Incomplete information bargaining with applications to mergers, investment, and vertical integration*

Simon Loertscher[†] Leslie M. Marx[‡]

May 24, 2021

Abstract

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 agents, 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 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.

Keywords: bargaining power, countervailing power, vertical integration, investment
JEL Classification: D44, D82, L41

*We are grateful to three anonymous reviewers and two editors of this journal, Liran Einav and Sylvain Chassang, 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, 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, Australiasian Microeconomic Theory Seminars, 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 Annual Competition Law, Economics & Policy Conference in Pretoria, MACCI Workshop on Mergers and Antitrust, and 19th Annual IIOC 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 & Economics Eminent Research Scholar Grant, and Australian Research Council Discovery Project Grant DP200103574 is also gratefully acknowledged.

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

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

1 Introduction

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 agents 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 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 agents' 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

[1]The notion that vertical integration improves outcomes remains influential in antitrust. A case in point is the 2020 update of the U.S. DOJ and 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" (U.S. DOJ and FTC, 2020, pp. 2, 11).

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 of a market has been controversial since its inception,[2] it has popular appeal and has influenced antitrust policies and regulation.[3] Our paper thus 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 agents make noncontractible and nonobservable investments that improve their own type distributions, efficient incomplete information bargaining implies efficient 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 agents 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 agents' 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.[4] We generalize Williams' approach of maximizing an objective that assigns differential weights

[2]Galbraith (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).

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

[4]Gresik and Satterthwaite (1989) also allow for multiple suppliers and multiple buyers, but they restrict attention to identical value distributions, identical cost distributions, and all distributions with a common support. Our setup has similarities with the optimal auction setting of Myerson (1981), with the important difference that our setup has two-sided private information.

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.[5] In particular, inherent to the independent private values setting is the key economic tradeoff between rent extraction and social surplus. We defer further discussion of the literature to Section 8.

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 2 introduces the setup. In Section 3, we provide a model of incomplete information bargaining. Section 4 derives results pertaining to horizontal mergers, and Section 5 derives results for vertical integration. Section 6 analyzes investment incentives. We discuss extensions in Section 7 and related literature in Section 8. Section 9 concludes the paper. Formal mechanism design results and longer proofs are relegated to appendices.

2 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 $[\underline{c}, \overline{c}]$ 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 $[\underline{v}, \overline{v}]$ and density f_i that is positive on the interior of the support. The problem is trivial if $\overline{v} \leq \underline{c}$ because then it is never ex post efficient to have any trade. Therefore, we assume that $\overline{v} > \underline{c}$. We assume that $G_1, \dots, G_{n^S}$, $F_1, \dots, F_{n^B}$, $k_1^S, \dots, k_{n^S}^S$, and $k_1^B, \dots, k_{n^B}^B$ 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

[5] For experimental results consistent with the incomplete information bargaining, see Valley et al. (2002, Fig. 3.A). See Larsen (2021) on the first-best and second-best frontiers for wholesale used cars.

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 quasilinear preferences. The payoff of supplier j with type c_j when producing $q \in \{0, \ldots, k_j^S\}$ 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 $q \in \{0, \ldots, k_i^B\}$ units of the good and making the monetary payment m is $v_i q - m$.

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.[6] 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 impossible if and only if $\underline{v} < \overline{c}$.[7] We refer to the case with $\underline{v} < \overline{c}$ as the case with *overlapping supports* and the case with $\underline{v} \geq \overline{c}$ as the case of *nonoverlapping 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[8]

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

which we assume to be increasing.[9] 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. For $a \in [0, 1]$, we define the a-weighted virtual cost functions and the a-weighted virtual value functions by[10]

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

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.

[6]To 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 we assumed that the agent with all the bargaining power organized the exchange.

[7]We provide a derivation of this result in Appendix A.

[8]If $g_j(\underline{c}) = 0$, then define $\Gamma_j(\underline{c}) \equiv \lim_{c \downarrow \underline{c}} \Gamma_j(c)$. If $g_j(\overline{c}) = 0$, then $\Gamma_j(\overline{c}) = \infty$. Likewise, if $f_i(\overline{v}) = 0$, then define $\Phi_i(\overline{v}) \equiv \lim_{v \uparrow \overline{v}} \Phi_i(v)$. If $f_i(\underline{v}) = 0$, then $\Phi_i(\underline{v}) = -\infty$.

[9]The assumption of increasing virtual type functions can be relaxed through the use of "ironing."

[10]This departs from the more 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.

3 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.[11] For 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.

3.1 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 $Q_j^S : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to \{0, \ldots, k_j^S\}$ and $Q_i^B : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to \{0, \ldots, k_i^B\}$, 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 $\mathbf{M}^S : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to \mathbb{R}^{n^S}$ and $\mathbf{M}^B : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to \mathbb{R}^{n^B}$, maps types to the payments to the suppliers and the payments from the buyers.[12] 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 *incentive compatible* if it is in the best interest of every agent to report its type truthfully to the mechanism and is *individually rational* if each agent, for every possible type, is weakly better off participating in the mechanism than walking away, where we normalize the payoffs of not trading and of walking away—that is,

[11] For evaluating the merits of buyer power arguments, the standard oligopoly models of Cournot and Bertrand are "dead on arrival" because one side of the market—typically, buyers—is characterized by price-taking behavior and hence has no bargaining or market power. The assumption of efficiency embedded in generalized Nash bargaining preempts any social-surplus-increasing effects of changes in the bargaining weights because the outcome is efficient both before and after the change. As noted by Ausubel et al. (2002, p. 1934), the results of Myerson and Satterthwaite (1983) imply that the search for efficiency is "fruitless." Indeed, axiomatic bargaining approaches that stipulate efficient bargaining rule out transaction costs by assumption. In light of the Coase Theorem, which puts the question of transaction costs on center stage in economics, this limits the value of the approach.

[12] Any 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.

the value of the outside option—to zero.[13] A direct mechanism has *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 Appendix A.

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

$$U^S_{j;\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c}) \equiv M^S_j(\mathbf{v}, \mathbf{c}) - c_j Q^S_j(\mathbf{v}, \mathbf{c}),$$

and

$$U^B_{i;\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c}) \equiv v_i Q^B_i(\mathbf{v}, \mathbf{c}) - M^B_i(\mathbf{v}, \mathbf{c}).$$

The budget surplus generated by the mechanism is

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

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

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

To capture bargaining power, we endow the agents with bargaining weights $\mathbf{w} = (\mathbf{w}^S, \mathbf{w}^B)$, where $w^S_j \in [0, 1]$ is supplier j's bargaining weight and $w^B_i \in [0, 1]$ 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

$$W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c}) \equiv \sum_{i \in \mathcal{N}^B} w^B_i U^B_{i;\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c}) + \sum_{j \in \mathcal{N}^S} w^S_j U^S_{j;\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c}), \tag{1}$$

and assume that the market maximizes $\mathbb{E}_{\mathbf{v},\mathbf{c}}[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c})]$, subject to incentive compatibility, individual rationality, and the constraint of *no deficit*:

$$\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[R_{\mathbf{M}}(\mathbf{v}, \mathbf{c})\right] \geq 0. \tag{2}$$

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,

[13]In 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 Appendix A in terms of Bayesian incentive compatibility and interim individual rationality, one could also use the ex post versions of those constraints.

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 Appendix A). Thus, we have

$$\mathbb{E}_{\mathbf{v},\mathbf{c}}[M_j^S(\mathbf{v},\mathbf{c})] = \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\Gamma_j(c_j) Q_j^S(\mathbf{v},\mathbf{c})\right] + \hat{u}_j^S(\bar{c}) \tag{3}$$

and

$$\mathbb{E}_{\mathbf{v},\mathbf{c}}[M_i^B(\mathbf{v},\mathbf{c})] = \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\Phi_i(v_i) Q_i^B(\mathbf{v},\mathbf{c})\right] - \hat{u}_i^B(\underline{v}), \tag{4}$$

where $\hat{u}_j^S(c) \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}_{-j}}[U^S_{j;\mathbf{Q},\mathbf{M}}(\mathbf{v},\mathbf{c})]$ and $\hat{u}_i^B(v) \equiv \mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}}[U^B_{i;\mathbf{Q},\mathbf{M}}(\mathbf{v},\mathbf{c})]$.

It is possible that multiple agents have the maximum bargaining weight and that a mechanism exists that maximizes weighted welfare subject to incentive compatibility and individual rationality and satisfies the no-deficit constraint with slack. If this is the case, then the bargaining weights pin down the allocation rule but not the payments because the expected budget surplus can be allocated among any of the agents with the maximum bargaining weight as lump sum payments without affecting the value of the objective or the incentive constraints. Because a complete specification of the outcome of incomplete information bargaining needs to account for these eventualities, we assume that there are tie-breaking shares $(\boldsymbol{\eta}^S, \boldsymbol{\eta}^B) \in [0,1]^{n^S+n^B}$ 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 tie-breaking shares.[14]

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, $\mathbb{E}_{\mathbf{v},\mathbf{c}}[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(\mathbf{v},\mathbf{c})]$. Notice that, because we evaluate outcomes according to expected welfare $\mathbb{E}_{\mathbf{v},\mathbf{c}}[W_{\mathbf{Q}}(\mathbf{v},\mathbf{c})]$, 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 con-

[14]For example, one might apply equal sharing or distribute the surplus according to Nash bargaining weights. Like with Nash bargaining weights, the tie-breaking shares have no social surplus effects.

straint.[15]

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 large numbers of agents,[16] 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.

3.2 Allocation rule for incomplete information bargaining

The Lagrangian associated with maximizing expected weighted welfare (1) subject to the no-deficit constraint (2) can be written as $\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(\mathbf{v},\mathbf{c}) + \rho R_{\mathbf{M}}(\mathbf{v},\mathbf{c})\right]$, where ρ is the Lagrange multiplier on the no-deficit constraint.[17] Using (3) and (4), the Lagrangian can be rewritten as

$$\mathbb{E}_{\mathbf{v},\mathbf{c}}\Bigg[\sum_{i\in\mathcal{N}^B} w_i^B \overbrace{(v_i - \Phi_i(v_i))Q_i^B(\mathbf{v},\mathbf{c})}^{\text{buyer } i\text{'s surplus}} + \sum_{j\in\mathcal{N}^S} w_j^S \overbrace{(\Gamma_j(c_j) - c_j)Q_j^S(\mathbf{v},\mathbf{c})}^{\text{supplier } j\text{'s surplus}} \qquad (5)$$
$$\underbrace{+\rho\left(\sum_{i\in\mathcal{N}^B}\Phi_i(v_i)Q_i^B(\mathbf{v},\mathbf{c}) - \sum_{j\in\mathcal{N}^S}\Gamma_j(c_j)Q_j^S(\mathbf{v},\mathbf{c})\right)}_{\text{budget surplus}}\Bigg]$$

plus the term $\sum_{i\in\mathcal{N}^B}(w_i^B - \rho)\hat{u}_i^B(\underline{v}) + \sum_{j\in\mathcal{N}^S}(w_j^S - \rho)\hat{u}_j^S(\overline{c})$, which can be treated parametrically.[18] The buyer, supplier, and budget surpluses identified in (5) are the parts of the

[15]The first-best allocation rule is monotone and hence permits incentive compatible implementation. Without the no-deficit constraint, individual rationality is trivial to satisfy.

[16]Hatfield et al. (forth.) 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.

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

[18]Denoting by $\hat{q}_i^x(z)$ and $\hat{m}_i^x(z)$ the interim expected quantity and payment of agent i when its type is z for $x \in \{B,S\}$, which are formally defined in (13) in Appendix A.1, we have $\hat{u}_i^B(\underline{v}) = \hat{q}_i^B(\underline{v})\underline{v} - \hat{m}_i^B(\underline{v})$ and $\hat{u}_j^S(\overline{c}) = \hat{m}_j^S(\overline{c}) - \hat{q}_j^S(\overline{c})\overline{c}$. Consequently, no matter what the pointwise maximizer implies for $\hat{q}_i^B(\underline{v})$ and $\hat{q}_j^S(\overline{c})$, one can choose achieve any value for $\hat{u}_i^B(\underline{v})$ and $\hat{u}_j^S(\overline{c})$ by appropriately varying $\hat{m}_i^B(\underline{v})$ and $\hat{m}_j^S(\overline{c})$, respectively.

respective surpluses that vary with the allocation rule and exclude the fixed terms.

Given the Lagrange multiplier ρ, the allocation rule that maximizes (5) 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 $\boldsymbol{\Gamma}_j^a(c)$ denote the constant vector $(\Gamma_j^a(c), \ldots, \Gamma_j^a(c))$ with k_j^S elements and denote by $\boldsymbol{\Gamma}^{\mathbf{a}}(\mathbf{c}) \equiv (\boldsymbol{\Gamma}_j^{a_j}(c_j))_{j \in \mathcal{N}^S}$ the merged list of these weighted virtual costs. Analogously, let $\boldsymbol{\Phi}_i^a(v)$ denote the constant vector $(\Phi_i^a(v), \ldots, \Phi_i^a(v))$ with k_i^B elements and denote by $\boldsymbol{\Phi}^{\mathbf{a}}(\mathbf{v}) \equiv (\boldsymbol{\Phi}_i^{a_i}(v_i))_{i \in \mathcal{N}^B}$ 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 (5) is maximized when the quantity traded $q^*(\rho)$ is the largest element of $\{0, 1, \ldots, K\}$ such that the $q^*(\rho)$ lowest elements of $\boldsymbol{\Gamma}^{\mathbf{w}/\rho}(\mathbf{c})$ are less than or equal to the $q^*(\rho)$ greatest elements of $\boldsymbol{\Phi}^{\mathbf{w}/\rho}(\mathbf{v})$.[19] Defining $\Gamma^*(\rho)$ to be the $q^*(\rho)$-th lowest element of $\boldsymbol{\Gamma}^{\mathbf{w}/\rho}(\mathbf{c})$ and $\Phi^*(\rho)$ to be the $q^*(\rho)$-th highest element of $\boldsymbol{\Phi}^{\mathbf{w}/\rho}(\mathbf{v})$, 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 that trade from those that do not. We denote the set of *inframarginal* suppliers and buyers, respectively, by

$$\mathcal{N}_I^S(\rho) \equiv \{j \in \mathcal{N}^S \mid \Gamma_j^{w_j^S/\rho}(c_j) < \Gamma^*(\rho)\} \text{ and } \mathcal{N}_I^B(\rho) \equiv \{i \in \mathcal{N}^B \mid \Phi_i^{w_i^B/\rho}(v_i) > \Phi^*(\rho)\}.$$

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. Because weighted virtual values are weakly less than the associated values and weighted virtual costs are weakly greater than the associated costs, $\Phi^*(\rho)$ is weakly greater than the value of the marginal buyer who would trade under the first-best and $\Gamma^*(\rho)$ is weakly less than the cost of the marginal seller who would trade under the first-best.

With this in hand, we are in a position to describe the allocation rule that maximizes (5) 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 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.

[19]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.

It only remains to specify the solution value $\rho^{\mathbf{w}}$ 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), $\rho^{\mathbf{w}}$ 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 $\rho^{\mathbf{w}}$ is the shadow price of the no-deficit constraint, we have $\rho^{\mathbf{w}} \geq \max \mathbf{w}$.[20] Thus, we can define as $\rho^{\mathbf{w}}$ as the smallest value of ρ greater than or equal to $\max \mathbf{w}$ such that the "budget surplus" term in (5) 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* $\mathbf{w}$*,* $\mathbf{Q}^{\mathbf{w}}$*, is defined by*

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

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

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

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

Proof. See Appendix B.

An immediate implication of Lemma 1 is that the total quantity traded in incomplete information bargaining never exceeds the first-best level. This follows from the observation made above that $\Phi^*(\rho)$ is weakly larger than the value and $\Gamma^*(\rho)$ is weakly less than the cost of the marginal agents who would trade under the first-best.

3.3 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

[20] In addition, because a positive expected budget surplus is always possible given our assumption that $\overline{v} > \underline{c}$, the shadow price is finite.

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

$$\pi^{\mathbf{w}} \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{i\in\mathcal{N}^B}\Phi_i(v_i)Q_i^{\mathbf{w},B}(\mathbf{v},\mathbf{c}) - \sum_{j\in\mathcal{N}^S}\Gamma_j(c_j)Q_j^{\mathbf{w},S}(\mathbf{v},\mathbf{c})\right]. \tag{6}$$

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

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

Of course, if $\rho^{\mathbf{w}} > \max\mathbf{w}$, then the no-deficit constraint binds, implying that $\pi^{\mathbf{w}} = 0$ and that the question of how to allocate the budget surplus is moot. In contrast, if $\rho^{\mathbf{w}} = \max\mathbf{w}$, then $\pi^{\mathbf{w}} \geq 0$. In this case, weighted welfare is maximized when $\pi^{\mathbf{w}}$ 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

$$\hat{u}_j^S(\overline{c};\mathbf{w},\boldsymbol{\eta}) = \eta_j^S\pi^{\mathbf{w}} \quad\text{and}\quad \hat{u}_i^B(\underline{v};\mathbf{w},\boldsymbol{\eta}) = \eta_i^B\pi^{\mathbf{w}}, \tag{7}$$

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 tie-breaking shares $\boldsymbol{\eta}$ is then given by the expected buyer and supplier payoffs implied by the allocation rule $\mathbf{Q}^{\mathbf{w}}$ given in Lemma 1 and interim expected payoffs to agents' worst-off types given by (7). 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*

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

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

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

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 (18) in Appendix B).

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 Appendix B.

As we show in 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.

