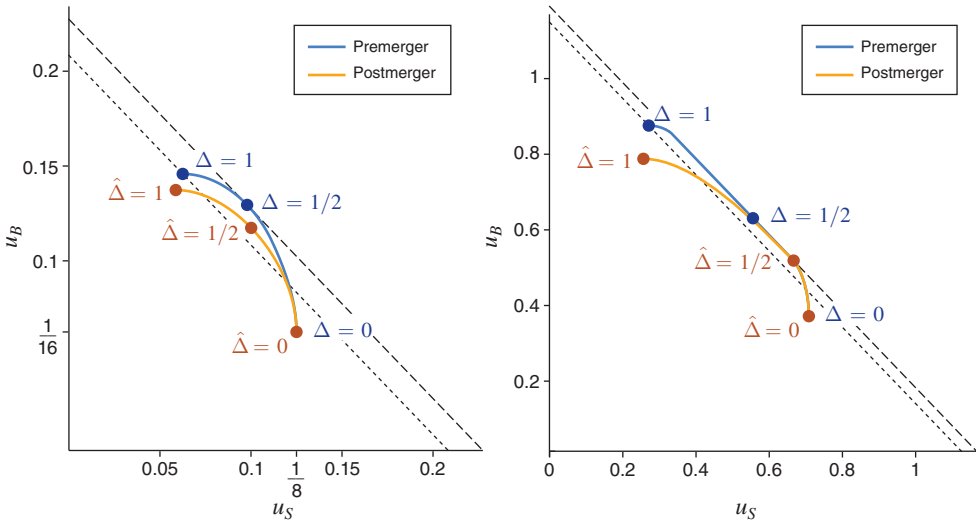


FIGURE 2. WILLIAMS FRONTIER PRE- AND POSTMERGER FOR 1 BUYER AND 2 PRE-MERGER SUPPLIERS

Notes: Williams frontier for the case of one pre-merger buyer and two symmetric pre-merger suppliers (blue), all with capacity of one, and the frontier following the merger of the two suppliers (orange). The pre-merger frontier is defined by $(w^S, w^B) = (1 - \Delta, 1 - \Delta)$ with $\Delta \in [0, 1]$, and the postmerger frontier is defined by $(\hat{w}^S, \hat{w}^B) = (1 - \hat{\Delta}, \hat{\Delta})$ with $\hat{\Delta} \in [0, 1]$. We assume equal tiebreaking shares among the firms with the maximum bargaining weight. Panel A assumes that the buyer's and pre-merger suppliers' types are uniformly distributed on $[0, 1]$. Panel B assumes that the pre-merger suppliers' costs are uniformly distributed on $[0, 1]$ and that the buyer's value is uniformly distributed on $[1, 2]$.

where shifts in bargaining power have no social surplus consequences, or rely on oligopoly models in which firms on one side of the market (typically buyers) are assumed to be price takers and so have no bargaining power.

In contrast, as stated in Corollary 1, with incomplete information, a change in bargaining weights has an impact on social surplus because the efficiency of the mechanism varies with bargaining weights. Consequently, a merger that results in buyer-side and supplier-side bargaining powers moving closer together increases social surplus if the bargaining-power effects outweigh the productive-power effects of consolidation. This possibility offers comfort to a competition authority that places weight on social surplus; however, an authority that is focused on consumer surplus (and uses buyer surplus as a proxy for consumer surplus—for which the analysis in Section VI provides a foundation) would never be swayed by claims of equalization of bargaining power because then a merger is bad for the buyers for two reasons: competition among suppliers is reduced *and* the remaining suppliers have increased bargaining power.¹⁸

¹⁸In the mirror setting with a powerful supplier and a merger among buyers with little bargaining power, the merger plus equalization of bargaining weights could increase social surplus *and* buyer surplus, to the comfort of a competition authority that weights buyer surplus. Indeed, consistent with this, Kirkwood (2012, p. 1523) argues that “neither courts nor enforcement agencies have ever objected to a buy-side merger on the ground that it would create countervailing power.”

As an example, Figure 2, panel A shows a case in which a merger of suppliers reduces social surplus if the buyer has all the bargaining power both before and after the merger, but a merger increases social surplus if the buyer and suppliers' bargaining weights are equalized after the merger. Indeed, Figure 2, panel B provides an example in which an equalization of bargaining power induces the first best in the postmerger market—specifically, if the buyer has all the bargaining weight prior to the merger, then the pre-merger outcome is not the first best, but with symmetric bargaining weights in the postmerger market, the outcome is the first best. In addition, in the example of Figure 2, panel B, when the pre-merger market is efficient, a merger causes that market to become inefficient unless the postmerger market has symmetric bargaining weights—the postmerger Williams frontier touches the first-best frontier only when $\hat{\Delta} = 12$.

Transposing the roles of buyers and suppliers in Figure 2 provides an example of how consolidation among buyers that equalizes bargaining power between buyers and a dominant supplier can increase welfare. The ability of buyer mergers to equalize bargaining weights is borne out in Decarolis and Rovigatti (2021), which shows that consolidation among online advertising intermediaries has increased their buyer power, countervailing Google's significant market power in online search.

We summarize with the following result:

COROLLARY 3: A merger between two symmetric suppliers or two symmetric buyers that does not alter bargaining weights or shares and reduces social surplus is more harmful to social surplus than a merger between the same two firms that equalizes the bargaining weights between the two sides of the market. Moreover, the effects of equalizing bargaining weights associated with a merger can be so strong that the first best is possible after the merger when it was not possible before the merger.

A merger of suppliers that does not alter bargaining weights can either harm or benefit a nonmerging supplier, and the effect can vary with the nonmerging supplier's bargaining weight. Further, even if a merger harms a nonmerging supplier when that nonmerging supplier has low bargaining weight, it might benefit the nonmerging supplier when the nonmerging supplier has high bargaining weight. Considering the equalization of bargaining weights, an increase in the bargaining weight of the merged entity harms the nonmerging suppliers, potentially reversing what would have been a beneficial effect of the merger on a nonmerging supplier. In addition, the benefits to social surplus associated with an increase in the bargaining weight for the merged entity (up to the level of the buyer's bargaining weight) varies with the bargaining power of outside suppliers. For example, if the equalization of bargaining weights between the buyers and merging suppliers also equalizes their bargaining weights with those of the nonmerging suppliers, then the effect is stronger than if the nonmerging suppliers have a bargaining weight of zero.

Our analysis allows us to identify necessary conditions for a defense of a merger by suppliers based on the equalization of bargaining power. First, as mentioned, the objective of the merger review would need to include the promotion of social surplus, and not just buyer surplus. Second, the side of the market on which the merger occurs would need to have less bargaining power than the other side, so

that an increase in the merging parties' bargaining power is a movement toward the equalization of bargaining power. Third, the side opposite the merger would need to retain at least some bargaining power following the merger—for example, following a supplier merger, buyer power would need to diminish, but not vanish—so that society is not simply trading dominant buyers for dominant suppliers.