3.4 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 $W^* \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}}[W_{\mathbf{Q}^*}(\mathbf{v}, \mathbf{c})]$ the value of the planner's objective under the planner's optimal allocation rule, which we denote by $\mathbf{Q}^*$, and by $W^{\mathbf{w}} \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}}[W_{\mathbf{Q}^{\mathbf{w}}}(\mathbf{v}, \mathbf{c})]$ the value of the planner's objective under the allocation rule chosen by the market, denoted $\mathbf{Q}^{\mathbf{w}}$, 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 $w \in (0, 1]$.[21] Hence, for any $w \in (0, 1]$, we have

[21]To see this, note that $W^{(w,\dots,w)}_{\mathbf{Q},\mathbf{M}}(\mathbf{v}, \mathbf{c}) = w(W_{\mathbf{Q}}(\mathbf{v}, \mathbf{c}) - R_{\mathbf{M}}(\mathbf{v}, \mathbf{c}))$, 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 it has no effect on the allocation rule $\mathbf{Q}^{\mathbf{w}}$, which depends on w divided by that multiplier.

$W^* = W^{(w,\dots,w)}$.

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 $[\underline{v},\overline{v}]^{n^B} \times [\underline{c},\overline{c}]^{n^S}$. 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 when 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,

$$\left(K\text{-th lowest of } \{\Gamma_j^{w_j^S/\rho^{\mathbf{w}}}(\overline{c})\}_{j\in\mathcal{N}^S}\right) \leq \left(K\text{-th highest of } \{\Phi_i^{w_i^B/\rho^{\mathbf{w}}}(\underline{v})\}_{i\in\mathcal{N}^B}\right), \tag{8}$$

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.[22]

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

(i) the full-trade condition (8) *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 for all* $i \in \mathcal{N}^B$, $F_i = F$*.*

Proof. See Appendix B.

Proposition 3 provides conditions on bargaining weights and primitives such that the planner does not benefit from an equalization of bargaining weights. That said, there exist asymmetric bargaining weights such that the planner benefits from an equalization of bargaining weights, i.e., there exist $\mathbf{w}$ such that $W^{\mathbf{w}} < W^*$, unless $n^B = n^S = 1$, $\overline{c} \leq \Phi_1(\underline{v})$, and $\underline{v} \geq \Gamma_1(\overline{c})$.[23] This shows that, quite generally, equalization of bargaining power increases

[22]That 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 tie-breaking.

[23]These 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 $v < \overline{v}$ and $\Gamma_1(c) > c$

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 $\boldsymbol{\eta}_{\mathbf{w}}$ denote the tiebreaking shares that specify an equal division of any budget surplus among the agents with the maximum bargaining weight in $\mathbf{w}$,[24] and define the minimum and maximum total supplier payoffs by

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

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

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

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

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

the *Williams frontier.*

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 $\Delta \in [0,1]$, $\mathbf{w}_\Delta$ by $\mathbf{w}_\Delta^S \equiv (1-\Delta,\ldots,1-\Delta)$ and $\mathbf{w}_\Delta^B \equiv (\Delta,\ldots,\Delta)$, and letting $\tilde{u}_j^S(\Delta) \equiv u_j^S(\mathbf{w}_\Delta,\boldsymbol{\eta}_{\mathbf{w}_\Delta})$ and $\tilde{u}_i^B(\Delta) \equiv u_i^B(\mathbf{w}_\Delta,\boldsymbol{\eta}_{\mathbf{w}_\Delta})$, we have:

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

Proof. See Appendix B.

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

[24] That 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 $\max \mathbf{w}$.

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

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 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 $\overline{v} = \overline{c}$, because in that case first-best is not possible.[25]

A natural but somewhat elusive question is under what conditions on the primitives the first-best becomes possible for some bargaining weights. One might expect, for example, that the performance of the mechanism improves as the market becomes thicker in the sense that there are more ex ante identical buyers and sellers. While this has been shown to be true for the case with identical supports, single-unit traders, and identical distributions on each side of the market (see Gresik and Satterthwaite, 1989), it follows from the above that the first-best remains elusive with identical supports irrespective of assumptions on distributions, maximum demands, and capacities.[26] Away from the assumption of nonoverlapping supports, which guarantees that the first-best is possible with equal bargaining weights, it is therefore difficult to formulate general conditions under which the first-best is achievable.

That said, if there is a gap between the supports, that is, $\underline{v} > \overline{c}$, then the Williams frontier follows the first-best frontier for a range of bargaining weights that are sufficiently symmetric.[27] 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),

[25] With 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 multi-unit 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 $\overline{v} \leq \overline{c}$ also implies that the least efficient types on each side never trade.

[26] With nonidentical but overlapping supports, Williams (1999, Theorem 4 and Table 1) provides conditions under which the first-best is possible with single-unit traders and identical distributions on each side of the market. These conditions require that the number of agents on one side of the market is much larger than on the other. The nature of these conditions is captured in and around our Proposition 8 on vertical integration.

[27] Sufficiently symmetric bargaining weights are not required for the first-best. For example, in the bilateral trade setting in which the buyer has all the bargaining power, we get the first-best if $\underline{v} \geq \Gamma(\overline{c})$. This holds, for example, if the supplier's cost is uniformly distributed on $[0, 1]$ and the buyer's distribution has $\underline{v} \geq 2$.

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(b) is a case where reoptimizing cannot improve outcomes because at both endpoints of that segment, the mechanism is already first-best.

Another natural question is whether there is any empirical evidence in support of the Pareto efficiency of bargaining outcomes. The analysis in Larsen (2021), which is based on used-car transactions in a bilateral trade setup, shows that the Williams frontier is a reasonable albeit imperfect approximation to real-world bargaining outcomes. As mentioned in footnote 17, the incomplete information bargaining approach is amenable to the incorporation of additional transaction costs, which can be modeled by adding the constraint that the mechanism's revenue be no less than some positive number. We do not pursue this father here because of space constraints and because, without a specific application in mind, it is not clear what these additional transaction costs would be. Assumptions that imply that bargaining outcomes are Pareto efficient, subject to the constraints imposed by incomplete information, appear to be a good starting point as they are in line with those underlying the foundations of complete information bargaining.

(a) Overlapping supports (b) Nonoverlapping supports

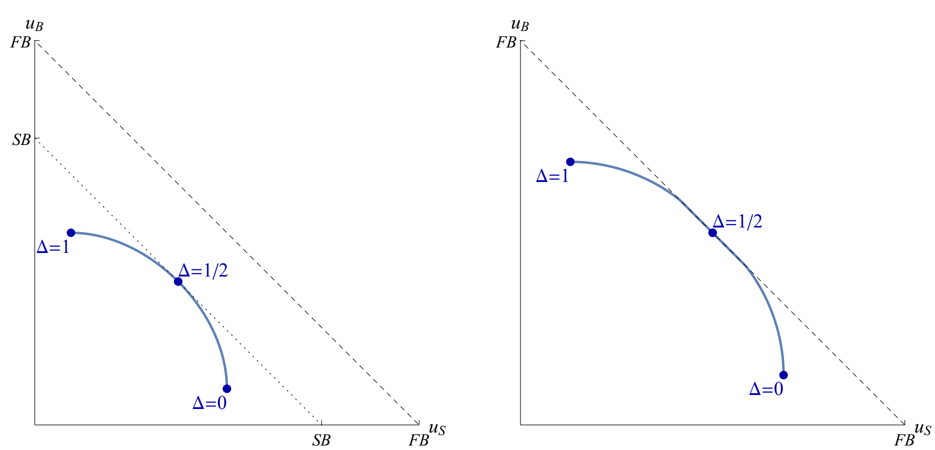

Figure 1: Williams frontier $\mathcal{F}$ for the case of 1 single-unit buyer and 1 single-unit supplier. FB denotes the first-best social surplus, and SB denotes the second-best social surplus. All types are uniformly distributed. Panel (a) assumes that $[\underline{c}, \overline{c}] = [\underline{v}, \overline{v}] = [0, 1]$. Panel (b) assumes that $[\underline{c}, \overline{c}] = [0, 1]$ and $[\underline{v}, \overline{v}] = [1.1, 2.1]$, in which case first-best and second-best total surplus are the same.

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 towards 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.[28] Considering bargaining between powerful insurance companies and doctors for the supply of health services, allowing physician joint ventures might equalize bargaining power and improve social surplus.[29]

An effect similar to that described in Corollary 1 arises in the partnership literature, where social surplus is increased by equalizing *ownership shares* rather than by equalizing bargaining weights. For example, as first observed by Cramton et al. (1987), when all agents draw their values independently from the same distribution, ex post efficient reallocation is possible if all agents have equal shares, and is impossible if one agent has full ownership. However, the paths through which these gains in social surplus are achieved, and the gains themselves, are different in the two approaches. In the partnership literature, the allocation rule is kept fixed at the ex post efficient one, but agents' ownership shares are allowed to change. The revenue of the mechanism increases as ownership shares (or, more generally, agents' worst-off types) become more similar, eventually permitting the first-best without running a deficit.[30] In contrast, in incomplete information bargaining with bargaining weights, the worst-off types of all agents are always the same (the lowest type for a buyer and the highest type for a supplier), and so is the revenue of the mechanism, which is zero. What changes as the bargaining weights change is the allocation rule, which transitions from, say, the buyer-optimal one via the second-best to the one that is optimal for the suppliers.

[28] See footnote 3.

[29] For example, the U.S. DOJ and FTC (1996) state: "The Agencies will not challenge, absent extraordinary circumstances, a non-exclusive 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.

[30] With identical distributions, equal shares imply equal worst-off types, which somewhat camouflages the point that the driving force for possibility is the equalization of worst-off types; see, for example, Che (2006) for a proof that with equal worst-off types, ex post efficiency is possible.

Moreover, because, for example with identical supports, the second-best is different from first-best in the model with bargaining weights, equalization of bargaining weights yields, in general, less social surplus than equalization of ownership shares (or worst-off types) in a partnership model.

4 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.

4.1 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 $k^S_{i,j}$, and for a merger of buyers i and j, we denote the merged entity's value distribution by $F_{i,j}$ and its capacity by $k^B_{i,j}$. 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 post-merger 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 $k^S_1 = k^S_2 = K^B$. 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, in line with Farrell and Shapiro (1990),[31] we assume that there are no synergies associated with a merger beyond the ability to rationalize production

[31] This is the approach also taken by, for example, Salant et al. (1983), Perry and Porter (1985), Waehrer (1999), Dalkir et al. (2000), and Loertscher and Marx (2019).

or consumption between the component firms.[32] 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 cost for K^S units from $F_{1,2}(v) = F_1(v)F_2(v)$.

4.2 Mergers that do not affect the bargaining weights

As we now show, in the case of a merger of "large" firms modeled as described above, a merger that does not alter bargaining weights or shares weakly reduces expected weighted welfare, and strictly so if the merging firms do not have the maximum bargaining weight. Of course, a merger that creates a merged entity that can produce at least some units at a cost below that of the pre-merger firms could increase expected weighted welfare, but below we provide conditions on the relation between the pre-merger and post-merger firms' type distributions and capacities such that a merger weakly reduces expected weighted welfare.

We say that a merger *does not alter bargaining weights or shares* if all nonmerging agents retain their pre-merger bargaining weights and shares in the post-merger market and the merging agents have the same bargaining weight in the pre-merger market, which is then inherited by the merged entity, and the share of the merged entity is equal to the sum of the shares of the merging agents.

Proposition 5. *A horizontal merger that does not alter bargaining weights or shares, and that involves suppliers i and j with $k_i^S = k_j^S = K^B$, weakly reduces expected weighted welfare (strictly if $w_{i,j} < \max \mathbf{w}$; and with equality for all agents if $\mathbf{w}_{-\{i,j\}} = \mathbf{0}$), and analogously for a merger of buyers i and j with $k_i^B = k_j^B = K^S$.*

Proof. See Appendix B.

A number of results follow from Proposition 5. It implies that a horizontal merger involving agents with maximum capacity reduces social surplus when all agents have the same bargaining weight and harms any nonmerging agent that has all the bargaining power. Further, the proposition implies that two maximum capacity agents on the same side of the market that are the only agents with bargaining power have no incentive to merge, which is an effect that is depicted in Figure 2 below. 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 multi-unit demand and supply, and bargaining power that is not restricted to be with the buyer. As in Loertscher and Marx (2019), a merger need not be profitable

[32] Additional merger-related synergies can be incorporated along the lines laid out in Loertscher and Marx (2019).

for the merging suppliers—when the buyers have sufficient bargaining power, the resulting more aggressive behavior against the merged entity because of its stronger type distribution can outweigh the benefits to the merging suppliers from the elimination of competition.

Building on Proposition 5, we see that while antitrust authorities seem inclined to look more favorably upon a supplier merger when there are powerful buyers,[33] in the absence of bargaining power effects, a merger of suppliers that face powerful buyers can actually be worse for social surplus than a merger of powerful suppliers:

Proposition 6. *Consider a horizontal merger of suppliers i and j with $k_i^S = k_j^S = K^B$ that does not alter bargaining weights. If the pre-merger market is:*
(i) efficient, then the merger decreases expected social surplus unless $w_i = w_j = \max \mathbf{w}$;
(ii) inefficient and $G_i = G_j$, then the merger decreases expected social surplus by more if the buyers have all the bargaining power than if the merging suppliers have all the bargaining power.

Proof. See Appendix B.

Under the conditions of Proposition 6, concerns regarding the welfare consequences of a merger are heightened when merging suppliers face a powerful buyer.

4.3 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.[34] It finds support in the empirical literature (Ho and Lee, 2017; Bhattacharyya and Nain, 2011; Decarolis and Rovigatti, forth.) 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, where shifts in bargaining power have no social surplus consequences, or rely on oligopoly models in which agents on one side of the market (typically buyers) are assumed to be price-takers and so have no bargaining power.

[33]"The Agencies consider the possibility that powerful buyers may constrain the ability of the merging parties to raise prices. ... However, the Agencies do not presume that the presence of powerful buyers alone forestalls adverse competitive effects flowing from the merger" (*U.S. Guidelines*, p. 27).

[34]As 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).

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. As an example, Figure 2(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.[35] Indeed, Figure 2(b) provides an example in which an equalization of bargaining power induces the first-best in the post-merger 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 post-merger market, the outcome is the first-best. In addition, in the example of Figure 2(b), when the pre-merger market is efficient, a merger causes that market to become inefficient unless the post-merger market has symmetric bargaining weights—the post-merger Williams frontier touches the first-best frontier only when $\hat{\Delta} = 1/2$.

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. This is consistent with the empirical analysis of Decarolis and Rovigatti (forth.) showing that consolidation among online advertising intermediaries has increased their buyer power, countervailing Google's significant market power in online search.[36]

We summarize with the following result:

Corollary 2. *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 agents 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.*

The above results, particularly Proposition 6 and Corollary 2, highlight the double-edged sword of allowing mergers of suppliers when buyers are powerful—there is the possibility that post-merger equalization of bargaining weights will increase social surplus, but if that does

[35] In the example of Figure 2(a), merger plus equalization of bargaining weights decreases buyer surplus by -0.028 and increases social surplus by 0.009, so for a competition authority to credit a countervailing power defense, it would need to place a weight of at least 75% on social surplus versus buyer surplus.

[36] For other empirical work related to countervailing power, see Ellison and Snyder (2010) and the cites therein.

(a) Social-surplus-enhancing equalization of bargaining weights

(b) First-best-achieving equalization of bargaining weights

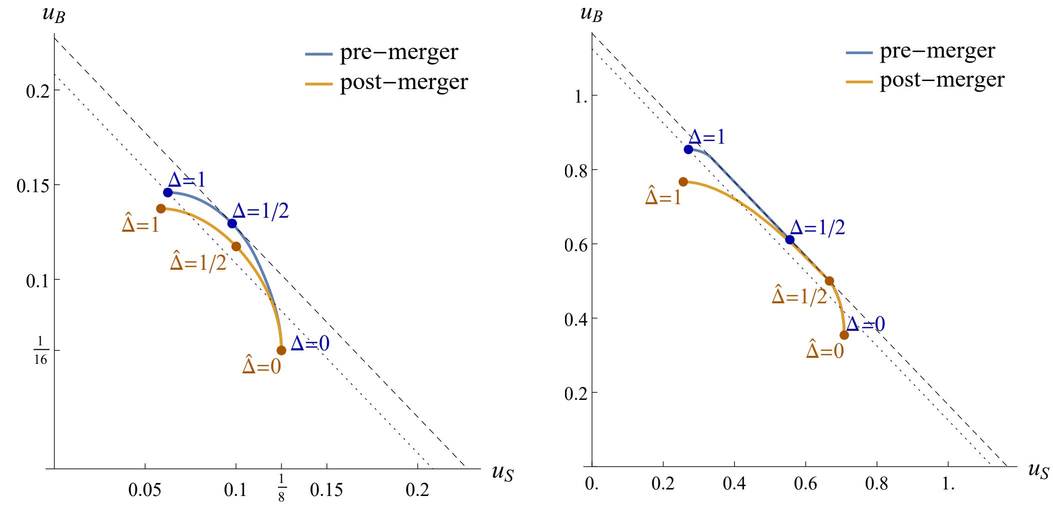

Figure 2: Williams frontier for the case of 1 pre-merger buyer and 2 symmetric pre-merger suppliers (blue), all with capacity of one, and the frontier following the merger of the two suppliers (orange). The pre-merger frontier is defined by $(\mathbf{w}^S, w^B) = (1,1,1)$ and $\Delta \in [0,1]$, and the post-merger frontier is defined by $(\hat{w}^S, \hat{w}^B) = (1,1)$ and $\hat{\Delta} \in [0,1]$. We assume equal tiebreaking shares among the agents 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 $[1,2]$.

not occur, then the merger harms social surplus, including possibly harming the buyers *and* the merging suppliers.