This offers an interpretation of and rationale for the European Commission (EC) merger guidelines, which state that “it is not sufficient that buyer power exists prior to the merger, it must also exist and remain effective following the merger. This is because a merger of two suppliers may reduce buyer power if it thereby removes a credible alternative” (EC 2004, para. 67).¹⁹ Our conclusions are consistent with that view insofar as the buyers must have power before a supplier merger and retain at least some power after the merger in order for a defense based on the equalization of bargaining power to make economic sense.

The necessary conditions for a defense based on the equalization of bargaining power raise the question of how one would ascertain that a firm has bargaining power. Such an evaluation will depend on the specifics of the problem at hand. For example, if a market is characterized as a k -double auction, then evidence of buyer power would be that transactions always occur at the buyer's price.²⁰ In the case of a procurement auction, evidence consistent with buyer power and inconsistent with its absence includes (i) the buyer using procurement methods that result in suppliers other than the lowest-cost supplier winning, such as handicaps or preferences; (ii) the distribution of reserve prices differing across the markets if the buyer purchases in separate markets; (iii) one observing with positive probability ties in procurement outcomes and randomization over winners.²¹ Analogous conditions apply for an analysis of supply power.

IV. Vertical Integration

We now analyze vertical integration between a buyer and a supplier. Throughout this section, to simplify the analysis, we consider settings in which the short side of the market has only one firm before integration, that is, $\min\{n^B, n^S\} = 1$, and all firms have single-unit demand and capacities. The assumption that $\min\{n^B, n^S\} = 1$ ensures that the trading position of the vertically integrated firm does not depend on type realizations. If $n^B = 1$, it can only buy, and if $n^S = 1$, it

¹⁹The EC merger guidelines also state, “Countervailing buyer power in this context should be understood as the bargaining strength that the buyer has vis-à-vis the seller in commercial negotiations due to its size, its commercial significance to the seller and its ability to switch to alternative suppliers” (EC 2004, para. 64).

²⁰This property does not hinge on particular distributional assumptions. For $k = 1$, the buyer's and supplier's optimal bids are $\Gamma_1^{-1}(v)$ and c , respectively, while for $k = 0$, they are v and $\Phi_1^{-1}(c_1)$. Hence, for $k = 1$ ($k = 0$) the k -double auction is the mechanism that is optimal for the buyer (supplier) for any distributions F_1 and G_1 with positive densities on their supports. If Φ_1 or Γ_1 is not monotone, one would replace the virtual type function with its ironed counterparts and the inverse with the generalized inverse (Myerson 1981).

²¹The background for these conditions is as follows: (i) a buyer with power discriminates among heterogeneous suppliers based on their virtual costs; (ii) a buyer without power would optimally set a reserve equal to $\min\{v, \bar{c}\}$, so even if suppliers in the different markets draw their types from different distributions, the distribution of reserves would be the same across the markets as long as the buyer's values for the goods in the markets are drawn from the same distribution and the suppliers' supports do not vary; (iii) for a buyer with power, this outcome arises when suppliers draw their costs from distributions that are identical but do not satisfy regularity, that is, their virtual costs are not monotone and so the optimal mechanism involves “ironing,” while a buyer without power purchases from the lowest-cost supplier.

can only sell, provided that it trades.²² Consequently, if buyer 1 and supplier 1 vertically integrate, then the integrated firm's willingness to pay will be $\min\{v_1, c_1\}$ if $n^B = 1$ and its cost for selling will be $\max\{v_1, c_1\}$ if $n^S = 1$. We say that a market is *one-to-one* if $\min\{n^B, n^S\} = \max\{n^B, n^S\}$ and *one-to-many* otherwise. We also assume that following vertical integration, the integrated entity can efficiently solve its internal agency problem, which is a standard assumption.

Consider first a setting with overlapping supports pre-integration (i.e., $\underline{v} < \bar{c}$). Because the first best is then impossible when there is only one buyer and one supplier, we have the following result.

PROPOSITION 6: *If supports overlap, then vertical integration increases social surplus when the number of firms is sufficiently small, regardless of bargaining weights.*

As reflected in Proposition 6, in a one-to-one market, vertical integration increases social surplus and enables the first best by essentially eliminating a Myerson-Satterthwaite problem. However, as we show next, vertical integration can also create a Myerson-Satterthwaite problem. In particular, if the pre-integration market has nonoverlapping supports, then the first best is possible in the pre-integration market and, indeed, occurs if the pre-integration bargaining weights are symmetric. In that case, vertical integration cannot possibly increase social surplus, and in some cases decreases social surplus by inducing the integrated firm to source internally for some type realizations when an outside supplier has a lower cost.

PROPOSITION 7: *Assuming a one-to-many pre-integration market with $n^B = 1 < n^S$ and $\underline{c} < \underline{v}$ and symmetric bargaining weights pre-integration, then, regardless of post-integration bargaining weights, vertical integration*

- (i) *cannot increase social surplus if $\bar{c} < \bar{v}$ and \underline{v} is sufficiently large, and decreases social surplus if $\bar{c} \leq \underline{v}$ (nonoverlapping supports);*
- (ii) *cannot increase social surplus if $G_j = G$ for all $j \in \mathcal{N}^S$ and n^S is sufficiently large.*

PROOF:

See online Appendix B.

Proposition 7 provides clear-cut conditions under which there is no efficiency rationale for vertical integration with equal bargaining weights before integration. If the buyer's and the suppliers' supports have sufficiently small overlap, that is, for $\bar{c} < \bar{v}$ and \underline{v} sufficiently large, then vertical integration cannot increase social surplus simply because the first best is already achieved without integration, and

²²This assumption substantially simplifies the derivation of, say, the second-best mechanism. Without it, the vertically integrated firm may, depending on type realizations, optimally sell, buy, or not trade at all. The analysis of problems of this kind is complicated by the fact that the integrated firm's worst-off type becomes endogenous, requiring techniques such as those developed by Loertscher and Wasser (2019). While interesting and relevant, it seems best to leave this analysis for future work.

it *decreases* social surplus when supports are nonoverlapping because it creates a Myerson-Satterthwaite problem. Likewise, with *ex ante* symmetric suppliers and $\underline{v} > \underline{c}$, there is no efficiency rationale for vertical integration if the supply side is sufficiently competitive. The first part of Proposition 7 follows by setting $\underline{v} = \bar{c}$, the second from Williams (1999, Section 3). Varying the overlap of the supports by changing \underline{v} is a way of capturing the somewhat loose notion of how much private information there is. Viewed from this angle, the first part of Proposition 7 says that if there is little private information, then there are no gains from vertical integration. Put differently, private information is necessary for an efficiency rationale for vertical integration. The second part says that vertical integration is less likely to increase social surplus in otherwise highly competitive environments, which resonates with intuition and insights from oligopoly models (see, e.g., Riordan 1998, Loertscher and Reisinger 2014).²³ With little competition, rival suppliers have large markups and react to vertical integration by reducing markups *and* quantities. In contrast, in highly competitive markets, price is already close to marginal costs, so the outside suppliers can essentially only reduce their quantities. This leads to an increase in consumer price and thereby to the perhaps paradoxical result that vertical integration is anticompetitive in otherwise competitive environments.