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.

Corollary 2 demonstrates that there is the possibility of defending a merger on grounds that it will equalize bargaining power because a merger that causes bargaining weights to shift in favor of the merging parties may improve expected social surplus despite the other adverse effects. Of course, a merger of, for example, suppliers that also induces an increase in suppliers' bargaining power is bad for the buyers for two reasons: competition among suppliers is reduced *and* the remaining suppliers have increased bargaining power. Thus, the review of a supplier merger based on a buyer-surplus standard would never be swayed by claims of equalization of bargaining power. In contrast, merger review based on a social-surplus standard may well be. In 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."

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 just 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 towards 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 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 Guidelines*, para. 67).[37] 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 an agent 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 trans-

[37]The 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 Guidelines*, para. 64).

actions always occur at the buyer's price.[38] In the case of a procurement auction, evidence consistent with buyer power and inconsistent with its absence includes: (i) the buyer uses procurement methods that result in suppliers other than the lowest-cost supplier winning, such as handicaps or preferences; (ii) the distribution of reserve prices is different across the markets if the buyer purchases in separate markets; (iii) one observes with positive probability ties in procurement outcomes and randomization over winners.[39] Analogous conditions apply for an analysis of supply power.

4.4 General conditions for weighted welfare reducing mergers

In general, a merger removes an independent agent and creates a merged entity that draws its type from a distribution that can differ from the distributions of the pre-merger firms. We now consider conditions on the merged entity's type distribution that are sufficient for the result, as in Proposition 5, that a merger reduces expected weighted welfare. To develop intuition, consider the case of a supplier merger in which, as in Proposition 5, the merged entity draws its constant marginal cost from the distribution of the minimum of the two merging suppliers' marginal costs. Then one can essentially transfer to the pre-merger market the allocation rule of any incentive compatible post-merger mechanism by replacing the type of a merged entity that combines suppliers i and j with $\min\{c_i, c_j\}$ and allocating the merged entity's quantity to the merging supplier i or j with the lower cost. Using threshold payments, the budget surplus, not accounting for fixed payments, is then greater in the pre-merger market because the competition between suppliers i and j reduces the threshold payments to those suppliers. This means that the post-merger incomplete information bargaining mechanism is feasible in the pre-merger market—indeed has strictly greater budget surplus not accounting for fixed payments. If it also gives (weakly) greater expected weighted welfare, then it follows by a form of revealed preference argument that expected weighted welfare under the (optimized) pre-merger mechanism is (weakly) greater

[38] This 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).)

[39] The 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.

than under the post-merger mechanism.[40]

We provide general conditions for this argument to apply in the following lemma. As stated in the lemma, we require that the post-merger distribution, $G_{i,j}$ in the case of a supplier merger, is equal to the distribution of some nondecreasing function of h the merging agents' types, c_i and c_j for a supplier merger. The remaining conditions ensure that one can rank the threshold payments of the merging agents relative to the threshold payment of the merged entity.

Lemma 3. *A merger of suppliers i and j that does not alter bargaining weights or shares weakly reduces expected weighted welfare if the merged entity's cost distribution $G_{i,j}$ and capacity $k^S_{i,j}$ are such that there exists continuous, nondecreasing function $h : [\underline{c}, \overline{c}]^2 \to [\underline{c}, \overline{c}]$ satisfying: (i) for all $z \in [\underline{c}, \overline{c}]$,*

$$\Pr_{c_i,c_j}\left(h(c_i, c_j) \leq z\right) = G_{i,j}(z),$$

(ii) $\min\{c_1, c_2\} \leq h(c_i, c_j)$, (iii) $k^S_i < k^S_{i,j} \Rightarrow c_j \leq h(c_i, c_j)$, and (iv) $k^S_j < k^S_{i,j} \Rightarrow c_i \leq h(c_i, c_j)$; and analogously for a merger of buyers.

Proof. See Appendix B.

Condition (i) in Lemma 3 is sufficient to permit the construction of a mechanism in the pre-merger market that has the same interim expected allocations for the nonmerging firms as the post-merger incomplete information bargaining mechanism. Then, as long as $h(c_i, c_j) \geq c_i$ whenever agent i trades, the sum of the pre-merger agents' threshold payments is no greater than the threshold payment of the merged entity in the post-merger market. Conditions (ii)–(iv) guarantee this. Thus, under the conditions of Lemma 3, one can construct a pre-merger mechanism that mimics the post-merger allocation and payments, but with weakly lower payments to merging suppliers (weakly higher payments to merging buyers), resulting in an incentive compatible, individually rational, no-deficit mechanism for the pre-merger market that has the same or greater expected weighted welfare. Optimizing that mechanism for the pre-merger market only reinforces the result that expected weighted welfare is greater pre merger than post merger.

[40]If the merging suppliers do not have the maximum bargaining weight, then pre-merger expected weighted welfare can be increased by distributing the savings from reduced payments to the merging suppliers to agents with higher bargaining weights, yielding the result that expected weighted welfare is greater pre-merger. If the merging suppliers have the maximum bargaining weight and all other agents have lower bargaining weights, then one can achieve the same expected weighted welfare in the pre-merger market, and potentially greater expected weighted welfare once the mechanism is optimized for the pre-merger market. If the merging suppliers have the maximum bargaining weight and all other agents have a bargaining weight of zero, then no further optimization is possible, and so the merger has no effect on expected weighted welfare.

5 Vertical integration

We now analyze vertical integration between a buyer and a supplier. We begin by highlighting the stark results that are available in specific settings and then analyze comparative statics for broader settings.

Throughout this section, to simplify the analysis, we consider settings in which the short side of the market has only one agent before integration, that is, $\min\{n^B, n^S\} = 1$, and all agents 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 can only sell, provided that it trades.[41] 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.[42]

5.1 Efficiency rationales for vertical integration

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

Proposition 7. *If supports overlap, then vertical integration increases social surplus when the number of agents is sufficiently small, regardless of bargaining weights.*

As reflected in Proposition 7, 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. This leaves the question of whether vertical integration is always neutral. The

[41] This 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.

[42] This assumption can be rationalized, for example, on the grounds that integration slackens the individual rationality constraints within the integrated firm. For an examination of incentive problems within an integrated firm, see Legros and Newman (2013).

following proposition shows that the answer is negative—in the post-integration market, the integrated firm sources internally for some type realizations when an outside supplier has a lower cost.

Proposition 8. *Assume a one-to-many pre-integration market with* $n^B = 1 < n^S$ *and* $\underline{v} > \underline{c}$ *and symmetric bargaining weights pre-integration. If* $\overline{c} < \overline{v}$*, then vertical integration cannot increase social surplus for* $\underline{v}$ *sufficiently large, regardless of post-integration bargaining weights. Similarly, if* $\underline{v} > \underline{c}$ *and* $G_j = G$ *for all* $j \in \mathcal{N}^S$*, then vertical integration cannot increase social surplus for* n^S *sufficiently large, regardless of post-integration bargaining weights.*

Proof. See Appendix B.

Proposition 8 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 $\overline{c} < \overline{v}$ and $\underline{v}$ sufficiently large, then vertical integration cannot increase social surplus simply because the first-best is already achieved without integration. 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 8 follows by setting $\underline{v} = \overline{c}$, 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 8 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).[43] With little competition, rival suppliers have large markups and react to vertical integration by reducing markups *and* quantities. In contrast, with in highly competitive markets, price is already close to marginal costs. Consequently, 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.

[43] Riordan (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.

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 $\overline{c}$ 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 $\overline{c} < \overline{v}$, then vertical integration cannot increase social surplus for n^B sufficiently large.

Propositions 7 and 8 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 analogue to 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 agents. 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.[44]

The results of Propositions 7 and 8 are robust in that they do not depend on specific assumptions about distributions or beliefs of agents. 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 vertical contracting and integration, whose predictions rely on assumptions about beliefs, feasible contracts, and/or market power.[45] Maybe more importantly, 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 et al. (2021), this is a matter of substance rather than taste.

Of course, our results do rely, inevitably, on support assumptions. Indeed, for $\underline{v} \geq \overline{c}$ and equal bargaining weights pre-integration, vertical integration cannot possibly increase social surplus. This suggests that, as the overlap of supports becomes smaller, the social surplus gains from vertical integration may decrease as well. To formalize and substantiate this notion, fix $[\underline{c}, \overline{c}] = [0, 1]$ and $\overline{v} = 1$, and define the *maximum gain from vertical integration*

[44]This 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.

[45]For 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).

associated with $\underline{v} \in [0,1)$ by $\mathcal{G}(\underline{v}) \equiv (W^{FB}(\underline{v}) - W^{SB}(\underline{v}))/W^{FB}(\underline{v})$, where $W^{FB}(\underline{v})$ and $W^{SB}(\underline{v})$ denote first-best and second-best social surplus, respectively, as a function of $\underline{v}$. The amount $\mathcal{G}(\underline{v})$ provides only an upper bound for the gain from vertical integration because vertical integration does not necessarily make the first-best possible when it is not possible absent vertical integration. Then we have:[46]

Proposition 9. *Assuming $n^B = 1$ and symmetric suppliers, $\mathcal{G}(\underline{v})$ decreases in $\underline{v}$ whenever $\mathcal{G}(\underline{v}) > 0$.*

Proof. See Appendix B.

Proposition 9 provides a monotonicity result relating differences in supports to the maximum gain from vertical integration. Reduced overlap of the supports reduces the maximum gain from vertical integration. Intuitively, the social benefit from vertical integration is reduced when gains from trade are more certain because then market-based transactions between nonintegrated firms work better.

5.2 Comparative statics for vertical integration

In this section, we begin by considering the possibility, captured by Proposition 7 that with overlapping supports, the social surplus effects of vertical integration depend, in general, on the number of firms. Specifically, consider the case of one buyer and multiple suppliers, each of which draws it cost from the same distribution. We know from Williams (1999) that the first-best is possible if $\overline{v} > \overline{c}$ and n^S is large enough.[47] Because vertical integration induces the buyer's willingness to pay to be $y = \min\{v, c\}$, the support of y is $[\min\{\underline{v}, \underline{c}\}, \overline{c}]$. The results of Williams (1999) for this case imply that the first-best is not possible. Hence, vertical integration is socially harmful whenever n^S and the supports are such that the first-best is possible without vertical integration. When $\underline{c} = \underline{v}$ and $\overline{c} = \overline{v}$, the first-best is not possible absent vertical integration, nor with vertical integration if $n^S > 1$ (see, e.g., Williams, 1999).

So the question arises whether vertical integration is more or less likely to be socially harmful if n^S is larger. The intuition and insights from oligopoly models, discussed above, together with the point that the double markup under linear pricing in those models corresponds to having to pay information rents to both the buyer and the suppliers in incomplete information bargaining, suggests that vertical integration is more likely to be harmful the larger is n^S. However, establishing this in general is challenging and beyond the scope of

[46]Not surprisingly, a result analogous to Proposition 9 obtains for the case of one single-unit supplier if one fixes the buyers' support at $[0,1]$ and $\underline{c} = 0$ and varies $\overline{c} \in (0,1]$. In that case, the maximum gain from vertical integration is decreasing in $\overline{c}$.

[47]See also Makowski and Mezzetti (1993).

the present paper. That said, Figure 3(a) illustrates a case in which the oligopoly intuition carries over to our setup. In particular, in a one-to-many market with overlapping supports, vertical integration eliminates a double-markup (of information rents), but it also makes the outside market less competitive, and possibly less efficient, and affects the virtual type function of the integrated firm. As the number of nonintegrated suppliers grows large, the probability that the vertically integrated firm sources internally goes to zero, and the outside market is close to efficient (see Appendix E for details). Because all effects become small, it is hard to prove general results analytically. As shown in Figure 3(a), for the case of uniformly distributed types on $[0, 1]$, the change in social surplus due to vertical integration is nonmonotone in the number of outside suppliers and, in the limit, approaches zero from below.

(a) Social surplus effect of vertical integration given $[\underline{c}, \overline{c}] = [\underline{v}, \overline{v}]$ as a function of n^S (illustration of Proposition 7)

(b) Social surplus effect of vertical integration given $[\underline{c}, \overline{c}] = [0, 1]$ and $\overline{v} = 1.2$ as a function of $\underline{v}$ (illustration of Proposition 8)

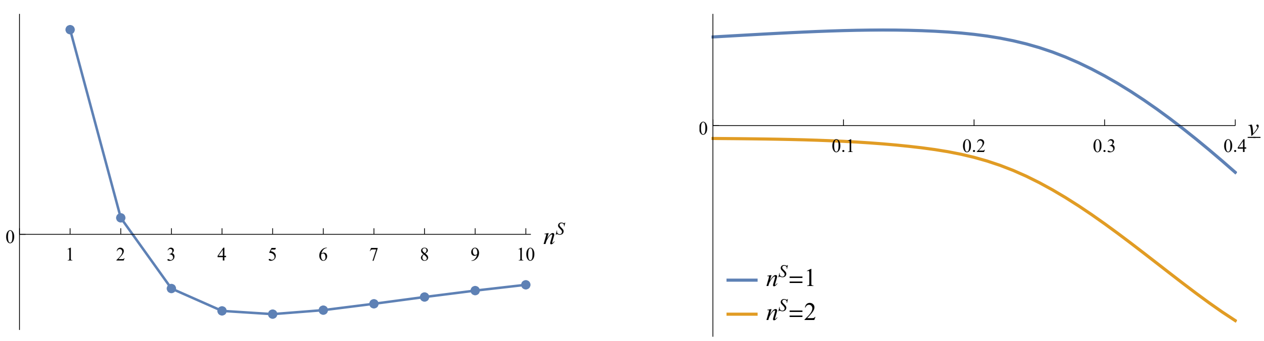

Figure 3: Change in expected social surplus as a result of vertical integration in a market with one buyer and multiple suppliers with single-unit demand and supply, where n^S is the number of independent suppliers after vertical integration (i.e., the pre-integration market has $n^S + 1$ suppliers). Pre-integration suppliers' costs are uniformly distributed on $[\underline{c}, \overline{c}] = [0, 1]$, the pre-integration buyer's value is uniformly distributed on $[\underline{v}, \overline{v}]$, and bargaining weights are symmetric. In panel (a), $[\underline{v}, \overline{v}] = [0, 1]$, and in panel (b), $\overline{v} = 1.2$ with $\underline{v}$ varying as indicated.

We provide comparative statics related to Proposition 8 in Figure 3(b), where we illustrate that given $\overline{c} < \overline{v}$, as $\underline{v}$ increases, eventually vertical integration decreases expected social surplus (again, see Appendix E for details). While Proposition 8 is straightforward to prove by taking the case of $\underline{v} = \overline{c}$, Figure 3 provides examples in which having as few as two independent suppliers in the post-integration market is sufficient for vertical integration to reduce social surplus even when $\underline{v} = \underline{c}$. This emphasizes the salient possibility of anticompetitive vertical integration in a variety of settings.

Interesting and challenging (and still open) issues arise with vertical integration in many-to-many settings. As discussed in and around footnote 41, in that case, the integrated firm may be a buyer or a supplier vis-à-vis the outside firms in the post-integration market, or not trade at all.

6 Investment

We now extend the model to allow for investment.

6.1 Setup and results

Investment incentives feature prominently, and at times controversially, in concurrent policy debates,[48] 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 extend our model by adding investment as an action taken by each agent prior to the realization of private information, where investment improves an agent'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 and, as a consequence, obtains hold-up and inefficient investment. In contrast, in our setting, incomplete information protects agents from hold-up, and investments are efficient if and, under additional assumptions, only if bargaining is efficient.

We extend the model to allow the suppliers and buyers to improve (or more generally change) their type distributions by investing. Investments do not change the supports of the distributions. 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.[49] 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 agents are not affected by

[48]For example, related to the 2017 Dow-DuPont merger, the U.S. 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 non-large 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).

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

actual investments. We assume that the buyers and suppliers first simultaneously make their investments and then bargaining takes place.

We first consider the planner's problem of determining investments when the allocation rule is first-best. 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 $W^{FB}(\mathbf{v},\mathbf{c}) \equiv \sum_{i\in\mathcal{N}^B} v_i Q_i^{FB,B}(\mathbf{v},\mathbf{c}) - \sum_{j\in\mathcal{N}^S} c_j Q_j^{FB,S}(\mathbf{v},\mathbf{c})$. We let $\overline{\mathbf{e}}$ denote first-best investments, which are a solution to the planner's first-best investment problem, given by $\max_{\mathbf{e}} \mathbb{E}_{\mathbf{v},\mathbf{c}|\mathbf{e}} \left[W^{FB}(\mathbf{v},\mathbf{c})\right] - \sum_{i\in\mathcal{N}^B} \Psi_i^B(e_i^B) - \sum_{j\in\mathcal{N}^S} \Psi_j^S(e_j^S)$.

Now consider the agents' 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 agent. Thus, for the case considered here in which the first-best is possible, we can, without loss of generality, focus on expected utilities for the Vickrey-Clarke-Groves (VCG) mechanism. Given a type realization $(\mathbf{v},\mathbf{c})$, supplier i's VCG payoff is $W^{FB}(\mathbf{v},\mathbf{c}) - W^{FB}(\mathbf{v},\overline{c},\mathbf{c}_{-i})$, plus possibly a constant. Likewise, the buyer i's payoff is $W^{FB}(\mathbf{v},\mathbf{c}) - W^{FB}(\underline{v},\mathbf{v}_{-i},\mathbf{c})$, plus possibly a constant.