Of course, with symmetric pre-integration bargaining weights, analogous results hold for $n^S = 1 < n^B$, in which case $\underline{c} < \underline{v}$ and \bar{c} sufficiently small imply that vertical integration decreases social surplus, regardless of post-integration bargaining weights, and if $F_i = F$ for all $i \in \mathcal{N}^B$ and $\bar{c} < \bar{v}$, then vertical integration cannot increase social surplus for n^B sufficiently large.

Propositions 6 and 7 provide conditions under which vertical integration either always increases or always decreases social surplus. At the heart of both results is the fact that the efficiency of the price-formation process is endogenous in incomplete information bargaining. The elimination of a Myerson-Satterthwaite problem through vertical integration is the incomplete information analog to eliminating the classic double mark-up problem. In contrast to the complete information literature, however, there is now a new effect, namely that trade becomes less efficient for the nonintegrated firms. Further, it is possible for this latter effect to dominate so that the market as a whole is made less efficient as a result of vertical integration.²⁴

The results of Propositions 6 and 7 are robust in that they do not depend on specific assumptions about distributions or beliefs of firms. Indeed, because there is always a dominant strategy implementation of the incomplete information bargaining mechanism, beliefs play no role. Moreover, we obtain social surplus decreasing vertical integration without imposing *any* restrictions on the contracting space and without invoking exertion of market power by any player (above and beyond requiring individual-rationality and incentive-compatibility constraints to be satisfied). These are noticeable differences relative to the post-Chicago school literature on

²³Riordan (1998) shows that vertical integration by a dominant firm that faces a competitive fringe is anticompetitive in that it raises the equilibrium price that downstream consumers face. Loertscher and Reisinger (2014) extend Riordan's model by replacing the assumption of a competitive fringe with a Cournot oligopoly and show that vertical integration is more likely to be harmful the larger is the number of firms in the Cournot oligopoly.

²⁴This occurs, for example, with $n^B = 1$, $n^S = 2$, and symmetric bargaining weights if F is uniform on $[0, 1]$ and for $j \in \{1, 2\}$, $G_j(c) = c^{1/10}$, also with support $[0, 1]$. Then vertical integration causes expected social surplus to decrease from 0.4827 to 0.4815.

vertical contracting and integration, whose predictions rely on assumptions about beliefs, feasible contracts, and/or market power.²⁵ Maybe more important, in our incomplete information setting, any benefits and costs of vertical integration are pinned to the primitives of the problem, which contrasts with complete information settings, where these hinge on restrictions on the contracting space. As argued persuasively by Choné, Linnemer, and Vergé (2021), this is a matter of substance rather than taste.

Of course, our results do rely, inevitably, on support assumptions. In the case of nonoverlapping supports, the first best is possible before but not after vertical integration. Complete information settings correspond to the limit case when all supports become singletons.

V. Investment

Investment incentives feature prominently, and at times controversially, in concurrent policy debates,²⁶ and they have been at center stage in the theory of the firm since Grossman and Hart (1986) and Hart and Moore (1990) (G-H-M hereafter). To account for investment, we now extend our model by adding investment as an action taken by each firm prior to the realization of private information, where investment improves a firm's type distribution. We show that the results under incomplete information differ starkly from those obtained in the G-H-M literature. This literature stipulates complete information and efficient bargaining, i.e., a bargaining process that achieves the first best, and, as a consequence, results in hold-up and inefficient investment. In contrast, in our setting, incomplete information protects firms from hold-up, and investments are efficient if and, under additional assumptions, only if bargaining is efficient.

Investments do not change the supports of the distributions, which is in line with the assumption that types are private information insofar as an agent's investment does not affect its worst-off type. Supplier $j \in \mathcal{N}^S$ making investment e_j^S incurs cost $\Psi_j^S(e_j^S)$, and buyer $i \in \mathcal{N}^B$ making investment e_i^B incurs cost $\Psi_i^B(e_i^B)$. Consistent with G-H-M, we assume that investments are not contractible.²⁷ Thus, bargaining only depends on equilibrium investments and does not vary with off-the-equilibrium-path investments. One implication of this is that the interim expected payments to the worst-off types of firms are not affected by actual investments. We assume that the buyers and suppliers first simultaneously make their investments and then bargaining takes place.

²⁵For an overview, see Riordan (2008). On the sensitivity of complete information vertical contracting results to assumptions of "symmetric," "passive," and "wary" beliefs see, e.g., McAfee and Schwartz (1994).

²⁶For example, related to the 2017 Dow-DuPont merger, the US DOJ's *Competitive Impact Statement* identifies reduced innovation as a key concern (<https://www.justice.gov/atr/case-document/file/973951/download>, pp. 2, 10, 15, 16). Interestingly, countervailing power was also an issue in the Dow-DuPont merger. The European Commission analyzed "whether significant buyer power exists to compensate for any potential added market power from the Parties" and the parties argued that they "face substantial countervailing bargaining power by their sophisticated customers, namely distributors and agricultural cooperatives," but the EC concluded that "the limited countervailing buyer power would be insufficient to off-set the anticompetitive concerns raised by the Transaction given that nonlarge customers do not have buyer power" (EC CASE M.7932 – Dow/DuPont, http://ec.europa.eu/competition/mergers/cases/decisions/m7932_13668_3.pdf, paras. 434, 528, 3565).

²⁷This assumption also prevents the mechanism from using harsh punishments for deviations from any prescribed investment level.

We first consider the planner’s problem of determining investments when the allocation rule is first best. Using the definition of first-best welfare W^{FB} in (5), the first-best investments, $\bar{\mathbf{e}}$, solve the planner’s first-best investment problem: $\max_{\mathbf{e}} E_{\mathbf{v}, \mathbf{c} | \mathbf{e}} [W^{FB}(\mathbf{v}, \mathbf{c})] - \sum_{i \in \mathcal{N}^B} \Psi_i^B(e_i^B) - \sum_{j \in \mathcal{N}^S} \Psi_j^S(e_j^S)$.

Now consider the firms’ incentives to invest when incomplete information bargaining is such that the first best is possible (see, e.g., Proposition 3 for conditions under which this is the case without symmetric bargaining weights). By the payoff equivalence theorem, it follows that, up to a constant, any incentive compatible mechanism generates the same interim and consequently the same ex ante expected utility for every firm. Thus, for the case considered here in which the first best is possible, we can, without loss of generality, focus on expected utilities in the VCG mechanism. Given a type realization (\mathbf{v}, \mathbf{c}) , supplier j ’s VCG payoff is $W^{FB}(\mathbf{v}, \mathbf{c}) - W^{FB}(\mathbf{v}, \bar{\mathbf{c}}, \mathbf{c}_{-j})$, plus possibly a constant. Likewise, buyer i ’s payoff is $W^{FB}(\mathbf{v}, \mathbf{c}) - W^{FB}(\underline{\mathbf{v}}, \mathbf{v}_{-i}, \mathbf{c})$, plus possibly a constant.