Taking expectations over $(\mathbf{v},\mathbf{c})$, and noticing that $W^{FB}(\mathbf{v},\overline{c},\mathbf{c}_{-j})$ 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 agents choose investments $\overline{\mathbf{e}}_{-j}$, is $\max_{e_j^S} \mathbb{E}_{\mathbf{v},\mathbf{c}|e_j^S,\overline{\mathbf{e}}_{-j}} \left[W^{FB}(\mathbf{v},\mathbf{c})\right] - \Psi_j^S(e_j^S)$. An analogous optimization problem applies to buyer i's choice of e_i^B, noting that $W^{FB}(\underline{v},\mathbf{v}_{-i},\mathbf{c})$ 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 $\overline{\mathbf{e}}$ is a Nash equilibrium if incomplete information bargaining permits the first-best. This proves the first part of Proposition 10 below.

Under additional conditions, the converse is also true, that is, $\overline{\mathbf{e}}$ being a Nash equilibrium outcome in the game in which agents' 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 $i \in \mathcal{N}^B$,

$$\Psi_j^{S\prime}(0) = \Psi_i^{B\prime}(0) = 0, \text{ and for all } e > 0,\ \Psi_j^{S\prime}(e), \Psi_i^{B\prime}(e) > 0 \text{ and } \Psi_j^{S\prime\prime}(e), \Psi_i^{B\prime\prime}(e) > 0; \quad (9)$$

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

$$\frac{\partial G_j(c;e)}{\partial e} > 0 \quad \text{and} \quad \frac{\partial F_i(v;e)}{\partial e} < 0; \tag{10}$$

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

$$G_j(c; \overline{e}_j^S) \equiv G(c), \tag{11}$$

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

$$F_i(v; \overline{e}_i^B) \equiv F(v). \tag{12}$$

Conditions (9)–(10) imply that the first-best investments $\overline{\mathbf{e}}$ 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.[50]

Proposition 10. *First-best investments are a Nash equilibrium outcome of the simultaneous investment game if incomplete information bargaining is efficient. Conversely, assuming that* (9)–(10) *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 Appendix B.

As shown in Proposition 10, when incomplete information bargaining is efficient, the agents' Nash equilibrium investment choices are first-best investments. Because private information protects agents from hold-up,[51] efficient incomplete information bargaining implies efficient investments.[52] Intuitively, given that the allocation rule is efficient, each agent is the

[50]Proposition 10 connects to the equivalence result of Hatfield et al. (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.

[51]Lauermann (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.

[52]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

residual claimant to the surplus that its investment generates. Anticipating that this will be the case once types are realized, each agent's incentives are also aligned with the planner's at the investment stage because each agent'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 10 with Corollary 2 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 2, Proposition 10 shows that it can also improve investments to the first-best level. Proposition 10 thus provides an additional channel—investments—through which change in bargaining power can increase social surplus.

While Proposition 10 focuses on investments that improve agents' own types, the first part of Proposition 10 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.[53] Our result does not hold if, for example, investment generates externalities, e.g., if there are technology spillovers across suppliers or if investment increases the buyer's value regardless of its trading partner.

6.2 Investment and vertical integration

Using Proposition 10, we can connect investment with vertical integration. We assume that vertical integration does not affect the cost of investment for the integrated firm, so if buyer i and supplier j integrate and invest $e_i^B + e_j^S$, the cost of investment is $\Psi_i^B(e_i^B) + \Psi_j^S(e_j^S)$. With one buyer and one supplier in the pre-integration market and overlapping supports, incomplete information bargaining is inefficient, which under conditions (9) and (10), implies that equilibrium investments are inefficient. But, by assumption, the allocation is efficient after vertical integration, which by Proposition 10 implies that investments are efficient after vertical integration. Thus, with overlapping supports, vertical integration promotes efficient investment insofar as there is an equilibrium with efficient investments after integration but not before. In contrast, with, say, one buyer and two or more symmetric suppliers and nonoverlapping supports, incomplete information bargaining is efficient for some bargaining weights, including symmetric ones, without vertical integration, which implies that

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."

[53]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.

investments are efficient without vertical integration. But following vertical integration, incomplete information bargaining is inefficient, and so, under (9) and (10), and investments are no longer efficient. In this case, vertical integration disrupts efficient investment insofar as there is no equilibrium with efficient investments after integration whereas there was one before integration.

Corollary 3. *Assuming that* (9) *and* (10) *hold, for a one-to-one pre-integration market with overlapping supports, vertical integration promotes efficient investment; but for a one-to-many pre-integration market with nonoverlapping supports, if at least one of (ii)–(iv) above holds, then vertical integration disrupts efficient investment if bargaining is efficient prior to vertical integration (which occurs, for example, with symmetric bargaining weights).*

6.3 Comparative statics for investment

We now analyze how equilibrium investments are affected by bargaining power and by the extent to which the supports of the value and cost distributions overlap. To analyze investment effects, we parameterize the agents' 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.

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. Specifically, 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 agent's investment e has cost $e^2/2$. Relegating the details to Appendix E, we illustrate the effects of bargaining power and the distributional supports on equilibrium investment in Figure 4.

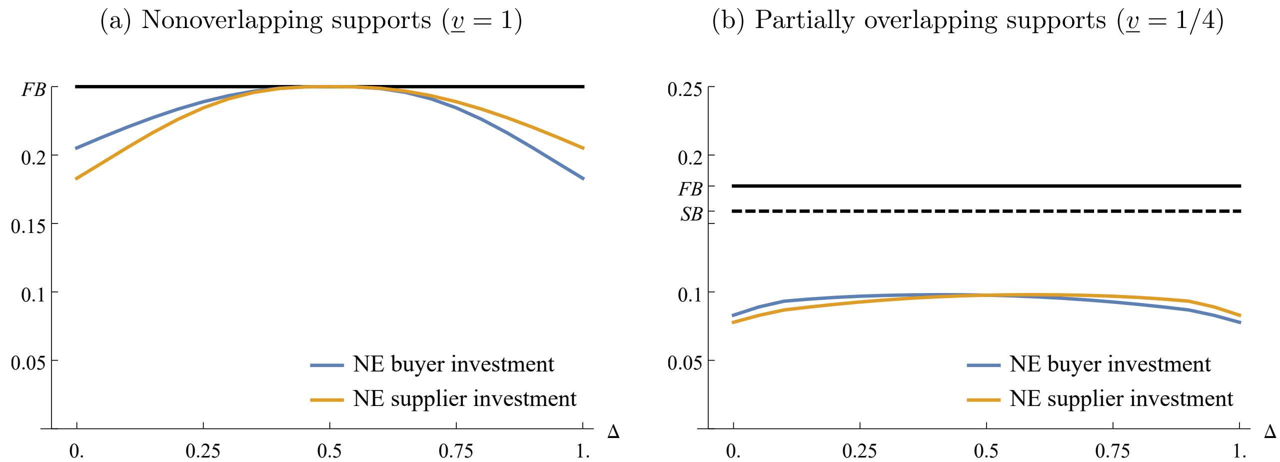

Figure 4: Nash equilibrium investments with bargaining weights $(w^S, w^B) = (1-\Delta, \Delta)$ for buyer distributions with varying supports. Assumes the linear virtual type setup for bilateral trade with $F(v) = 1 - (1 + \underline{v} - v)^{1.25-e_B}$, where $e_B \in [0, 1.25)$ is the buyer's investment, 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.

As shown in Figure 4, each agent's equilibrium investment is maximized away from extreme bargaining weights. This points to an additional benefit of the equalization of bargaining weights; namely, it has the potential to improve the efficiency of investment, in some cases to the first-best.

7 Extensions and discussion

The model presented here can accommodate a range of extensions. We first provide an extension to incorporate price-taking downstream consumers. Then we briefly discuss extensions that allow variation in agents' outside options and that allow buyers to have preferences over suppliers, in which case bargaining externalities arise naturally, with details contained in Appendix C. Also in Appendix C, we show 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. Finally, we briefly discuss extensive-form and axiomatic foundations for incomplete information bargaining, with details contained in Appendix D.

7.1 Downstream consumers

Thus far, our model does not include a mass of downstream consumers beyond the buyers that participate in incomplete information bargaining. Because competition authorities commonly put weight on the welfare of final consumers,[54] it seems important to 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. (Whether it is the only supplier in that market does not matter for these purposes; to rule out externalities in the bargaining mechanism, what is important is that none of the other agents participating in the mechanism is active in this market.) The input that the buyers procure can be interpreted as reducing the marginal cost of production in that market or (by and large) equivalently as improving the quality of the product.

A natural and tractable model in which the buyer's value and consumer surplus move in the same direction is one in which the buyer's private information relates to the size of the market. Specifically, let $D_i(p)$ be the decreasing demand function in market i when the mass of consumers is equal to 1, normalize units so that $D_i(0) = 1$, and assume that $D_i(\overline{p}_i) = 0$ for some finite price $\overline{p}_i$. If we assume that the buyer is a retailer that has constant marginal costs of production, which we normalize to zero without loss of generality, then it is optimal for the retailer to set a uniform price p_i^*, assuming each consumer is privately informed about its value and has single-unit demand, where $p_i^* \in \arg\max_p pD_i(p)$. Accordingly, consumer surplus in the market with a mass 1 of consumers and a good of quality 1 is $CS_i^* = \int_{p_i^*}^{\overline{p}_i} D_i(p)dp - D_i(p_i^*)p_i^*$, and if the mass of consumers in the market is ω, demand at price p is $\omega D_i(p)$, implying that consumer surplus is ωCS_i^*, while the retailer's profit is $\omega p_i^* D_i(p_i^*)$. Letting $\gamma_i = \int_{p_i^*}^{\overline{p}_i} D_i(p)dp/(p_i^* D_i(p_i^*)) - 1 > 0$, where the inequality follows because D_i is decreasing, we have $CS_i^* = \gamma_i p_i^* D_i(p_i^*)$. In words, γ_i captures how much of retailer i's profit is "passed through" to downstream consumers in market i.

If the quality of the retailer's good improves by the commonly known parameter $\delta > 0$ when the retailer buys the input, then consumer surplus increases by $\omega\delta CS_i^*$ while the retailer's willingness to pay for the quality increment is $v_i = \omega\delta p_i^* D(p_i^*)$ if the mass of consumers is ω. Accordingly, the increase in consumer surplus in market i if the retailer obtains the input is $\gamma_i v_i$. In other words, the larger is the buyer's willingness to pay, the larger is the consumer surplus effect of this buyer obtaining the input.[55]

[54] 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.

[55] Alternatively, and largely equivalently, one could think of the input as decreasing the constant marginal cost of production from $c > 0$ to $c - \delta$ with $\delta < c$. As with the quality improvements, $\Delta CS^* = \gamma\Delta\Pi^*$ for

With downstream consumers, both the social planner who aims at maximizing equally weight social surplus and an 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

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

plus $\sum_{i\in\mathcal{N}^B}(1-\rho)\hat{u}_i^B(\underline{v}) + \sum_{j\in\mathcal{N}^S}(1-\rho)\hat{u}_j^S(\overline{c})$. This is the same as (5) with $\mathbf{w} = \mathbf{1}$ and the addition of the γ_i's. By contrast, the Lagrangian for the authority with a consumer surplus objective is simply

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

plus $\sum_{i\in\mathcal{N}^B}(1-\rho)\hat{u}_i^B(\underline{v}) + \sum_{j\in\mathcal{N}^S}(1-\rho)\hat{u}_j^S(\overline{c})$.

If the "pass-through" of 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 planner nor the authority will discriminate among buyers in the weights that they attach to them. In this case, any discrimination in their mechanisms will be due to the no-deficit constraint being binding, which implies that there is some discrimination based on virtual values, provided $F_i \neq F_h$ for some $i \neq h$ with $i, h \in \mathcal{N}^B$. Simply having a larger expected market is not a reason for discrimination in this case because 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 authority would discriminate across downstream markets.[56] This is true for the social planner even if the no-deficit constraint does not bind.[57]

---

some $\gamma > 0$.

[56] A simple specification that allows for the γ_i's to vary across market is to assume that for all i, for $D_i(p) = (1-p)^{a_i}$ for $p \in [0,1]$ with $a_i > 0$. This implies $p_i^* = \frac{1}{1+a_i}$ and hence $D_i(p_i^*) = \left(\frac{a_i}{1+a_i}\right)^{a_i}$. At $\omega = 1$, $v_i = \frac{1}{a_i}\left(\frac{a_i}{1+a_i}\right)^{1+a_i}$ while $CS_i^* = \int_{\frac{1}{1+a_i}}^1 (1-p)^a dp - v_i = \frac{1}{1+a_i}v_i$. Thus, $\gamma_i = \frac{1}{1+a_i}$.

[57] For the authority with a consumer surplus standard, the no-deficit constraint always binds because it does not put any weight on the buyers' and the suppliers' welfare.

7.2 Bargaining breakdown

A pervasive feature of real-world bargaining is that negotiations often break down.[58] Anecdotal examples range from the U.S. government shut down, to the British coal miners' and the U.S. air traffic controllers' strikes in the 1980s, to failures to form coalition governments in countries with proportional representation. Providing systematic evidence of bargaining breakdown, Backus et al. (2020) analyze a data set covering 25 million observations of bilateral negotiations on eBay and find a breakdown probability of roughly 55 percent. More generally, when firms bargain over essential inputs, such as medical equipment for hospitals or computer chips for manufacturers, bargaining breakdown will typically not mean that the firms stop trading with each other, but rather that the latest, quality-improved version of the input is not traded.

Because with incomplete information, bargaining breakdown occurs on the equilibrium path, one can use observed bargaining breakdown frequencies as a moment to match in empirical research rather than as an error.[59] While specifics will, of course, depend on the available data and on the econometric approach, we provide an illustration in Appendix E.

7.3 Variation in agents' outside options

Our model is amenable to introducing variation in agents' outside options, which occupy center stage in complete information bargaining. In the incomplete information setting, outside options can affect an agent's cost of participating in the mechanism independently of whether the agent trades and can affect its value or cost distribution by shifting its support.[60] We provide a brief discussion here, with details contained in Appendix C.

The comparative statics with respect to increasing an agent's participation cost are intuitive and largely the same as in models with complete information because it increases the agent's share of the surplus that is created; in contrast to complete information models, it may decrease expected social surplus because of distortions to the allocation rule required

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

[59] With incomplete information, bilateral bargaining can break down on the equilibrium path for three reasons. First, it may be that the buyer's value is below the supplier's cost, but because of private information, the two parties do not know this before they sit down at the negotiating table, so bargaining begins but then breaks down. Second, with unequal bargaining power, incentives for rent extraction may lead more powerful agents to impose sufficiently aggressive thresholds for trade that breakdown results. Third, because of impossibility theorems, even if the buyer's value exceeds the supplier's cost, the constraints imposed by incentive compatibility, individual rationality, and no deficit may prevent ex post efficient trade from taking place.

[60] Of course, to the extent that outside options affect bargaining weights, the comparative statics are those discussed above.

to cover larger outside options.

The effects of changing an agent's *production-relevant* outside option are more nuanced. For example, as a supplier's outside option improves, the support of its cost distribution shifts upwards by the amount of the improvement, with the result that higher costs become more likely. Hence, the supplier will tend to be less likely to trade. However, under the assumption of monotone hazard rates, this effect is partly (but not completely) offset because, for a given cost realization, the supplier's weighted virtual cost is lower than before the increase in the outside option. This implies that, ex post, given the same cost realization, the supplier is treated more favorably after the outside option increases. This is in line with intuition gleaned from complete information models. But from an ex ante perspective, the increase in the outside option reduces the supplier's expected payoff from incomplete information bargaining because overall it makes the supplier less likely to trade and thereby decreases the supplier's ex ante expected payoff. Moreover, as a supplier's cost distribution worsens, the revenue constraint faced by the mechanism becomes tighter, which further tends to worsen the agent's bargaining outcome.

7.4 Implementation

In many cases, economists have achieved greater comfort with models of price-formation processes when the literature has shown that there exists a noncooperative game that, at least under some assumptions, has an equilibrium outcome that is the same as the outcome delivered by the model under consideration. Indeed, this comfort often extends well beyond the narrow confines of the foundational game. For example, the existence of microfoundations are regularly invoked to support empirical estimation of a model even when the data-generation process does not conform to the extensive-form game providing the microfoundation.[61]

In light of this, it is perhaps useful to note that, as mentioned above and discussed in Appendix D.1, for the case of one supplier, one buyer, and uniformly distributed types, the k-double auction of Chatterjee and Samuelson (1983) provides an extensive-form game that delivers the same outcomes as incomplete information bargaining. In addition, as we show in Appendix D.2, our approach has axiomatic foundations analogous to those that underpin Nash bargaining. Further, intermediaries like eBay, Amazon, and Alibaba play a prominent trade in organizing markets and, as we show in Appendix D.3, provide micro-foundations for incomplete information bargaining. Specifically, building on the model of Loertscher and Niedermayer (2019), for general distributions, one buyer, and any number of suppliers, the incomplete information bargaining outcome arises in equilibrium in an extensive-form

[61]For example, a model based on Nash bargaining might be estimated even when it is clear that alternating-offers bargaining is not a good description of the bargaining process used in reality.

game involving a buyer, suppliers, and a fee-setting broker. This is reminiscent of the role of intermediaries in the wholesale used car market as described by Larsen (2021). There, auction houses run auctions, facilitate further bargaining in the substantial number of cases in which the auction does not result in trade, and collect fees from traders.

8 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. Of particular interest to industrial organization and antitrust economics,[62] it also has the feature that, quite generally, there is a tradeoff between allocating efficiently and extracting rents. This tradeoff is at the heart of both industrial organization and Myerson's optimal auction. This tradeoff 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 tradeoff between profit and social surplus.[63]

There has also been a recent upsurge of interest in bargaining (see, for example, Backus et al., 2020, 2019; Zhang et al., forth.; Larsen, 2021), and buyer power (see, for example, Snyder, 1996; Nocke and Thanassoulis, 2014; Caprice and Rey, 2015; Loertscher and Marx, 2019; Decarolis and Rovigatti, forth.). For example, the latter find empirical evidence that increased buyer power has reduced Google's online search revenues. Larsen and Zhang (2018)

[62]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. The prevalence of procurement in business-to-business transactions provides clear motivation for tools that embrace the incomplete information that drives the price formation process in such markets. Further, the observance of bargaining breakdown in real-world settings is also consistent with the presence of asymmetric information as documented by Backus et al. (2020) and difficult to reconcile with complete information as is the case with the evidence discussed and presented by Backus et al. (2019).

[63]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 tradeoff between profit and social surplus. Without private values, additional and, therefore, in some sense arbitrary, restrictions may be required to maintain tractability and/or the tradeoff between profit and social surplus (Mezzetti, 2004, 2007). Notwithstanding recent progress, with multi-dimensional private information and multiple agents, the optimal mechanism is not known (see, e.g., Daskalakis et al., 2017). With discrete types, there is no payoff equivalence theorem. In other words, the mechanism is not pinned down by the allocation rule.

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.[64] Bargaining has also come to the forefront of the empirical IO literature, in particular in analyses of bundling and vertical integration such as Crawford and Yurukoglu (2012) and Crawford et al. (2018). Collard-Wexler et al. (2019) and Rey and Vergé (2019) provide recent theoretical foundations for the widely used Nash-in-Nash bargaining model.[65] 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,[66] and that the probability of breakdown can, under suitable assumptions, be used to estimate distributions. Ausubel et al. (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.[67] Ordover et al. (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 et al. (2016) show that, while vertical integration solves hold-up problems for the merging parties, it may also create or exacerbate problems for

[64] They provide conditions under which their estimation procedure works well, including a demonstration based on the k-double auction where they estimate both k and the agents' type distributions, interpreting k as a bargaining weight.

[65] 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 et al. (2019)), and provide a micro-foundation for Nash-in-Nash.

[66] As stated by Holmström and Myerson (1983, p. 1809), "Some economists, following Coase have ... argued that we should expect to observe efficient allocations in any economy where there is complete information and bargaining costs are small. However, this positive aspect of efficiency does not extend to economies with incomplete information."

[67] 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 vertical mergers lead to some efficiencies.

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 Grossman and Hart (1986) and Hart and Moore (1990).[68] As discussed, the predictions could hardly differ more starkly because with incomplete (complete) information efficient bargaining implies efficient (inefficient) investment.[69] 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 et al., 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.

9 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 multi-unit 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 tradeoff 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, that the effects of outside options can differ relative to complete information setups, and that bargaining externalities arise naturally.

Our paper shows that an economic agent's strength or weakness has two dimensions that

[68] Nocke and Thanassoulis (2014) provide model within the paradigm of efficient, complete information bargaining in which bargaining power can mitigate frictions due to credit constraints.

[69] The 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 et al. (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.

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, 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. Such differences do not reflect differences in bargaining power as commonly understood. No one would explain that economy airfares are lower than business airfares because of economy customers' 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.[70] The distinction between productive strength and bargaining power brought to light in the present paper may prove useful in that regard.

[70]Byrne et al. (2021) find evidence that bargained prices for retail electricity contracts are affected by an individual's willingness to search and bargain and the extent to which they reveal that they are informed about market prices.

References

Allain, M.-L., C. Chambolle, and P. Rey (2016): "Vertical Integration as a Source of Hold-up," *Review of Economic Studies*, 83, 1–25.

Australian Competition & Consumer Commission (1999): "Merger Guidelines," https://www.accc.gov.au/.

Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002): "Bargaining with Incomplete Information," in *Handbook of Game Theory*, ed. by R. Aumann and S. Hart, Elsevier Science B.V., vol. 3, 1897–1945.

Backus, M., T. Blake, B. Larsen, and S. Tadelis (2020): "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions," *Quarterly Journal of Economics*, 135, 1319–1361.

Backus, M., T. Blake, and S. Tadelis (2019): "On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining," *Journal of Political Economy*, 127, 1599–1628.

Bhattacharyya, S. and A. Nain (2011): "Horizontal Acquisitions and Buying Power: A Product Market Analysis," *Financial Economics*, 99, 97–115.

Börgers, T. (2015): *An Introduction to the Theory of Mechanism Design*, Oxford University Press.

Boyd, S. and L. Vandenberghe (2004): *Convex Optimization*, Cambridge University Press.

Byrne, D. P., L. A. Martin, and J. S. Nah (2021): "Price Discrimination, Search, and Negotiation in an Oligopoly: A Field Experiment in Retail Electricity," University of Melbourne.

Caprice, S. and P. Rey (2015): "Buyer Power from Joint Listing Decision," *Economic Journal*, 125, 1677–1704.

Chatterjee, K. and W. Samuelson (1983): "Bargaining under Incomplete Information," *Operations Research*, 31, 835–851.

Che, Y.-K. (2006): "Beyond the Coasian Irrelevance: Asymmetric Information," Unpublished Lecture Notes, Columbia University.

Che, Y.-K. and D. Hausch (1999): "Cooperative Investments and the Value of Contracting," *American Economic Review*, 89, 125–147.

Choné, P., L. Linnemer, and T. Vergé (2021): "Double Marginalization and Vertical Integration," CESifo Working Paper No. 8971.

Clarke, E. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 11, 17–33.

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019): "Nash-in-Nash: A Microfoundation for Applied Work," *Journal of Political Economy*, 127, 163–195.

Cramton, P., R. Gibbons, and P. Klemperer (1987): "Dissolving a Partnership Efficiently." *Econometrica*, 55, 615–632.

Crawford, G., R. Lee, M. Whinston, and A. Yurukoglu (2018): "The Welfare Effects of Vertical Integration in Multichannel Television Markets," *Econometrica*, 86, 891–954.

Crawford, G. and A. Yurukoglu (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102, 643–685.

Crawford, G. S. (2014): "Cable Regulation in the Internet Era," in *Economic Regulation and Its Reform: What Have We Learned?*, ed. by N. L. Rose, University of Chicago Press, 137–193.

Crémer, J. and R. McLean (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53, 345–361.

——— (1988): "Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.

Dalkir, S., J. W. Logan, and R. T. Masson (2000): "Mergers in Symmetric and Asymmetric

Noncooperative Auction Markets: The Effects on Prices and Efficiency," *International Journal of Industrial Organization*, 18, 383–413.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2017): "Strong Duality for a Multiple-Good Monopolist," *Econometrica*, 85, 735–767.

DeCarolis, F. and G. Rovigatti (forth.): "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising," *American Economic Review*.

Delacrétaz, D., S. Loertscher, L. Marx, and T. Wilkening (2019): "Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade," *Journal of Economic Theory*, 416–454.

Ellison, S. F. and C. M. Snyder (2010): "Countervailing Power in Wholesale Pharmaceuticals," *Journal of Industrial Economics*, 58, 32–53.

European Commission (2004): "Guidelines on the Assessment of Horizontal Mergers Under the Council Regulation on the Control of Concentrations Between Undertakings," Official Journal of the European Union 2004/C 31/03.

Farrell, J. and C. Shapiro (1990): "Horizontal Mergers: An Equilibrium Analysis," *American Economic Review*, 80, 107–126.

Federico, G., G. Langus, and T. Valletti (2017): "A Simple Model of Mergers and Innovation," *Economics Letters*, 157, 136–140.

——— (2018): "Horizontal Mergers and Product Innovation," *International Journal of Industrial Organization*, 59, 1–23.

Fudenberg, D. and J. Tirole (1991): *Game Theory*, MIT Press, Cambridge, Massachusetts.

Galbraith, J. K. (1954): "Countervailing Power," *American Economic Review (Papers and Proceedings)*, 44, 1–6.

Green, J. and J.-J. Laffont (1977): "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." *Econometrica*, 45, 427–438.

Gresik, T. and M. 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, 304–332.

Grossman, S. J. and O. D. Hart (1986): "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691–719.

Groves, T. (1973): "Incentives in Teams," *Econometrica*, 41, 617–631.

Hart, O. and J. Moore (1990): "Property Rights and the Nature of the Firm." *Journal of Political Economy*, 98, 1119–1158.

Hart, O. and J. Tirole (1990): "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity: Microeconomics*, 205–286.

Hatfield, J. W., F. Kojima, and S. D. Kominers (2018): "Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence," SSRN Working Paper.

Hatfield, J. W., S. D. Kominers, R. Lowery, and J. M. Barry (forth.): "Collusion in Markets with Syndication," *Journal of Political Economy*.

Hemphill, S. C. and N. L. Rose (2018): "Mergers that Harm Sellers," *Yale Law Journal*, 127, 2078–2109.

Ho, K. and R. Lee (2017): "Insurer Competition in Health Care Markets," *Econometrica*, 85, 379–417.

Holmström, B. (1979): "Groves' Scheme on Restricted Domains," *Econometrica*, 47, 1137 – 1144.

Holmström, B. and R. B. Myerson (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, 1799–1819.

Jullien, B. and Y. Lefouili (2018): "Horizontal Mergers and Innovation," *Journal of Competition Law & Economics*, 14, 364–392.

Kirkwood, J. B. (2012): "Powerful Buyers and Merger Enforcement," *Boston University Law Review*, 92, 1485–1559.

Krishna, V. (2010): *Auction Theory*, Elsevier, second ed.

Larsen, B. J. (2021): "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," *Review of Economic Studies*, 88, 851–882.

Larsen, B. J. and A. Zhang (2018): "A Mechanism Design Approach to Identification and Estimation," Working Paper, Stanford University.

Lauermann, S. (2013): "Dynamic Matching and Bargaining Games: A General Approach," *American Economic Review*, 103, 663–689.

Legros, P. and A. F. Newman (2013): "A Price Theory of Vertical and Lateral Integration," *Quarterly Journal of Economics*, 128, 725–770.

Loertscher, S. and L. M. Marx (2019): "Merger Review for Markets with Buyer Power," *Journal of Political Economy*, 127, 2967–3017.

Loertscher, S. and C. 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, 101–106.

Loertscher, S. and A. Niedermayer (2019): "An Optimal Pricing Theory of Transaction Fees in Thin Markets," Working Paper, University of Melbourne.

Loertscher, S. and M. Reisinger (2014): "Market Structure and the Competitive Effects of Vertical Integration," *RAND Journal of Economics*, 45, 471–494.

Loertscher, S. and M. H. Riordan (2019): "Make and Buy: Outsourcing, Vertical Integration, and Cost Reduction," *American Economic Journal: Microeconomics*, 11, 105–123.

Loertscher, S. and C. Wasser (2019): "Optimal Structure and Dissolution of Partnerships," *Theoretical Economics*, 14, 1063–1114.

Makowski, L. and C. Mezzetti (1993): "The Possibility of Efficient Mechanisms for Trading and Indivisible Object," *Journal of Economic Theory*, 59, 451–465.

Maskin, E. and J. Riley (1984): "Optimal Auctions with Risk Averse Buyers," *Econometrica*, 52, 1473–1518.

Matthews, S. (1984): "On the Implementability of Reduced Form Auctions," *Econometrica*, 52, 1519–1522.

McAfee, R. P. and M. Schwartz (1994): "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *American Economic Review*, 84, 210–230.

Mezzetti, C. (2004): "Mechanism Design with Interdependent Valuations: Efficiency." *Econometrica*, 72, 1617–1626.

——— (2007): "Mechanism Design with Interdependent Valuations: Surplus Extraction." *Economic Theory*, 31, 473–488.

Milgrom, P. (1987): "Adverse Selection without Hidden Information," Working Paper #8742, University of California Berkeley.

Mussa, M. and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–317.

Myerson, R. and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Mylovanov, T. and T. Tröger (2014): "Mechanism Design by an Informed Principal: Private Values with Transferable Utility," *Review of Economic Studies*, 81, 1668–1707.

Nocke, V. and P. Rey (2018): "Exclusive Dealing and Vertical Integration in Interlocking Relationships," *Journal of Economic Theory*, 177, 183–221.

Nocke, V. and J. Thanassoulis (2014): "Vertical Relations under Credit Constraints," *Journal of the European Economic Association*, 12, 337–367.

OECD (2007): "Competition and Efficient Usage of Payment Cards 2006," *OECD Policy Roundtables*, DAF/COMP(2006)32.

——— (2011): "Competition and Efficient Usage of Payment Cards," *OECD Journal: Competition Law and Policy*, 2009.

Ordover, J. A., G. Saloner, and S. C. Salop (1990): "Equilibrium Vertical Foreclosure," *American Economic Review*, 80, 127–142.

Perry, M. K. and R. H. Porter (1985): "Oligopoly and the Incentive for Horizontal Merger," *American Economic Review*, 75, 219–227.

Rey, P. and T. Vergé (2019): "Secret Contracting in Multilateral Relations," TSE Working Paper.

Riordan, M. (2008): "Competitive Effects of Vertical Integration," in *Handbook of Antitrust Economics*, ed. by P. Buccirossi, MIT Press.

Riordan, M. H. (1998): "Anticompetitive Vertical Integration by a Dominant Firm," *American Economic Review*, 88, 1232–1248.

Riordan, M. H. and S. C. Salop (1995): "Evaluating Vertical Mergers: A Post-Chicago Approach," *Antitrust Law Journal*, 63, 513–568.

Rogerson, W. P. (1992): "Contractual Solutions to the Hold-up Problem," *Review of Economic Studies*, 59, 777–793.

Salant, S. W., S. Switzer, and R. J. Reynolds (1983): "Losses From Horizontal Merger: The Effects of an Exogenous Change in Industry Structure on Cournot-Nash Equilibrium," *Quarterly Journal of Economics*, 98, 185–199.

Salinger, M. A. (1988): "Vertical Mergers and Market Foreclosure," *Quarterly Journal of Economics*, 103, 345–356.

Schmitz, P. W. (2002): "Simple Contracts, Renegotiation under Asymmetric Information, and the Hold-up Problem," *European Economic Review*, 46, 169–188.

Segal, I. and M. D. Whinston (2011): "A Simple Status Quo that Ensures Participation with Application to Efficient Bargaining," *Theoretical Economics*, 6, 109–125.

Snyder, C. M. (1996): "A Dynamic Theory of Countervailing Power," *RAND Journal of Economics*, 27, 747–769.

——— (2008): "Countervailing Power," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, London: Palgrave Macmillan, 1188–1189.

Stigler, G. J. (1954): "The Economist Plays with Blocs," *American Economic Review (Papers and Proceedings)*, 44, 7–14.

U.S. Department of Justice and the Federal Trade Commission (1996): "Statement of Antitrust Enforcement Policy in Health Care," https://www.justice.gov/atr/public/guidelines/1791.htm.

——— (2010): "Horizontal Merger Guidelines," Https://www.justice.gov/atr/horizontal-merger-guidelines-08192010.

——— (2020): "Vertical Merger Guidelines," https://www.ftc.gov/system/files/documents/reports/us-department-justice-federal-trade-commission-vertical-merger-guidelines/vertical_merger_guidelines_6-30-20.pdf.

Valley, K., L. Thompson, R. Gibbons, and M. H. Bazerman (2002): "How Communication

Improves Efficiency in Bargaining Games," *Games and Economic Behavior*, 38, 127–155.

Vickrey, W. (1961): "Counterspeculation, Auction, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

Waehrer, K. (1999): "Asymmetric Private Values Auctions with Application to Joint Bidding and Mergers," *International Journal of Industrial Organization*, 17, 437–452.

Williams, S. R. (1987): "Efficient Performance in Two Agent Bargaining," *Journal of Economic Theory*, 41, 154–172.

——— (1999): "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14, 155–180.

Zhang, X., P. Manchanda, and J. Chu (forth.): "'Meet Me Halfway': The Costs and Benefits of Bargaining," *Marketing Science*.

Appendices are for online publication only

Online Appendix

Incomplete information bargaining with applications to mergers, investment, and vertical integration

Simon Loertscher and Leslie M. Marx

May 24, 2021

A Appendix: Mechanism design foundations

In this appendix, we first define and develop the mechanism design concepts relevant for our analysis (Appendix A.1) and then apply these concepts to derive the Myerson-Satterthwaite impossibility result and the second-best mechanism (Appendix A.2).

A.1 Concepts and derivations

Take as given a direct mechanism $\langle \mathbf{Q}, \mathbf{M} \rangle$, where for $i \in \mathcal{N}^B$ and $j \in \mathcal{N}^S$, $Q_i^B : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to [0, k_i^B]$, $Q_j^S : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to [0, k_j^S]$, and $M_i^B, M_j^S : [\underline{v}, \overline{v}]^{n^B} \times [\underline{c}, \overline{c}]^{n^S} \to \mathbb{R}$. Given reports $(\mathbf{v}, \mathbf{c})$, $Q_i^B(\mathbf{v}, \mathbf{c})$ is the quantity received by buyer i, $Q_j^S(\mathbf{v}, \mathbf{c})$ is the quantity provided by supplier j, $M_i^B(\mathbf{v}, \mathbf{c})$ is the payment from buyer i to the mechanism, and $M_j^S(\mathbf{v}, \mathbf{c})$ is the payment from the mechanism to supplier j. By the Revelation Principle, the focus on direct mechanisms is without loss of generality.