Taking expectations over (\mathbf{v}, \mathbf{c}) , and noticing that $W^{FB}(\mathbf{v}, \bar{\mathbf{c}}, \mathbf{c}_{-j})$ is independent of supplier j ’s type and its distribution, and so independent of e_j^S , it follows that each supplier j ’s problem at the investment stage, taking as given that the other firms choose investments $\bar{\mathbf{e}}_{-j}$, is $\max_{e_j^S} E_{\mathbf{v}, \mathbf{c} | e_j^S, \bar{\mathbf{e}}_{-j}} [W^{FB}(\mathbf{v}, \mathbf{c})] - \Psi_j^S(e_j^S)$. An analogous optimization problem applies to buyer i ’s choice of e_i^B , noting that $W^{FB}(\underline{\mathbf{v}}, \mathbf{v}_{-i}, \mathbf{c})$ is independent of buyer i ’s type and its distribution, and so independent of e_i^B . It then follows that the planner’s solution $\bar{\mathbf{e}}$ is a Nash equilibrium if incomplete information bargaining permits the first best. This proves the first part of Proposition 8 below.

Under additional conditions, the converse is also true; that is, $\bar{\mathbf{e}}$ being a Nash equilibrium outcome in the game in which firms’ first-stage investments are followed by incomplete information bargaining implies that bargaining is efficient. Given investments \mathbf{e} , for $j \in \mathcal{N}^S$, let $G_j(\cdot; e_j^S)$ and for $i \in \mathcal{N}^B$, let $F_i(\cdot; e_i^B)$ denote supplier j ’s and buyer i ’s type distributions, respectively, with virtual type functions assumed to be monotone. Sufficient conditions for the converse to hold are, for all $j \in \mathcal{N}^S$ and $i \in \mathcal{N}^B$,

$$(10) \quad \Psi_j^S(0) = \Psi_i^B(0) = 0, \text{ and for all } e > 0, \Psi_j^S(e), \Psi_i^B(e) > 0$$

$$\text{and } \Psi_j^{S''}(e), \Psi_i^{B''}(e) > 0;$$

for all $c \in (\underline{c}, \bar{c})$ and $v \in (\underline{v}, \bar{v})$,

$$(11) \quad \frac{\partial G_j(c; e)}{\partial e} > 0 \text{ and } \frac{\partial F_i(v; e)}{\partial e} < 0;$$

and either (i) the type distributions have overlapping supports, $\underline{v} < \bar{c}$, (ii) $K^B = K^S$, (iii) $K^B < K^S$ and for all $j \in \mathcal{N}^S$ and $c \in [\underline{c}, \bar{c}]$,

$$(12) \quad G_j(c; \bar{e}_j^S) \equiv G(c),$$

or (iv) $K^B > K^S$ and for all $i \in \mathcal{N}^B$ and $v \in [\underline{v}, \bar{v}]$,

$$(13) \quad F_i(v; \bar{e}_i^B) \equiv F(v).$$

Conditions (10) and (11) imply that the first-best investments \bar{e} are positive and determined by first-order conditions. This allows one to show that when first-best investments are a Nash equilibrium, the total number of trades under incomplete information bargaining is the same as under the first best. Given any one of the remaining conditions (i)–(iv), one can show further that it is the same set of buyers and suppliers that trade in the Nash equilibrium as under the first best.²⁸

PROPOSITION 8: *First-best investments are a Nash equilibrium outcome of the simultaneous investment game if incomplete information bargaining is efficient. Conversely, assuming that (10) and (11) and at least one of (i)–(iv) above holds, if first-best investments are a Nash equilibrium outcome, then incomplete information bargaining is efficient.*

PROOF:

See online Appendix B.

As shown in Proposition 8, when incomplete information bargaining is efficient, the firms' Nash equilibrium investment choices are first-best investments. Because private information protects firms from hold-up,²⁹ efficient incomplete information bargaining implies efficient investments.³⁰ Intuitively, given that the allocation rule is efficient, each firm is the residual claimant to the surplus that its investment generates. Anticipating that this will be the case once types are realized, each firm's incentives are also aligned with the planner's at the investment stage because each firm's and the planner's reward from investment are the same. Further, under additional conditions, any inefficiency in bargaining results in inefficient investments.

Combining Proposition 8 with Corollary 3 allows us to connect investment with the equalization of bargaining power. While the equalization of bargaining power can increase social surplus holding investments fixed, as in Corollary 3, Proposition 8 shows that it can also improve investments to the first-best level. Proposition 8 thus provides an additional channel—investments—through which changes in bargaining power can increase social surplus.

While Proposition 8 focuses on investments that improve firms' own types, the first part of Proposition 8 continues to hold if, for example, there is a single buyer and each supplier can invest in the "quality" of its product, thereby increasing the value of its product to the buyer.³¹ Our result does not hold if, for example, investment generates

²⁸ Proposition 8 connects to the equivalence result of Hatfield, Kojima, and Kominers (2018), which links efficient dominant-strategy mechanisms under incomplete information with efficient investments, and to earlier work by Milgrom (1987) and Rogerson (1992). A difference is that the no-deficit constraint in our setting may preclude the first best.

²⁹ Lauerermann (2013) finds that private information protects against hold-up in a dynamic search model, finding that it is easier/possible to converge to Walrasian efficiency with private information, but otherwise hold-up prevents convergence to efficiency. This is consistent with our results, interpreting search as investment.

³⁰ In a setup where efficient bargaining is possible because of shared ownership (rather than the absence of any allocation-relevant private information), Schmitz (2002, p. 176) notes that "Intuitively, ... a party's ex ante expected utility from an ex post efficient mechanism is (up to a constant) equal to the total expected surplus, so that each party is residual claimant on the margin from his or her point of view."

³¹ This result contrasts with that of Che and Hausch (1999), who study a contracting setup in which investments by suppliers in cost reduction are efficient, but investments by suppliers that benefit the buyer need not be. Importantly, however, there is no incomplete information at the price-formation stage in their model.

externalities, e.g., if there are technology spillovers across suppliers or if investment increases the buyer's value regardless of its trading partner. Using Proposition 8, one can connect investment with vertical integration. As shown in online Appendix F, depending on conditions, vertical integration can either promote or disrupt efficient investment.

We now illustrate how equilibrium investments are affected by bargaining power and by the extent to which the supports of the value and cost distributions overlap. We parameterize the firms' type distributions and allow investment to affect the distributional parameter in a way that improves the distribution in a first-order stochastic dominance sense, where investment results in a dominating distribution for buyers and a dominated distribution for suppliers.

Specifically, we consider a bilateral trade setup with linear virtual types. We hold fixed the support of the supplier's distribution at $[0, 1]$ and let the support of the buyer's distribution be $[\underline{v}, \underline{v} + 1]$, where we vary \underline{v} from 0 to 1. We fix $X > 0$ and consider a supplier type distribution of $G_{e_S}(c) \equiv c^{X-e_S}$ with support $[0, 1]$, where $e_S \in [0, X]$ is the supplier's investment, and a buyer type distribution of $F_{e_B}(v) \equiv 1 - (1 + \underline{v} - v)^{X-e_B}$ with support $[\underline{v}, \underline{v} + 1]$, where $e_B \in [0, X]$ is the buyer's investment. We assume that each firm's investment e has cost $e^2/2$. Relegating the details to online Appendix E, we illustrate the effects of bargaining power and the distributional supports on equilibrium investment in Figure 3. As the figure shows, each firm's equilibrium investment is maximized away from extreme bargaining weights. This illustrates the additional benefit of the equalization of bargaining weights mentioned above; namely, that it has the potential to improve the efficiency of investment, in some cases to the first best.