Let $\hat{q}_i^B(z)$ be the buyer i's expected quantity if it reports z and all other agents report truthfully, and let $\hat{m}_i^B(z)$ be buyer i's expected payment if it reports z and all other agents report truthfully:

$$\hat{q}_i^B(z) = \mathbb{E}_{\mathbf{v}_{-i}, \mathbf{c}}[Q_i^B(z, \mathbf{v}_{-i}, \mathbf{c})] \quad \text{and} \quad \hat{m}_i^B(z) = \mathbb{E}_{\mathbf{v}_{-i}, \mathbf{c}}[M_i^B(z, \mathbf{v}_{-i}, \mathbf{c})]. \tag{13}$$

Define $\hat{q}_j^S$ and $\hat{m}_j^S$ analogously, where $\hat{m}_j^S$ is the expected payment to supplier j. Because we assume independent draws, these interim expected quantities and payments depend only on the report z and not on the reporting agent's true type. The expected payoff of buyer i with type v that reports z is then $\hat{q}_i^B(z)v - \hat{m}_i^B(z)$, and the expected payoff of supplier j with type c that reports z is $\hat{m}_j^S(z) - \hat{q}_j^S(z)c$.

Key constraints

The mechanism is *incentive compatible* for buyer i if for all $v, z \in [\underline{v}, \overline{v}]$,

$$\hat{u}_i^B(v) \equiv \hat{q}_i^B(v)v - \hat{m}_i^B(v) \geq \hat{q}_i^B(z)v - \hat{m}_i^B(z), \tag{14}$$

and is *incentive compatible* for supplier j if for all $c, z \in [\underline{c}, \overline{c}]$,

$$\hat{u}_j^S(c) \equiv \hat{m}_j^S(c) - \hat{q}_j^S(c)c \geq \hat{m}_j^S(z) - \hat{q}_j^S(z)c. \tag{15}$$

Individual rationality is satisfied for buyer i if for all $v \in [\underline{v}, \overline{v}]$, $\hat{u}_i^B(v) \geq 0$, and for supplier j if for all $c \in [\underline{c}, \overline{c}]$, $\hat{u}_j^S(c) \geq 0$. The mechanism satisfies the *no-deficit* condition if

$$\mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} M_i^B(\mathbf{v}, \mathbf{c}) - \sum_{j \in \mathcal{N}^S} M_j^S(\mathbf{v}, \mathbf{c}) \right] \geq 0.$$

Interim expected payoffs

Standard arguments (see, e.g., Krishna, 2010, Chapter 5.1) proceed as follows:

Focusing on buyer i, incentive compatibility implies that

$$\hat{u}_i^B(v) = \max_{z \in [\underline{v}, \overline{v}]} \hat{q}_i^B(z) v - \hat{m}_i^B(z),$$

i.e., $\hat{u}_i^B$ is a maximum of a family of affine functions, which implies that $\hat{u}_i^B$ is convex and so absolutely continuous and differentiable almost everywhere in the interior of its domain.[1] In addition, incentive compatibility implies that $\hat{u}_i^B(z) \geq \hat{q}_i^B(v) z - \hat{m}_i^B(v) = \hat{u}_i^B(v) + \hat{q}_i^B(v)(z - v)$, which for $\varepsilon > 0$ implies

$$\frac{\hat{u}_i^B(v + \varepsilon) - \hat{u}_i^B(v)}{\varepsilon} \geq \hat{q}_i^B(v)$$

and for $\varepsilon < 0$ implies

$$\frac{\hat{u}_i^B(v + \varepsilon) - \hat{u}_i^B(v)}{\varepsilon} \leq \hat{q}_i^B(v),$$

so taking the limit as ε goes to zero, at every point v where $\hat{u}_i^B$ is differentiable, $\hat{u}_i^{B\prime}(v) = \hat{q}_i^B(v)$. Because $\hat{u}_i^B$ is convex, this implies that $\hat{q}_i^B(v)$ is nondecreasing. Because every absolutely continuous function is the definite integral of its derivative,

$$\hat{u}_i^B(v) = \hat{u}_i^B(\underline{v}) + \int_{\underline{v}}^{v} \hat{q}_i^B(t) dt,$$

which implies that, up to an additive constant, buyer i's expected payoff in an incentive-compatible direct mechanism depends only on the allocation rule. By an analogous argument, $\hat{u}_j^{S\prime}(c) = -\hat{q}_j^S(c)$, $\hat{q}_j^S(c)$ is nonincreasing, and

$$\hat{u}_j^S(c) = \hat{u}_j^S(\overline{c}) + \int_{c}^{\overline{c}} \hat{q}_j^S(t) dt.$$

[1] A function $h : [\underline{v}, \overline{v}] \to \mathbb{R}$ is absolutely continuous if for all $\varepsilon > 0$ there exists $\delta > 0$ such that whenever a finite sequence of pairwise disjoint sub-intervals (v_k, v_k') of $[\underline{v}, \overline{v}]$ satisfies $\sum_k (v_k' - v_k) < \delta$, then $\sum_k |h(v_k') - h(v_k)| < \varepsilon$. One can show that absolute continuity on compact interval $[a, b]$ implies that h has a derivative h' almost everywhere, the derivative is Lebesgue integrable, and that $h(x) = h(a) + \int_a^x h'(t) dt$ for all $x \in [a, b]$.

Mechanism budget surplus

Using the definitions of $\hat{u}_i^B$ and $\hat{u}_j^S$ in (14) and (15), we can rewrite these as

$$\hat{m}_i^B(v) = \hat{q}_i^B(v)v - \int_{\underline{v}}^{v} \hat{q}_i^B(t)dt - \hat{u}_i^B(\underline{v}) \tag{16}$$

and

$$\hat{m}_j^S(c) = \hat{q}_j^S(c)c + \int_{c}^{\overline{c}} \hat{q}_j^S(t)dt + \hat{u}_j^S(\overline{c}). \tag{17}$$

The expected payment by buyer i is then

$$\begin{aligned}
\mathbb{E}_v\left[\hat{m}_i^B(v)\right] &= \int_{\underline{v}}^{\overline{v}} \hat{m}_i^B(v)f_i(v)dv \\
&= \int_{\underline{v}}^{\overline{v}} \left(\hat{q}_i^B(v)v - \int_{\underline{v}}^{v} \hat{q}_i^B(t)dt\right) f_i(v)dv - \hat{u}_i^B(\underline{v}) \\
&= \int_{\underline{v}}^{\overline{v}} \hat{q}_i^B(v)vf_i(v)dv - \int_{\underline{v}}^{\overline{v}}\int_{t}^{\overline{v}} \hat{q}_i^B(t)f_i(v)dvdt - \hat{u}_i^B(\underline{v}) \\
&= \int_{\underline{v}}^{\overline{v}} \hat{q}_i^B(v)vf_i(v)dv - \int_{\underline{v}}^{\overline{v}} \hat{q}_i^B(t)\left(1 - F_i(t)\right)dt - \hat{u}_i^B(\underline{v}) \\
&= \int_{\underline{v}}^{\overline{v}} \hat{q}_i^B(v)\left(v - \frac{1 - F_i(v)}{f_i(v)}\right) f_i(v)dv - \hat{u}_i^B(\underline{v}) \\
&= \int_{\underline{v}}^{\overline{v}} \hat{q}_i^B(v)\Phi_i(v)f_i(v)dv - \hat{u}_i^B(\underline{v}) \\
&= \mathbb{E}_v\left[\hat{q}_i^B(v)\Phi_i(v)\right] - \hat{u}_i^B(\underline{v}),
\end{aligned}$$

where the first equality uses the definition of the expectation, the second uses (16), the third switches the order of integration, the fourth integrates, the fifth collects terms, the sixth uses the definition of the virtual value Φ_i, and the last equality uses the definition of the expectation. Similarly, using (17), the expected payment to supplier j is

$$\mathbb{E}_c\left[\hat{m}_j^S(c)\right] = \int_{\underline{c}}^{\overline{c}} \hat{m}_j^S(c)g_j(c)dc = \mathbb{E}_c\left[\hat{q}_j^S(c)\Gamma_j(c)\right] + \hat{u}_j^S(\overline{c}).$$

Thus, we have the result that in any incentive-compatible, interim individually-rational direct mechanism $\langle \mathbf{Q}, \mathbf{M} \rangle$, the mechanism's expected budget surplus is

$$\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{i \in N^B} \Phi_i(v_i)Q_i^B(\mathbf{v},\mathbf{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j(c_j)Q_j^S(\mathbf{v},\mathbf{c})\right] - \sum_{i \in N^B} \hat{u}_i^B(\underline{v}) - \sum_{j \in \mathcal{N}^S} \hat{u}_j^S(\overline{c}).$$

A.2 Myerson-Satterthwaite Redux

Impossibility result

For the purpose of making the paper self-contained, we provide a statement and proof of the impossibility theorem of Myerson and Satterthwaite (1983). Under the assumption of independent private values and the assumption that $\underline{v} < \overline{c}$, Myerson and Satterthwaite (1983) show that there is no mechanism satisfying incentive compatibility and individual rationality that allocates ex post efficiently and that does not run a deficit. Their result depends on $\underline{v} < \overline{c}$ because, without this assumption, ex post efficiency subject to incentive compatibility and individual rationality can easily be achieved without running a deficit. For example, the *posted price* mechanism that has the buyer pay $p = (\underline{v} + \overline{c})/2$ to the supplier achieves this.

By now, the proof of this result can be provided in a couple of lines (see, e.g, Krishna, 2010). Consider the dominant strategy implementation in which the buyer pays $p^B = \max\{c, \underline{v}\}$ and the supplier receives $p^S = \min\{v, \overline{c}\}$ whenever there is trade, and no payments are made otherwise. Notice that $\hat{u}^B(\underline{v}) = 0 = \hat{u}^S(\overline{c})$. Thus, the individual rationality constraints are satisfied. Further, notice that $p^B - p^S \leq 0$, with a strict inequality for almost all type realizations. This implies that the mechanism runs a deficit in expectation. By the payoff equivalence theorem, any other ex post efficient mechanism satisfying incentive compatibility and individual rationality will run a deficit of at least that size (and a larger one if one or both of the individual rationality constraints are slack).

To see how this impossibility result rests on the assumption that $\underline{v} < \overline{c}$, assume to the contrary that $\underline{v} \geq \overline{c}$. Then the mechanism described above continues to satisfy incentive compatibility and individual rationality, but for all type realizations $p^B = \underline{v} \geq \overline{c} = p^S$, which implies that the mechanism does not run a deficit.

Second-best mechanism

The impossibility result in the bilateral trade problem of Myerson and Satterthwaite raises the question as to what is the second-best mechanism, that is, the mechanism that maximizes equally weighted social surplus subject to incentive compatibility and individual rationality constraints and the constraint of no deficit. Denoting by F and G the buyer's and seller's distributions, which are assumed to exhibit increasing virtual type functions $\Phi(v) = v - \frac{1-F(v)}{f(v)}$ and $\Gamma(c) = c + \frac{G(c)}{g(c)}$, and using incentive compatibility, the second-best mechanism maximizes the equally weighted surplus of the buyer and the seller,[2]

$$\int_{\underline{v}}^{\overline{v}}\int_{\underline{c}}^{\overline{c}}(v-\Phi(v)+\Gamma(c)-c)Q(v,c)g(c)f(v)dcdv+\hat{u}^B(\underline{v})+\hat{u}^S(\overline{c}),$$

subject to the no-deficit constraint,

$$\int_{\underline{v}}^{\overline{v}}\int_{\underline{c}}^{\overline{c}}(\Phi(v)-\Gamma(c))Q(v,c)g(c)f(v)dcdv-\hat{u}^B(\underline{v})-\hat{u}^S(\overline{c})\geq 0,$$

and the individual rationality constraints

$$\hat{u}^B(\underline{v})\geq 0 \quad\text{and}\quad \hat{u}^S(\overline{c})\geq 0.$$

Letting ρ denote the Lagrange multiplier associated with the no-deficit constraint, which must be positive because of the Myerson-Satterthwaite impossibility result, and μ^B and μ^S be the multipliers on the individual rationality constraints, the Lagrangian can be written as

$$\begin{aligned}&\rho\int_{\underline{v}}^{\overline{v}}\int_{\underline{c}}^{\overline{c}}\left[v-\frac{\rho-1}{\rho}\frac{1-F(v)}{f(v)}-c-\frac{\rho-1}{\rho}\frac{G(c)}{g(c)}\right]Q(v,c)g(c)f(v)dcdv\\ +\;&(1-\rho+\mu^B)\hat{u}^B(\underline{v})+(1-\rho+\mu^S)\hat{u}^S(\overline{c}).\end{aligned}$$

The Lagrange multipliers, ρ, μ^B, and μ^S, must be nonnegative, and optimization with respect to $\hat{u}^B(\underline{v})$ and $\hat{u}^S(\overline{c})$ requires that $1-\rho+\mu^B=0$ and $1-\rho+\mu^S=0$, which cannot be satisfied if $\rho<1$. Therefore, we conclude that $\rho\geq 1$. Intuitively, if the shadow price of the no-deficit constraint is less than 1, the Lagrangian is maximized by running an infinite budget deficit and paying that out to the agents in fixed payments, violating primal feasibility.

Recalling that for $a\in[0,1]$ we define $\Phi^a(v)\equiv v-(1-a)\frac{1-F(v)}{f(v)}$ and $\Gamma^a(c)\equiv c+(1-a)\frac{G(c)}{g(c)}$, we can rewrite the Lagrangian as

$$\rho\int_{\underline{v}}^{\overline{v}}\int_{\underline{c}}^{\overline{c}}\left[\Phi^{1/\rho}(v)-\Gamma^{1/\rho}(c)\right]Q(v,c)g(c)f(v)dcdv+(1-\rho+\mu^B)\hat{u}^B(\underline{v})+(1-\rho+\mu^S)\hat{u}^S(\overline{c}).$$

It follows that for a given ρ, the Lagrangian is maximized with respect to Q pointwise by setting $Q(v,c)=1$ if $\Phi^{1/\rho}(v)\geq\Gamma^{1/\rho}(c)$ and $Q(v,c)=0$ otherwise. Because the virtual types are increasing, this pointwise maximizer, denoted $Q^\rho(v,c)$, is increasing in v and decreasing

[2] Myerson and Satterthwaite (1983) write the objective as $\int_{\underline{v}}^{\overline{v}}\int_{\underline{c}}^{\overline{c}}(v-c)Q(v,c)g(c)f(v)dcdv$, that is, without accounting for the payments from the buyer and to the supplier, but that does not affect the conclusions for the case that they consider because the budget surplus, not including fixed payments, is 0.

in c for any $\rho \geq 1$. Hence, there is an incentive compatible implementation.

Putting this together, the allocation rule of the second-best mechanism is given by Q^ρ with the smallest distortion, i.e., the smallest value of $\rho \in [1, \infty)$, such that the no-deficit constraint can be satisfied for some nonnegative fixed payments, i.e., such that $\int_{\underline{v}}^{\overline{v}} \int_{\underline{c}}^{\overline{c}} (\Phi(v) - \Gamma(c)) Q^\rho(v, c) g(c) f(v) dc dv \geq 0$.