VI. Extension: Downstream Consumers

Competition authorities commonly put weight on the welfare of final consumers.³² With that in mind, we now extend the model to incorporate final consumers. A natural and tractable way of doing this is to assume that each buyer in our model is a retailer that has exclusive access to a downstream market. To fix ideas, we assume that the good that is being procured in the incomplete information bargaining mechanism is an input that improves the quality of the product that buyers sell in their downstream markets and focus, for now, on the case with $k_i^B = 1$ for all $i \in \mathcal{N}^B$. Specifically, letting $P_i(Y)$ be the willingness to pay of a typical consumer in market i for the Y -th unit of a good of quality 1, we assume that the willingness to pay for the Y -th unit of a good of quality $\theta > 0$ is $\theta P_i(Y)$.

We normalize the quality of the good without improvement to 1 and assume that suppliers offer quality-improving inputs, so that trade between a buyer and a supplier increases the quality of the buyer's product to some commonly known $\theta > 1$. Thus, the incremental quality improvement of the input is $\theta - 1$. We assume further that the marginal cost of production of each buyer is zero and that the private information of each buyer i pertains to the mass $\omega_i > 0$ of identical consumers in its

³²While common, this practice has recently been challenged by Hemphill and Rose (2018), who argue that the mission of antitrust merger review is to protect the welfare of the merging firms' trading partners, whether they are purchasers or sellers.

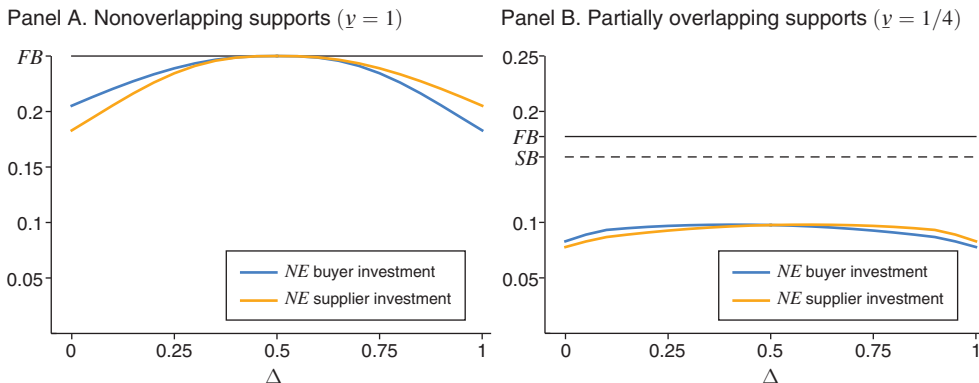


FIGURE 3. NASH EQUILIBRIUM INVESTMENTS

Notes: Nash equilibrium (NE) investments with bargaining weights $(w^S, w^B) = (1 - \Delta, \Delta)$ for buyer distributions with varying supports. Assumes the linear virtual type setup for bilateral trade with $F(v) = 1 - (1 + \underline{v} - v)^{1.25 - e_B}$, where $e_B \in [0, 1.25]$ is the buyer’s investment, and $G(c) = c^{1.25 - e_S}$, where $e_S \in [0, 1.25]$ is the supplier’s investment. Investment e has cost $e^2/2$. When $\underline{v} = 1$, we obtain $e^{FB} = e^{SB} = 0.25$, implying that first-best (and second-best) investment levels result in uniformly distributed types. For $\underline{v} = 1$, $\rho^{NE} = \max\{w_S, w_B\}$ for all bargaining weights, and for $\underline{v} = 1/4$, $\rho^{NE} > \max\{w_S, w_B\}$ for all bargaining weights.

downstream market. The inverse demand function $P_i(Y)$ is assumed to be decreasing and to be such that $YP_i(Y)$ has a unique maximum, with the maximizer being denoted by Y_i^* .

Given quality θ , in equilibrium, per capita consumer surplus in market i is $\theta \left[\int_0^{Y_i^*} P_i(y) dy - Y_i^* P_i(Y_i^*) \right]$. The increases in per capita consumer surplus and profit from the quality-improving input are thus $(\theta - 1) \left(\int_0^{Y_i^*} P_i(y) dy - Y_i^* P_i(Y_i^*) \right)$ and $(\theta - 1) Y_i^* P_i(Y_i^*)$, respectively. For a given realization of the mass of consumers in market i , ω_i , the buyer’s willingness to pay for the input is $v_i \equiv \omega_i (\theta - 1) Y_i^* P_i(Y_i^*)$, while a competition authority with a consumer surplus standard values the input at $\sigma_i \equiv \omega_i (\theta - 1) \left(\int_0^{Y_i^*} P_i(y) dy - Y_i^* P_i(Y_i^*) \right)$. Consequently, letting $\gamma_i \equiv \frac{\int_0^{Y_i^*} P_i(y) dy}{Y_i^* P_i(Y_i^*)} - 1$ be a market-specific constant, which is positive because $P_i(Y)$ is decreasing, we have $\sigma_i = \gamma_i v_i$.

With downstream consumers, both the social planner who aims at maximizing equally weighted social surplus and a competition authority whose objective is consumer surplus will take the γ_i s into account. The social planner will attach a weight of $\gamma_i + 1$ to buyer i ’s value (and a weight of 1 to each supplier), so that the planner’s Lagrangian becomes

$$E_{\mathbf{v}, \mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} (\gamma_i + 1) v_i Q_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} c_j Q_j^S(\mathbf{v}, \mathbf{c}) + (\rho - 1) \left(\sum_{i \in \mathcal{N}^B} \Phi_i(v_i) Q_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j(c_j) Q_j^S(\mathbf{v}, \mathbf{c}) \right) \right]$$

minus $\sum_{i \in \mathcal{N}^B} (\rho - 1) \hat{u}_i^B(\underline{v}) + \sum_{j \in \mathcal{N}^S} (\rho - 1) \hat{u}_j^S(\bar{c})$. This is the same as (6) with $\mathbf{w} = \mathbf{1}$ and the addition of consumer surplus, which is $\sum_{i \in \mathcal{N}^B} \gamma_i v_i Q_i^B(\mathbf{v}, \mathbf{c})$. By

contrast, the Lagrangian for a competition authority with a consumer surplus objective is simply

$$E_{\mathbf{v}, \mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} \gamma_i v_i Q_i^B(\mathbf{v}, \mathbf{c}) + \rho \left(\sum_{i \in \mathcal{N}^B} \Phi_i(v_i) Q_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j(c_j) Q_j^S(\mathbf{v}, \mathbf{c}) \right) \right]$$

minus $\sum_{i \in \mathcal{N}^B} \rho \hat{u}_i^B(\underline{v}) + \sum_{j \in \mathcal{N}^S} \rho \hat{u}_j^S(\bar{c})$.

If the “pass-through” of a retailer’s profit to consumer surplus is the same across markets, that is, $\gamma_i = \gamma$ for all $i \in \mathcal{N}^B$, then neither the social planner nor the competition authority discriminates among buyers in the weights that they attach to them. In this case, any discrimination in their mechanisms is due to the no-deficit constraint being binding, which implies that there is some discrimination based on virtual values, provided that for some buyers $i \neq h$ we have $F_i \neq F_h$. Simply having a larger expected market is not a reason for discrimination in this case because buyer i ’s market being larger than h ’s would be reflected in F_i and F_h . In contrast, when $\gamma_i \neq \gamma_h$, then both the social planner and the competition authority would discriminate across downstream markets. For the social planner, this is true even if the no-deficit constraint does not bind, while for a competition authority with a consumer surplus standard, the no-deficit constraint always binds. Last, if $k_i^B > 1$ for some i , this means that buyer i has demand for multiple quality-improving inputs, and the analysis above extends by replacing $\theta - 1$ with $k_i^B \theta - 1$.

The model with downstream consumers is a setting in which a social surplus maximizing planner favors buyers relative to suppliers and a competition authority with a consumer surplus standard exclusively puts weight on buyers’ values. In this sense, the analysis in this section provides a foundation for the use of buyer surplus as a proxy for consumer surplus.

The incomplete information bargaining framework can accommodate a range of other extensions, as we show in online Appendix C. These include allowing variation in agents’ outside options; allowing buyers to have preferences over suppliers, in which case bargaining externalities arise naturally; and showing that one can use the results of Delacrétaz et al. (2019) to generalize the setup to allow buyers to have preferences over suppliers, which also naturally leads to bargaining externalities.

VII. Related Literature

The independent private values setting with continuous distributions has the virtue that, for a given objective, the mechanism that maximizes this objective, subject to incentive-compatibility, individual-rationality, and no-deficit constraints, is well defined and pinned down (up to a constant in the payments) by the allocation rule, which is unique. Antitrust authorities regularly face the task of evaluating the competitive effects of mergers in settings in which the payments that merging firms receive for their products are determined through competitive procurements, exactly because buyers have incomplete information regarding the suppliers’ costs. Of particular interest to industrial organization and antitrust economics, this setting also has the feature that, quite generally, there is a trade-off between allocating efficiently and extracting rents. This trade-off is at the heart of both industrial organization and

Myerson's optimal auction. This trade-off is the reason why the Williams frontier is typically not identical to the 45-degree line and, therefore, the basis from which the possibility of social-surplus-increasing equalization of bargaining power emerges. Moreover, the aforementioned assumptions are essentially the *only* assumptions that permit a tractable approach that maintain the basic trade-off between profit and social surplus.³³

There has also been a recent upsurge of interest in bargaining (see, for example, Backus et al. 2020; Backus, Blake, and Tadelis 2019; Zhang, Manchanda, and Chu 2021; Larsen 2021; Byrne, Martin, and Nah 2021), and buyer power (see, for example, Snyder 1996, Nocke and Thanassoulis 2014, Caprice and Rey 2015, Loertscher and Marx 2019, Decarolis and Rovigatti 2021). Larsen and Zhang (2021) emphasize the value in abstracting away from the rules or extensive form of a game and instead focusing on outcomes, e.g., allocations and transfers, to estimate bargaining weights and distributions that can then be used for the analysis of counterfactuals. Bargaining has also come to the forefront of the empirical industrial organization literature, in particular in analyses of bundling and vertical integration such as Crawford and Yurukoglu (2012) and Crawford et al. (2018). Collard-Wexler, Gowrisankaran, and Lee (2019) and Rey and Vergé (2019) provide recent theoretical foundations for the widely used Nash-in-Nash bargaining model.³⁴ Ho and Lee (2017) apply this framework to the question of countervailing power by insurers when negotiating with hospitals and find evidence that consolidation among insurers improves their bargaining position vis-à-vis hospitals. Our paper contributes to this literature by showing, among other things, that in incomplete information models, bargaining breakdown occurs on the equilibrium path, and that the probability of breakdown can, under suitable assumptions, be used to estimate distributions (see online Appendix C). Ausubel, Cramton, and Deneckere (2002) explicitly account for inefficiencies in bargaining and focus on the second-best mechanisms introduced by Myerson and Satterthwaite (1983), as do we; however, they focus on the robustness of the Bayesian mechanism design setting in two-person bargaining, which appears not to be a central concern for applied work, given the frequent reliance on models based on Nash bargaining, in which agents literally know each other's types.

Consistent with our results, the literature on vertical integration and foreclosure also notes that a vertical merger that eliminates internal frictions may create or exacerbate external ones for the case in which buyers are competing downstream intermediaries. For an overview of the literature on the competitive effects of vertical integration, see Riordan and Salop (1995). As described there, the literature takes the view that most

³³ 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, among other things, 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 trade-off 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 trade-off between profit and social surplus (Mezzetti 2004, 2007). Notwithstanding recent progress, with multidimensional private information and multiple agents, the optimal mechanism is not known (see, e.g., Daskalakis, Deckelbaum, and Tzamos, 2017). With discrete types, there is no payoff equivalence theorem. In other words, the mechanism is not pinned down by the allocation rule.

³⁴ While the empirical literature examining multilateral bargaining focuses on fixed quantities or linear tariffs, Rey and Vergé (2019) allow for nonlinear tariffs, take into account the impact of these tariffs on downstream competition (placing it outside the approach of Collard-Wexler, Gowrisankaran, and Lee 2019), and provide a microfoundation for Nash-in-Nash.

vertical mergers lead to some efficiencies. Ordoover, Saloner, and Salop (1990) and Salinger (1988) show that vertical integration leads to an increase in rivals' (linear) prices and Hart and Tirole (1990) provide a similar insight in the context of secret contracting, without restriction to linear tariffs. Nocke and Rey (2018) and Rey and Vergé (2019) extend the latter insight to multiple strategic suppliers for Cournot and Bertrand downstream competition. Allain, Chambolle, and Rey (2016) show that, while vertical integration solves hold-up problems for the merging parties, it may also create or exacerbate problems for rivals.

The incomplete information approach also has implications for two-stage models in which investments precede bargaining, which have been at the center of attention in incomplete contracting models in the tradition of G-H-M. As discussed, the predictions could hardly differ more starkly because with incomplete (complete) information efficient bargaining implies efficient (inefficient) investment.³⁵ There has also been a recent upsurge of interest in industrial organization relating to market structure and the incentives to invest (see, e.g., Federico, Langus, and Valletti 2017, 2018; Jullien and Lefouili 2018; Loertscher and Marx 2019), onto which our paper—in particular, the results pertaining to mergers and vertical integration—sheds new light as well.

VIII. Conclusions

We provide an incomplete information bargaining model suitable for analyzing a range of important issues in industrial organization. In a methodological contribution, we show how one can allow multiple buyers and multiple suppliers, with multiunit demand and supply, while still maintaining the assumption of one-dimensional private information. In our setup, the social surplus increasing effect of an equalization of bargaining power arises naturally because of the inherent trade-off between social surplus and rent extraction: with independent private values, neither the mechanism that is optimal for buyers nor the one that is optimal for the suppliers is efficient in general, which opens the scope for increasing social surplus by making bargaining powers more equal. We show that socially harmful vertical integration arises naturally in our setting. We also examine the relation between the efficiency of incomplete information bargaining and the incentives to invest, which differs fundamentally from what obtains in complete information models that are based on the assumption that efficient trade is always possible. In extensions, we show that one can incorporate effects on downstream consumers.