B Appendix: Proofs

Proof of Lemma 1. We can write the Lagrangian that corresponds to (5) for the case with multiple buyers and suppliers, each with single-unit demand and supply, as

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} w_i^B (v_i - \Phi_i(v_i)) Q_i^B(\mathbf{v}, \boldsymbol{c}) + \sum_{j \in \mathcal{N}^S} w_j^S (\Gamma_j(c_j) - c_j) Q_j^S(\mathbf{v}, \boldsymbol{c}) \right] \\
& + \rho \mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} \Phi_i(v_i) Q_i^B(\mathbf{v}, \boldsymbol{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j(c_j) Q_j^S(\mathbf{v}, \boldsymbol{c}) \right] \\
= \; & \mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} \left[w_i^B v_i + (\rho - w_i^B) \Phi_i(v_i) \right] Q_i^B(\mathbf{v}, \boldsymbol{c}) + \sum_{j \in \mathcal{N}^S} \left[-w_j^S c_j - (\rho - w_j^S) \Gamma_j(c_j) \right] Q_j^S(\mathbf{v}, \boldsymbol{c}) \right] \\
= \; & \rho \mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} \left[v_i - \frac{\rho - w_i^B}{\rho} \frac{1 - F_i(v_i)}{f_i(v_i)} \right] Q_i^B(\mathbf{v}, \boldsymbol{c}) - \sum_{j \in \mathcal{N}^S} \left[c_j + \frac{\rho - w_j^S}{\rho} \frac{G_j(c_j)}{g_j(c_j)}) \right] Q_j^S(\mathbf{v}, \boldsymbol{c}) \right] \\
= \; & \rho \mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{i \in \mathcal{N}^B} \Phi_i^{w_i^B/\rho}(v_i) Q_i^B(\mathbf{v}, \boldsymbol{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j^{w_j^S/\rho}(c_j) Q_j^S(\mathbf{v}, \boldsymbol{c}) \right].
\end{aligned}
$$

It is then clear that the allocation rule defined in the statement of the lemma maximizes the Lagrangian pointwise subject to the feasibility constraints, which completes the proof. ■

Proof of Proposition 2. Recall that $\mathcal{M}$ is the set of incentive compatible, individually rational, no-deficit mechanisms. Let $\mathcal{C} \subset \mathbb{R}^{n^S + n^B}$ be the induced space of expected payoffs associated with $\mathcal{M}$. Note that $\mathcal{M}$ is convex,[3] and correspondingly $\mathcal{C}$ is also convex. The incomplete

[3]Given $(\mathbf{Q}^0, \mathbf{M}^0), (\mathbf{Q}^1, \mathbf{M}^1) \in \mathcal{M}$ and $\lambda \in [0,1]$ and defining $(\mathbf{Q}^\lambda, \mathbf{M}^\lambda)$ by for $j \in \mathcal{N}^S$, $Q_j^{\lambda,S}(\mathbf{v}, \mathbf{c}) \equiv (1-\lambda) Q_j^{0,S}(\mathbf{v}, \mathbf{c}) + \lambda Q_j^{1,S}(\mathbf{v}, \mathbf{c})$ and $M_j^{\lambda,S}(\mathbf{v}, \mathbf{c}) \equiv (1-\lambda) M_j^{0,S}(\mathbf{v}, \mathbf{c}) + \lambda M_j^{1,S}(\mathbf{v}, \mathbf{c})$, and similarly for $Q_i^{\lambda,B}$ and $M_i^{\lambda,B}$ with $i \in \mathcal{N}^B$, then we have $(\mathbf{Q}^\lambda, \mathbf{M}^\lambda) \in \mathcal{M}$. That is, each supplier j's expected payoff under $(\mathbf{Q}^\lambda, \mathbf{M}^\lambda)$, $\mathbb{E}_{\mathbf{v},\mathbf{c}}[M_j^{\lambda,S}(\mathbf{v}, \mathbf{c}) - c_j Q_j^{\lambda,S}(\mathbf{v}, \mathbf{c})]$, is the convex combination of its payoffs under the two component mechanisms, and analogously for buyer i. Consequently, incentive compatibility and individual rationality are satisfied, and the no-deficit constraint continues to be satisfied.

information bargaining mechanism solves

$$\max_{\mathbf{u}\in\mathcal{C}} \sum_{j\in\mathcal{N}^S} w_j^S u_j^S + \sum_{i\in\mathcal{N}^B} w_i^B u_i^B. \tag{18}$$

The solution to (18) is Pareto optimal and, by the dual characterization of maximal elements (see, e.g., Boyd and Vandenberghe, 2004, Chapter 2.6.3), any Pareto optimal $\tilde{\mathbf{u}}$ solves $\max_{\mathbf{u}\in\mathcal{C}} \tilde{\mathbf{w}}^T\mathbf{u}$ for some nonzero $\tilde{\mathbf{w}}$ satisfying $\tilde{\mathbf{w}} \geq \mathbf{0}$. Because we can rescale $\tilde{\mathbf{w}}$ by $1/\max\tilde{\mathbf{w}}$, there exists $\mathbf{w} \in [0,1]^{n^S+n^B}$ with $\mathbf{w} \neq \mathbf{0}$ such that $\tilde{\mathbf{u}}$ solves $\max_{\mathbf{u}\in\mathcal{C}} \mathbf{w}^T\mathbf{u}$. Because $\tilde{\mathbf{u}} \in \mathcal{C}$, there exists a mechanism $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle \in \mathcal{M}$ that generates payoffs $\tilde{\mathbf{u}}$. Letting

$$\tilde{\pi} \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{i\in\mathcal{N}^B} \Phi_i(v_i)\tilde{Q}_i^B(v,c) - \sum_{j\in\mathcal{N}^S} \Gamma_j(c_j)\tilde{Q}_j^S(\mathbf{v},\mathbf{c})\right],$$

which is nonnegative by virtue of $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle$ satisfying individual rationality and having no deficit, we can define $\boldsymbol{\eta} \in [0,1]^{n^S+n^B}$ with $\sum_{j\in\mathcal{N}^S} \eta_j^S + \sum_{i\in\mathcal{N}^B} \eta_i^B = 1$ by, for $j \in \mathcal{N}^S$,

$$\tilde{u}_j^S = \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[(\Gamma_j(c_j) - c_j)\tilde{Q}_j^S(\mathbf{v},\mathbf{c})\right] + \eta_j^S\tilde{\pi}$$

and for $i \in \mathcal{N}^B$,

$$\tilde{u}_i^B = \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[(v_i - \Phi_i(v_i))\tilde{Q}_i^B(\mathbf{v},\mathbf{c})\right] + \eta_i^B\tilde{\pi}.$$

If $w_i^x < \max\mathbf{w}$, we have $\eta_i^x = 0$ for else $\tilde{\mathbf{u}}$ would not maximize $\mathbf{w}^T\mathbf{u}$ over $\mathcal{C}$, which would be a contradiction. Thus, $\mathbf{w}$ and $\boldsymbol{\eta}$ satisfy the conditions to be bargaining weights and tie-breaking shares, and $u_j^S(\mathbf{w},\boldsymbol{\eta}) = \tilde{u}_j^S$ and $u_i^B(\mathbf{w},\boldsymbol{\eta}) = \tilde{u}_i^B$. This completes the proof. ■

Proof of Proposition 3. The discussion in the text shows that the planner's and market's outcomes coincide (up to fixed payments) if (i)–(iv) hold, implying that $W^{\mathbf{w}} = W^*$, and so there is no benefit from equalization of bargaining power. It remains to show that $W^{\mathbf{w}} < W^*$ if any one of these conditions fails.

Case 1. Suppose that $K^B \leq K^S$ and (8) fails to hold. (Analogous arguments apply if $K^B > K^S$ and (8) fails.) Then for an open set of types, not all of the n^B buyers trade under $\mathbf{Q}^{\mathbf{w}}$. Thus, in order for $\mathbf{Q}^{\mathbf{w}}$ and $\mathbf{Q}^*$ to coincide, they must agree on not only the ranking within buyers and within suppliers, but also the ranking across buyers and suppliers. Consistent with (ii)–(iv), suppose that all buyers have the same bargaining weight w^B, all suppliers have the same bargaining weight w^S, $w^S < w^B$, and all suppliers have the same distribution. (Analogous analysis applies if $w^S > w^B$ and all buyers have the same distribution.) Then the planner and market both rank the buyers the same and rank the suppliers the same, but they evaluate the buyers' virtual values using weight $w^B/\rho^{\mathbf{w}}$ and the suppliers' virtual costs using weight $w^S/\rho^{\mathbf{w}}$, where $w^B/\rho^{\mathbf{w}} > w^S/\rho^{\mathbf{w}}$. Because either $w^B/\rho^{\mathbf{w}} \neq 1/\rho^{\mathbf{1}}$ or $w^S/\rho^{\mathbf{w}} \neq 1/\rho^{\mathbf{1}}$ or both, $\mathbf{Q}^{\mathbf{w}}(\mathbf{v},\mathbf{c}) \neq \mathbf{Q}^*(\mathbf{v},\mathbf{c})$ for all $(\mathbf{v},\mathbf{c})$ in an open subset of $[\underline{v},\overline{v}]^{n^B} \times [\underline{c},\overline{c}]^{n^S}$.

Case 2. If either the buyers' weights are not equal or the suppliers' weights are not equal, then the planner and market rank the agents differently on that side of the market and so $\mathbf{Q}^{\mathbf{w}}(\mathbf{v},\mathbf{c}) \neq \mathbf{Q}^*(\mathbf{v},\mathbf{c})$ for all $(\mathbf{v},\mathbf{c})$ in an open subset of $[\underline{v},\overline{v}]^{n^B} \times [\underline{c},\overline{c}]^{n^S}$.

Case 3. Suppose that (i) and (ii) hold and that $w^S < w^B$, but that $G_1 \neq G_2$, so that (iii) fails. It follows that $1 \geq w^B/\rho^{\mathbf{w}} > w^S/\rho^{\mathbf{w}}$. Because $w^S/\rho^{\mathbf{w}} < 1$ and $G_1 \neq G_2$, the market's ranking of suppliers 1 and 2 based on their virtual costs differs from the ranking of their costs for (c_1, c_2) in an open subset of $[\underline{c},\overline{c}]^2$. Thus, $\mathbf{Q}^{\mathbf{w}}(\mathbf{v},\mathbf{c}) \neq \mathbf{Q}^*(\mathbf{v},\mathbf{c})$ for all $(\mathbf{v},\mathbf{c})$ in an open subset of $[\underline{v},\overline{v}]^{n^B} \times [\underline{c},\overline{c}]^{n^S}$.

Case 4. Suppose that (i) and (ii) hold and that $w^B < w^S$, but that $F_1 \neq F_2$, so that (iv) fails. It follows that $1 \geq w^S/\rho^{\mathbf{w}} > w^B/\rho^{\mathbf{w}}$. Because $w^B/\rho^{\mathbf{w}} < 1$ and $F_1 \neq F_2$, the market's ranking of buyers 1 and 2 based on their virtual values differs from the ranking of their values for (v_1, v_2) in an open subset of $[\underline{v},\overline{v}]^2$. Thus, $\mathbf{Q}^{\mathbf{w}}(\mathbf{v},\mathbf{c}) \neq \mathbf{Q}^*(\mathbf{v},\mathbf{c})$ for all $(\mathbf{v},\mathbf{c})$ in an open subset of $[\underline{v},\overline{v}]^{n^B} \times [\underline{c},\overline{c}]^{n^S}$. ■

Proof of Lemma 2. Given $u \in [\underline{u}_S, \overline{u}_S]$, $\omega(u)$ is defined by the mechanism that maximizes

$$\begin{aligned}
&\sum_{i\in\mathcal{N}^B} \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[(v_i - \Phi_i(v_i))Q_i^B(\mathbf{v},\mathbf{c}) + \hat{u}_i^B(\underline{v})\right] \\
=\; &\sum_{i\in\mathcal{N}^B} \left(\int_{[\underline{v},\overline{v}]^{n^B}} \int_{[\underline{c},\overline{c}]^{n^S}} (v_i - \Phi_i(v_i))Q_i^B(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v}) + \hat{u}_i^B(\underline{v})\right),
\end{aligned}$$

where $dG(\mathbf{c}) \equiv dG_1(c_1)\cdots dG_{n^S}(c_{n^S})$ and $dF(\mathbf{v}) \equiv dF_1(v_1)\cdots dF_{n^B}(v_{n^B})$, subject to the no-deficit constraint

$$\begin{aligned}
&\sum_{i\in\mathcal{N}^B} \left(\int_{[\underline{v},\overline{v}]^{n^B}} \int_{[\underline{c},\overline{c}]^{n^S}} \Phi_i(v_i)Q_i^B(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v}) - \hat{u}_i^B(\underline{v})\right) \\
&- \sum_{j\in\mathcal{N}^S} \left(\int_{[\underline{v},\overline{v}]^{n^B}} \int_{[\underline{c},\overline{c}]^{n^S}} \Gamma_j(c_j)Q_j^S(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v}) + \hat{u}_j^S(\overline{c})\right) \\
\geq\; & 0,
\end{aligned}$$

the individual rationality constraints

$$\text{for all } \; i \in \mathcal{N}^B, \; \hat{u}_i^B(\underline{v}) \geq 0 \; \text{ and for all } \; j \in \mathcal{N}^S, \; \hat{u}^S(\overline{c}) \geq 0,$$

and the constraint that total supplier surplus is at least u,

$$\sum_{j\in\mathcal{N}^S}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}(\Gamma_j(c_j)-c_j)Q_j^S(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})+\hat{u}_j^S(\overline{c})\right)\geq u.$$

Letting ρ denote the Lagrange multiplier associated with the no-deficit constraint, μ_i^B and μ_j^S be the multipliers on the individual rationality constraints, and γ be the multiplier on the constraint that total supplier surplus is at least u, the Lagrangian is

$$\begin{aligned}
&\sum_{i\in\mathcal{N}^B}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}(v_i-\Phi_i(v_i))Q_i^B(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})+\hat{u}_i^B(\underline{v})\right)\\
&+\rho\sum_{i\in\mathcal{N}^B}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}\Phi_i(v_i)Q_i^B(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})-\hat{u}_i^B(\underline{v})\right)\\
&-\rho\sum_{j\in\mathcal{N}^S}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}\Gamma_j(c_j)Q_j^S(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})+\hat{u}_j^S(\overline{c})\right)\\
&+\sum_{i\in\mathcal{N}^B}\mu_i^B\hat{u}_i^B(\underline{v})+\sum_{j\in\mathcal{N}^S}\mu_j^S\hat{u}_j^S(\overline{c})\\
&+\gamma\sum_{j\in\mathcal{N}^S}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}(\Gamma_j(c_j)-c_j)Q_j^S(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})+\hat{u}_j^S(\overline{c})\right)-\gamma u,
\end{aligned}$$

which we can rewrite as

$$\begin{aligned}
&\rho\sum_{i\in\mathcal{N}^B}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}\left(v_i-\frac{\rho-1}{\rho}\frac{1-F_i(v_i)}{f_i(v_i)}\right)Q_i^B(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})\right)\\
&-\rho\sum_{j\in\mathcal{N}^S}\left(\int_{[\underline{v},\overline{v}]^{n^B}}\int_{[\underline{c},\overline{c}]^{n^S}}\left(c_j+\frac{\rho-\gamma}{\rho}\frac{G_j(c_j)}{g_j(c_j)}\right)Q_j^S(\mathbf{v},\mathbf{c})dG(\mathbf{c})dF(\mathbf{v})\right)\\
&+\sum_{i\in\mathcal{N}^B}(1-\rho+\mu_i^B)\hat{u}_i^B(\underline{v})+\sum_{j\in\mathcal{N}^S}(\gamma-\rho+\mu_j^S)\hat{u}_j^S(\overline{c})-\gamma u.
\end{aligned}$$

The Lagrange multipliers, ρ, μ_i^B, μ_j^S, and γ must be nonnegative, and optimization with respect to $\hat{u}^B(\underline{v})$ and $\hat{u}^S(\overline{c})$ requires that $1-\rho+\mu_i^B=0$ and $\gamma-\rho+\mu_j^S=0$, which cannot be satisfied if $\rho<\max\{1,\gamma\}$. Therefore, we conclude that $\rho\geq\max\{1,\gamma\}$. In addition, because a positive expected budget surplus is always possible given our assumption that $\overline{v}>\underline{c}$, the shadow price ρ is finite.

Recalling that for $a\in[0,1]$ we define $\Phi_i^a(v)\equiv v-(1-a)\frac{1-F_i(v)}{f_i(v)}$ and $\Gamma_j^a(c)\equiv c+(1-a)\frac{G_j(c)}{g_j(c)}$, we can rewrite the Lagrangian as

$$\rho \int_{[\underline{v},\overline{v}]^{n^B}} \int_{[\underline{c},\overline{c}]^{n^S}} \left[\sum_{i \in \mathcal{N}^B} \Phi_i^{1/\rho}(v_i) Q_i^B(\mathbf{v},\mathbf{c}) - \sum_{j \in \mathcal{N}^S} \Gamma_j^{\gamma/\rho}(c_j) Q_j^S(\mathbf{v},\mathbf{c}) \right] dG(\mathbf{c}) dF(\mathbf{v})$$
$$+ \sum_{i \in \mathcal{N}^B} (1 - \rho + \mu_i^B) \hat{u}_i^B(\underline{v}) + \sum_{j \in \mathcal{N}^S} (\gamma - \rho + \mu_j^S) \hat{u}_j^S(\overline{c}) - \gamma u.$$

If the frontier has finite slope (i.e., for $u < \overline{u}_S$), then constraint qualification is satisfied (and γ is finite) and for given ρ and γ, the Lagrangian is maximized with respect to $\mathbf{Q}$ pointwise by setting $\mathbf{Q}$ equal to $\mathbf{Q}^{\hat{\mathbf{w}}}$ defined in Lemma 1 for $\hat{\mathbf{w}}$ defined by $\hat{w}_j^S \equiv \min\{\gamma, 1\}$ for all $j \in \mathcal{N}^S$ and $\hat{w}_i^B \equiv 1/\max\{1, \gamma\}$ for all $i \in \mathcal{N}^B$ (essentially we use weight γ for all suppliers and weight 1 for all buyers, but rescaled so that the weights are all in $[0, 1]$). If the frontier has infinite slope (i.e., $u = \overline{u}_S$), then we can instead define the frontier as maximizing total expected supplier surplus subject to a lower bound on total expected buyer surplus, in which case constraint qualification is satisfied, and the analogous analysis gives $w_j^S = 1$ for all $j \in \mathcal{N}^S$ and $w_i^B = 0$ for all $i \in \mathcal{N}^B$. Thus, we conclude that for any given u the bargaining weights that maximize the sum of buyers' utilities over the set of mechanisms in $\mathcal{M}$ such that the sellers' utilities sum to at least u must be uniform across buyers and uniform across suppliers. Because one can always rescale bargaining weights, for example by the sum of the buyer weight and the supplier weight, it is without loss of generality to restrict attention to bargaining weights with buyer weight $\Delta \in [0, 1]$ and supplier weight $1 - \Delta$. ■

Proof of Proposition 4. First note that $\omega(u)$ is decreasing in u because a decrease in u relaxes a binding constraint. Thus, the frontier $\mathcal{F}$ is decreasing in (u_S, u_B) space. Turning to the question of concavity, we show that $\omega(u)$ is concave, that is, $\omega(u_\lambda) \geq \lambda\omega(u_0) + (1 - \lambda)\omega(u_1)$, where $u_\lambda = \lambda u_0 + (1 - \lambda)u_1$ and $\lambda \in [0, 1]$. Let $P(u)$ be the problem of maximizing the sum of the buyers' utilities over the set of mechanisms in $\mathcal{M}$ such that the sellers' utilities sum to at least u, which is a convex set. Denote the corresponding mechanism by $\langle \mathbf{Q}_u, \mathbf{M}_u \rangle$ and the associated sum of buyers' utilities by $U_B(u)$. Because the mechanisms $\langle \mathbf{Q}_{u_0}, \mathbf{M}_{u_0} \rangle$ and $\langle \mathbf{Q}_{u_1}, \mathbf{M}_{u_1} \rangle$ are feasible for $P(u_\lambda)$, the designer could choose the mechanism $\langle \mathbf{Q}_{u_0}, \mathbf{M}_{u_0} \rangle$ with probability λ and the mechanism $\langle \mathbf{Q}_{u_1}, \mathbf{M}_{u_1} \rangle$ with probability $1 - \lambda$, thereby generating payoffs of u_λ for the sellers and $\lambda U_B(u_0) + (1 - \lambda) U_B(u_1)$ for the buyers. Of course, the convex combination need not be optimal for $P(u_\lambda)$ and so $U_B(u_\lambda) \geq \lambda U_B(u_0) + (1 - \lambda) U_B(u_1)$. That is, $U_B(u)$ is concave.