Our paper shows that an economic agent's strength or weakness has two dimensions that are, conceptually, independent. The first one reflects the agent's productivity. Is the agent likely to have a high value if it is a buyer or a low cost if it is a supplier? The second dimension captures the agent's bargaining power, that is,

³⁵The tight connection between incentives for efficient investment and efficient allocation in incomplete information models has its roots in the seminal works of Vickrey (1961), Clarke (1971), and Groves (1973) and the subsequent uniqueness results of Green and Laffont (1977) and Holmström (1979). Essentially, dominant strategy incentive compatibility under incomplete information requires each agent to be a price taker, and efficiency then further requires this price to be equal to the agent's social marginal product (or cost). As demonstrated by Milgrom (1987); Rogerson (1992); Segal and Whinston (2011); Hatfield, Kojima, and Kominers (2018); and Loertscher and Riordan (2019), this is precisely the set of conditions that have to be satisfied for incentives for investment to be aligned with efficiency.

its ability (or inability) to affect bargaining in its favor. For example, consider a supplier whose bargaining power allows it to make a take-it-or-leave-it offer to a buyer that depends on the realization of the supplier's cost. The supplier optimally customizes its offer to the productivity of the buyer, with a weaker buyer (in the sense of hazard rate dominance) receiving a lower offer on average. That the weak buyer receives a better offer than the strong buyer does not reflect differences in bargaining power between them as commonly understood since one would typically not explain that economy airfares are lower than business airfares by suggesting that economy customers have greater bargaining power. What is indicative of the relative bargaining powers is then not so much the level of prices, but rather the price-formation process itself. For example, in a bilateral trade setting, if the buyer (supplier) always makes the price offer, then one would conclude that the buyer (supplier) has all the bargaining power, indicating that there is scope for social benefits from an equalization of bargaining power. In contrast, if the buyer and supplier participate in a k -double auction with $k = 1/2$, then this may be indicative of equal bargaining powers, suggesting that there is no scope for equalization of bargaining power.

Avenues for future research are many. Among other things, developing a better understanding of what determines bargaining power would add considerable value. The distinction between productive strength and bargaining power brought to light in the present paper may prove useful in that regard.

REFERENCES

- Allain, Marie-Laure, Claire Chambolle, and Patrick Rey. 2016. "Vertical Integration as a Source of Hold-Up." *Review of Economic Studies* 83 (1): 1–25.
- Australian Competition and Consumer Commission (ACCC). 1999. *Merger Guidelines*. Canberra, Australia: Australian Competition and Consumer Commission.
- Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere. 2002. "Bargaining with Incomplete Information." In *Handbook of Game Theory*, Vol. 3, edited by Robert Aumann and Sergiu Hart, 1897–1945. Amsterdam: Elsevier.
- Backus, Matthew, Thomas Blake, and Steven Tadelis. 2019. "On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining." *Journal of Political Economy* 127 (4): 1599–1628.
- Backus, Matthew, Thomas Blake, Bradley Larsen, and Steven Tadelis. 2020. "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions." *Quarterly Journal of Economics* 135 (3): 1319–61.
- Bhattacharyya, Sugato, and Amrita Nain. 2011. "Horizontal Acquisitions and Buying Power: A Product Market Analysis." *Financial Economics* 99 (1): 97–115.
- Börger, Tilman. 2015. *An Introduction to the Theory of Mechanism Design*. Oxford: Oxford University Press.
- Boyd, Stephen, and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge, UK: Cambridge University Press.
- Byrne, David P., Leslie A. Martin, and Jia Sheen Nah. 2021. "Price Discrimination, Search, and Negotiation in an Oligopoly: A Field Experiment in Retail Electricity." Unpublished.
- Caprice, Stéphane, and Patrick Rey. 2015. "Buyer Power from Joint Listing Decision." *Economic Journal* 125 (589): 1677–1704.
- Chatterjee, Kalyan, and William Samuelson. 1983. "Bargaining under Incomplete Information." *Operations Research* 31 (5): 835–51.
- Che, Yeon-Koo, and Donald B. Hausch. 1999. "Cooperative Investments and the Value of Contracting." *American Economic Review* 89 (1): 125–47.
- Choné, Philippe, Laurent Linnemer, and Thibaud Vergé. 2021. "Double Marginalization and Vertical Integration." CESifo Working Paper 8971.
- Clarke, Edward H. 1971. "Multipart Pricing of Public Goods." *Public Choice* 11 (1): 17–33.

- Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee. 2019. "Nash-in-Nash": A Micro-foundation for Applied Work." *Journal of Political Economy* 127 (1): 163–95.
- Crawford, Greg, and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review* 102 (2): 643–85.
- Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu. 2018. "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica* 86 (3): 891–954.
- Crémer, Jacques, and Richard P. McLean. 1985. "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent." *Econometrica* 53 (2): 345–61.
- Crémer, Jacques, and Richard P. McLean. 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56 (6): 1247–57.
- Daskalakis, Constantinos, Alan Deckelbaum, and Christos Tzamos. 2017. "Strong Duality for a Multiple-Good Monopolist." *Econometrica* 85 (3): 735–67.
- Decarolis, Francesco, and Gabriele Rovigatti. 2021. "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising." *American Economic Review* 111 (10): 3299–3327.
- Delacrétaz, David, Simon Loertscher, Leslie M. Marx, and Tom Wilkening. 2019. "Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade." *Journal of Economic Theory* 179: 416–54.
- European Commission (EC). 2004. *Guidelines on the Assessment of Horizontal Mergers under the Council Regulation on the Control of Concentrations between Undertakings*. European Union: Official Journal of the European Union 2004/C 31/03.
- Federico, Giulio, Gregor Langus, and Tommaso Valletti. 2017. "A Simple Model of Mergers and Innovation." *Economics Letters* 157: 136–40.
- Federico, Giulio, Gregor Langus, and Tommaso Valletti. 2018. "Horizontal Mergers and Product Innovation." *International Journal of Industrial Organization* 59: 1–23.
- Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.
- Galbraith, John Kenneth. 1954. "Countervailing Power." *American Economic Review* 44 (2): 1–6.
- Green, Jerry, and Jean-Jacques Laffont. 1977. "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." *Econometrica* 45 (2): 427–38.
- Gresik, Thomas A., and Mark A. Satterthwaite. 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 (1): 304–32.
- Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94 (4): 691–719.
- Groves, Theodore. 1973. "Incentives in Teams." *Econometrica* 41 (4): 617–31.
- Hart, Oliver, and Jean Tirole. 1990. "Vertical Integration and Market Foreclosure." *Brookings Papers on Economic Activity: Microeconomics* 20: 205–86.
- Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy* 98 (6): 1119–58.
- Hatfield, John William, Fuhito Kojima, and Scott Duke Kominers. 2018. "Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence." SSRN 2544951.
- Hatfield, John William, Scott Duke Kominers, Richard Lowery, and Jordan M. Barry. 2020. "Collusion in Markets with Syndication." *Journal of Political Economy* 128 (10): 3779–3819.
- Hemphill, Scott C., and Nancy L. Rose. 2018. "Mergers that Harm Sellers." *Yale Law Journal* 127 (7): 2078–2109.
- Ho, Kate, and Robin S. Lee. 2017. "Insurer Competition in Health Care Markets." *Econometrica* 85 (2): 379–417.
- Holmström, Bengt. 1979. "Groves' Scheme on Restricted Domains." *Econometrica* 47 (5): 1137–44.
- Jullien, Bruno, and Yassine Lefouili. 2018. "Horizontal Mergers and Innovation." *Journal of Competition Law and Economics* 14 (3): 364–92.
- Kirkwood, John B. 2012. "Powerful Buyers and Merger Enforcement." *Boston University Law Review* 92 (5): 1485–1559.
- Krishna, Vijay. 2010. *Auction Theory*. 2nd ed. Amsterdam: Elsevier.
- Larsen, Bradley J. 2021. "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions." *Review of Economic Studies* 88 (2): 851–82.
- Larsen, Bradley J., and Anthony Lee Zhang. 2021. "Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry." Unpublished.
- Lauermann, Stephan. 2013. "Dynamic Matching and Bargaining Games: A General Approach." *American Economic Review* 103 (2): 663–89.
- Loertscher, Simon, and Cédric Wasser. 2019. "Optimal Structure and Dissolution of Partnerships." *Theoretical Economics* 14 (3): 1063–1114.

- Loertscher, Simon, and Claudio Mezzetti.** 2019. "The Deficit on Each Trade in a Vickrey Double Auction Is at Least as Large as the Walrasian Price Gap." *Journal of Mathematical Economics* 84: 101–06.
- Loertscher, Simon, and Leslie M. Marx.** 2019. "Merger Review for Markets with Buyer Power." *Journal of Political Economy* 127 (6): 2967–3017.
- Loertscher, Simon, and Markus Reisinger.** 2014. "Market Structure and the Competitive Effects of Vertical Integration." *RAND Journal of Economics* 45 (3): 471–94.
- Loertscher, Simon, and Michael H. Riordan.** 2019. "Make and Buy: Outsourcing, Vertical Integration, and Cost Reduction." *American Economic Journal: Microeconomics* 11 (1): 105–23.
- Maskin, Eric S., and John G. Riley.** 1984. "Optimal Auctions with Risk Averse Buyers." *Econometrica* 52 (6): 1473–518.
- Matthews, Steven A.** 1984. "On the Implementability of Reduced Form Auctions." *Econometrica* 52 (6): 1519–22.
- McAfee, R. Preston, and Marius Schwartz.** 1994. "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity." *American Economic Review* 84 (1): 210–30.
- Mezzetti, Claudio.** 2004. "Mechanism Design with Interdependent Valuations: Efficiency." *Econometrica* 72 (5): 1617–26.
- Mezzetti, Claudio.** 2007. "Mechanism Design with Interdependent Valuations: Surplus Extraction." *Economic Theory* 31 (3): 473–88.
- Milgrom, Paul.** 1987. "Adverse Selection without Hidden Information." Unpublished.
- Mussa, Michael, and Sherwin Rosen.** 1978. "Monopoly and Product Quality." *Journal of Economic Theory* 18 (2): 301–17.
- Myerson, Roger B.** 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1): 58–73.
- Myerson, Roger B., and Mark A. Satterthwaite.** 1983. "Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory* 29 (2): 265–81.
- Mylovanov, Tymofiy, and Thomas Tröger.** 2014. "Mechanism Design by an Informed Principal: Private Values with Transferable Utility." *Review of Economic Studies* 81 (4): 1668–707.
- Nocke, Volker, and John Thanassoulis.** 2014. "Vertical Relations under Credit Constraints." *Journal of the European Economic Association* 12 (2): 337–67.
- Nocke, Volker, and Patrick Rey.** 2018. "Exclusive Dealing and Vertical Integration in Interlocking Relationships." *Journal of Economic Theory* 177: 183–221.
- Organisation for Economic Co-operation and Development (OECD).** 2007. *Competition and Efficient Usage of Payment Cards 2006*. Paris: OECD.
- Organisation for Economic Co-operation and Development (OECD).** 2011. "Competition and Efficient Usage of Payment Cards." *OECD Journal: Competition Law and Policy* 11 (3): 7–99.
- Ordober, Janusz A., Garth Saloner, and Steven C. Salop.** 1990. "Equilibrium Vertical Foreclosure." *American Economic Review* 80 (1): 127–42.
- Rey, Patrick, and Thibaud Vergé.** 2019. "Secret Contracting in Multilateral Relations." Unpublished.
- Riordan, Michael H.** 2008. "Competitive Effects of Vertical Integration." In *Handbook of Antitrust Economics*, edited by Paolo Buccirossi, 145–82. Cambridge, MA: MIT Press.
- Riordan, Michael H.** 1998. "Anticompetitive Vertical Integration by a Dominant Firm." *American Economic Review* 88 (5): 1232–48.
- Riordan, Michael H., and Steven C. Salop.** 1995. "Evaluating Vertical Mergers: A Post-Chicago Approach." *Antitrust Law Journal* 63 (2): 513–68.
- Rogerson, William P.** 1992. "Contractual Solutions to the Hold-up Problem." *Review of Economic Studies* 59 (4): 777–93.
- Salinger, Michael A.** 1988. "Vertical Mergers and Market Foreclosure." *Quarterly Journal of Economics* 103 (2): 345–56.
- Schmitz, Patrick W.** 2002. "Simple Contracts, Renegotiation under Asymmetric Information, and the Hold-up Problem." *European Economic Review* 46 (1): 169–88.
- Segal, Ilya, and Michael D. Whinston.** 2011. "A Simple Status Quo That Ensures Participation with Application to Efficient Bargaining." *Theoretical Economics* 6 (1): 109–25.
- Snyder, Christopher M.** 2008. "Countervailing Power." In *The New Palgrave Dictionary of Economics*, edited by Steven N. Durlauf and Lawrence E. Blume, 1188–89. London: Palgrave Macmillan.
- Snyder, Christopher M.** 1996. "A Dynamic Theory of Countervailing Power." *RAND Journal of Economics* 27 (4): 747–69.
- Stigler, George J.** 1954. "The Economist Plays with Blocs." *American Economic Review* 44 (2): 7–14.
- US Department of Justice, and the Federal Trade Commission (DOJ and FTC).** 1996. *Statement of Antitrust Enforcement Policy in Health Care*. Washington, DC: US Department of Justice and the Federal Trade Commission.

- US Department of Justice, and the Federal Trade Commission (DOJ and FTC).** 2020. *Vertical Merger Guidelines*. Washington, DC: US Department of Justice and the Federal Trade Commission.
- Vickrey, William.** 1961. "Counterspeculation, Auction, and Competitive Sealed Tenders." *Journal of Finance* 16 (1): 8–37.
- Williams, Steven R.** 1987. "Efficient Performance in Two Agent Bargaining." *Journal of Economic Theory* 41 (1): 154–72.
- Williams, Steven R.** 1999. "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms." *Economic Theory* 14 (1): 155–80.
- Zhang, Xu, Puneet Manchanda, and Junhong Chu.** 2021. "'Meet Me Halfway': The Costs and Benefits of Bargaining." *Marketing Science*. <https://doi.org/10.1287/mksc.2021.1296>.