We are left to map this back to the original problem of maximizing over weights $\mathbf{w}$ rather than mechanisms. Observe that the payoff profiles constructed in the paragraph above are

Pareto undominated. Hence, by Proposition 2 there exists weights associated with each of those such that they are the payoffs generated by the incomplete information bargaining mechanism. Thus, $\omega(u)$ is concave.

We now turn to the issue of strict concavity. Given $u \in [\underline{u}_S, \overline{u}_S]$, let $\mathbf{w}(u)$ be bargaining weights that solve $P(u)$. By Lemma 2, $\mathbf{w}^B(u)$ and $\mathbf{w}^S(u)$ are symmetric. Assume that the first-best is achieved at most at one point on the frontier, which is necessarily the point associated with all buyers and suppliers having the same bargaining weight. This implies that

$$\underline{v} \leq \overline{c}. \tag{19}$$

We have shown that the frontier is concave to the origin. If it is not strictly concave, then there exists a linear portion of the frontier. We assume that the linear portion lies in the region in which the buyers' common bargaining weight is greater than the suppliers' common bargaining weight (an analogous argument applies if it lies in the region in which the suppliers' bargaining weight is greater). Without loss of generality, let the buyers' bargaining weight be 1 and the suppliers' weight be $(1-\Delta)/\Delta$ for $\Delta \in (1/2, 1]$ (essentially, we let the buyer weight be Δ and the supplier weight be $1-\Delta$ and rescale so that the maximum bargaining weight is 1). Thus, letting $u_j^S(\Delta)$ and $u_i^B(\Delta)$ denote supplier j's and buyer i's expected payoffs as a function of Δ, respectively, there exist Δ' and Δ'' with $1/2 < \Delta'' < \Delta' < 1$ and $\lambda \in (0,1)$ such that, letting $\Delta_\lambda \equiv \lambda\Delta' + (1-\lambda)\Delta''$, we have

$$\sum_{j \in \mathcal{N}^S} u_j^S(\Delta_\lambda) = \lambda \sum_{j \in \mathcal{N}^S} u_j^S(\Delta') + (1-\lambda) \sum_{j \in \mathcal{N}^S} u_j^S(\Delta''), \tag{20}$$

and

$$\sum_{i \in \mathcal{N}^B} u_i^B(\Delta_\lambda) = \lambda \sum_{i \in \mathcal{N}^B} u_i^B(\Delta') + (1-\lambda) \sum_{i \in \mathcal{N}^B} u_i^B(\Delta''). \tag{21}$$

Denote the Lagrange multipliers on the no-deficit constraint in the incomplete information bargaining mechanism associated with Δ', Δ_λ, and Δ'' by ρ', ρ_λ, and ρ'', respectively. It follows from the assumption that the buyer weight is fixed at 1 and that we are away from the first-best that the multipliers are greater than or equal to 1 and decreasing in Δ, i.e., $1 \leq \rho' < \rho_\lambda < \rho''$, which implies that

$$\frac{1}{\rho'} > \frac{1}{\rho_\lambda} > \frac{1}{\rho''}.$$

Define $y \equiv \max\{\underline{c}, \underline{v}\}$. Then $y \in [\underline{v}, \overline{v})$ and, using (19), $y \in [\underline{c}, \overline{c}]$, which says that y is in the range of the buyers' weighted virtual values and suppliers' weighted virtual costs. We have for all $i \in \mathcal{N}^B$,

$$\Phi_i^{1/\rho'^{-1}}(y) < \Phi_i^{1/\rho_\lambda{}^{-1}}(y) < \Phi_i^{1/\rho''^{-1}}(y).$$

Thus, for

$$v_i \in \left(\Phi_i^{1/\rho'^{-1}}(y), \Phi_i^{1/\rho_\lambda{}^{-1}}(y)\right), \tag{22}$$

we have

$$\Phi_i^{1/\rho''}(v_i) < \Phi_i^{1/\rho_\lambda}(v_i) < y < \Phi_i^{1/\rho'}(v_i).$$

Using the continuity of the weighted virtual cost functions, it follows that for an open set of types $(\mathbf{v}, \mathbf{c})$ such that for all $i \in \mathcal{N}^B$, v_i satisfies (22) and for all $j \in \mathcal{N}^S$, c_j is in a sufficiently tight neighborhood of $\Gamma_j^{\frac{1-\Delta'}{\Delta'}/\rho'^{-1}}(y)$, we have for all $i \in \mathcal{N}^B$ and $j \in \mathcal{N}^S$,

$$\Phi_i^{1/\rho''}(v_i) < \Phi_i^{1/\rho_\lambda}(v_i) < \Gamma_j^{\frac{1-\Delta'}{\Delta'}/\rho'}(c_j) < \Phi_i^{1/\rho'}(v_i). \tag{23}$$

Recalling from Proposition 1 the incomplete information bargaining allocation rule and letting $\mathbf{Q}^\Delta$ denote the unique incomplete information bargaining allocation rule associated with Δ, (23) implies that for an open set of type, $\mathbf{Q}^{\Delta''}$ and $\mathbf{Q}^{\Delta_\lambda}$ specify no trade, while $\mathbf{Q}^{\Delta'}$ has trade. Hence, the allocation rule $\mathbf{Q}^{\Delta_\lambda}$ is *not* a convex combination of the allocation rules $\mathbf{Q}^{\Delta'}$ and $\mathbf{Q}^{\Delta''}$. Because a convex combination of $\mathbf{Q}^{\Delta'}$ and $\mathbf{Q}^{\Delta''}$ implies a convex combination of the payoffs, it follows that $\mathbf{Q}^{\Delta_\lambda}$ does not induce a convex combination of the payoffs, a contradiction. Hence, the Williams frontier must be strictly concave.

If the Williams frontier coincides with the first-best frontier for Δ' and Δ'' with $\Delta' < \Delta''$, then the mechanism that is a convex combination of the mechanism corresponding to Δ' and the mechanism corresponding to Δ'' also achieves the first-best. By the definition of the first-best, no mechanism achieves greater social surplus, so for $\Delta \in (\Delta', \Delta'')$, the frontier must be linear, coinciding with the first-best frontier. ■

Proof of Lemma 3. Consider a merger of suppliers 1 and 2. Let $\langle \hat{\mathbf{Q}}, \hat{\mathbf{M}} \rangle$ be the post-merger incomplete information bargaining mechanism. Construct a pre-merger mechanism $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle$ that mimics the allocation rule of the post-merger mechanism as follows: define the allocation rule $\tilde{\mathbf{Q}}$ such that for supplier $j \in \mathcal{N}^S \backslash \{1, 2\}$ and buyer $i \in \mathcal{N}^B$,

$$\tilde{Q}_j^S(\mathbf{v}, \mathbf{c}) \equiv \hat{Q}_j^S(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}}) \text{ and } \tilde{Q}_i^B(\mathbf{v}, \mathbf{c}) \equiv \hat{Q}_i^B(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}}).$$

By the assumption that $\Pr_{c_i, c_j}(h(c_i, c_j) \leq z) = G_{i,j}(z)$, the interim expected allocations of the nonmerging agents are the same under $\hat{\mathbf{Q}}$ and $\tilde{\mathbf{Q}}$, and so their expected thresholds payments are the same as well.

For supplier 1, define the allocation rule by

$$\begin{aligned}\tilde{Q}_1^S(\mathbf{v},\mathbf{c}) \;&\equiv\; \min\left\{k_1^S, \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1,c_2), \mathbf{c}_{-\{1,2\}})\right\}\cdot \mathbf{1}_{c_1\le c_2}\\ &\quad +\max\left\{0, \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1,c_2), \mathbf{c}_{-\{1,2\}}) - k_2^S\right\}\cdot \mathbf{1}_{c_1> c_2},\end{aligned}$$

and for supplier 2, define

$$\begin{aligned}\tilde{Q}_2^S(\mathbf{v},\mathbf{c}) \;&\equiv\; \min\left\{k_2^S, \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1,c_2), \mathbf{c}_{-\{1,2\}})\right\}\cdot \mathbf{1}_{c_2\le c_1}\\ &\quad +\max\left\{0, \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1,c_2), \mathbf{c}_{-\{1,2\}}) - k_1^S\right\}\cdot \mathbf{1}_{c_2> c_1},\end{aligned}$$

which implies that

$$\tilde{Q}_1^S(\mathbf{v},\mathbf{c}) + \tilde{Q}_2^S(\mathbf{v},\mathbf{c}) \equiv \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1,c_2), \mathbf{c}_{-\{1,2\}}).$$

This mechanism is incentive compatible—the assumptions on h and the incentive compatibility of $\tilde{\mathbf{Q}}$ imply that $\tilde{Q}_1^S$ is nonincreasing in c_1 and $\tilde{Q}_2^S$ is nonincreasing in c_2.

We now show that the sum of the threshold types of suppliers 1 and 2 under $\tilde{\mathbf{Q}}$ are less than or equal to the threshold types of the merged entity under $\hat{\mathbf{Q}}$.

Suppose that $(\mathbf{v},\mathbf{c})$ is such that supplier 1 trades under $\tilde{\mathbf{Q}}$.

Case 1: $c_1 > c_2$. Because 1 trades even though it has the higher cost, it must be that $k_{1,2}^S > k_2^S$. It then follows by condition (iv) of the lemma that $c_1 \le h(c_1,c_2)$. Consider increasing supplier 1's report x above c_1. As x increases, $h(x,c_2)$ continuously weakly increases and so $\hat{Q}_{1,2}^S(\mathbf{v}, h(x,c_2), \mathbf{c}_{-\{1,2\}})$ weakly decreases, as does $\tilde{Q}_1^S(\mathbf{v}, x, \mathbf{c}_{-1})$. So supplier 1's threshold types are costs x such that $\hat{Q}_{1,2}^S$ decreases at $h(x,c_2)$, implying that supplier 1's threshold types for its q units are greater than or equal to c_1 and less than or equal to the threshold types for the merged entity's q units with the lowest threshold types.

Case 2: $c_1 \le c_2$. Then $c_1 \le h(c_1,c_2)$. As in case 1, as supplier 1's report x increases above c_1, $h(x,c_2)$ continuously weakly increases and so $\hat{Q}_{1,2}^S(\mathbf{v}, h(x,c_2), \mathbf{c}_{-\{1,2\}})$ weakly decreases, as does $\tilde{Q}_1^S(\mathbf{v}, x, \mathbf{c}_{-1})$. As supplier 1's report increases above c_2, supplier 1's quantity decreases from $\min\left\{k_1^S, \hat{Q}_{1,2}^S(\mathbf{v}, h(x^*,c_2), \mathbf{c}_{-\{1,2\}})\right\}$ to $\max\left\{0, \hat{Q}_{1,2}^S(\mathbf{v}, h(x^*,c_2), \mathbf{c}_{-\{1,2\}}) - k_2^S\right\}$, which is weakly smaller. It follows that supplier 1's threshold types for its q units are greater than or equal to c_1 and less than or equal to the threshold types for the merged entity's q units with the lowest threshold types.

A similar argument applies to supplier 2. Thus, the sum of the expected threshold payments of suppliers 1 and 2 is weakly less than the expected threshold payment of the merged entity, and individual rationality is satisfied for suppliers 1 and 2 under their threshold payments.

It then follows that $\tilde{\mathbf{Q}}$, augmented with a payment rule based on threshold payments and the apportionment of the expected budget surplus through fixed payments, is an IC, IR, no-deficit mechanism and has weakly greater expected weighted surplus than $\langle \hat{\mathbf{Q}}, \hat{\mathbf{M}} \rangle$ does in the post-merger market. Optimizing for the pre-merger market reinforces the result. ■

Proof of Proposition 5. Assume that $k_1^S = k_2^S = K^B$ and $w_1^S = w_2^S = w$, and consider a merger of suppliers 1 and 2, where w is also the bargaining weight of the merged entity. (Analogous analysis applies to a merger of buyers 1 and 2 with $k_1^B = k_2^B = K^S$ and $w_1^B = w_2^B = w_{1,2}^B = w$.) Let $\mathcal{N}^B$ and $\mathcal{N}^S$ denote the set of pre-merger buyers and suppliers, respectively. The merged entity draws its constant marginal cost $c_{1,2}$ for up to K^B units from $G_{1,2}(c) \equiv 1-(1-G_1(c))(1-G_2(c))$, which is the distribution of $\min\{c_1, c_2\}$. Denote the associated density by $g_{1,2}$ and weighted virtual cost function by $\Gamma_{1,2}^a(x) \equiv x + (1-a)\frac{G_{1,2}(x)}{g_{1,2}(x)}$.

Let $\langle \hat{\mathbf{Q}}, \hat{\mathbf{M}} \rangle$ be the incomplete information bargaining mechanism in the post-merger market following the merger of suppliers 1 and 2, and let $\hat{\rho}$ be the associated Lagrange multiplier and $\hat{\boldsymbol{\eta}}$ be the shares. As described in Lemma 1, $\hat{\mathbf{Q}}$ allocates trades to the buyers with the greatest weighted virtual values, $\Phi_i^{w_i^B/\hat{\rho}}(v_i)$ for $i \in \mathcal{N}^B$, and the suppliers with the smallest weighted virtual costs, $\Gamma_{1,2}^{w/\hat{\rho}}(c_{1,2})$ and for $j \in \mathcal{N}^S \backslash \{1,2\}$, $\Gamma_j^{w_j^S/\hat{\rho}}(c_j)$, and has the greatest number of trades such that the cutoff weighted virtual cost is less than or equal to the cutoff weighted virtual value.

Let $\hat{\pi}$ denote the expected budget surplus for the post-merger mechanism, not including fixed payments:

$$\hat{\pi} \equiv \mathbb{E}_{\mathbf{v},c_{1,2},\mathbf{c}_{-\{1,2\}}}\left[\sum_{i\in\mathcal{N}^B} \Phi_i(v_i)\hat{Q}_i^B - \sum_{j\in\mathcal{N}^S\backslash\{1,2\}} \Gamma_j(c_j)\hat{Q}_j^S - \Gamma_{1,2}(c_{1,2})\hat{Q}_{1,2}^S\right],$$

where we drop the argument $(\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}})$ on the allocation rule.

By the payoff equivalence theorem, we can, without loss, focus on payment rules based on threshold payments, which are the sum of an agent's threshold types for each unit traded, where the threshold type for a unit is the worst type (lowest value for a buyer and highest cost for a supplier) that the agent could report and still trade that unit. Specifically, for each buyer $i \in \mathcal{N}^B$, its payment to the mechanism $\hat{M}_i^B(\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}})$ is the sum of the threshold types for each unit that it trades (and zero if it does not trade) minus its fixed payment $\hat{\eta}_i^B\hat{\pi}$. For each supplier $i \in \mathcal{N}^S\backslash\{1,2\}$, its payment from the mechanism $\hat{M}_i^S(\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}})$ is the sum of the threshold types for each unit that it trades (and zero if it does not trade) plus its fixed payment $\hat{\eta}_i^S\hat{\pi}$. For the merged entity, its payment from the mechanism $\hat{M}_{1,2}^S(\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}})$

is the sum of the threshold types for each unit that it trades (and zero if it does not trade) plus its fixed payment $\hat{\eta}^S_{1,2}\hat{\pi}$.

Next, we apply $\langle \hat{\mathbf{Q}}, \hat{\mathbf{M}} \rangle$ to the pre-merger market by defining a pre-merger mechanism $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle$ that mimics the allocation rule of the post-merger mechanism and has threshold payments. Specifically, given reports for all of the pre-merger agents $(\mathbf{v}, \mathbf{c})$, define the allocation rule for supplier $j \in \mathcal{N}^S \backslash \{1, 2\}$ by

$$\tilde{Q}^S_j(\mathbf{v}, \mathbf{c}) \equiv \hat{Q}^S_j(\mathbf{v}, \min\{c_1, c_2\}, \mathbf{c}_{-\{1,2\}}),$$

and define buyer i's allocation rule by

$$\tilde{Q}^B_i(\mathbf{v}, \mathbf{c}) \equiv \hat{Q}^B_i(\mathbf{v}, \min\{c_1, c_2\}, \mathbf{c}_{-\{1,2\}}).$$

For suppliers 1 and 2, define the allocation rule by

$$\tilde{Q}^S_1(\mathbf{v}, \mathbf{c}) \equiv \begin{cases} \hat{Q}^S_{1,2}(\mathbf{v}, \min\{c_1, c_2\}, \mathbf{c}_{-\{1,2\}}) & \text{if } c_1 \leq c_2, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\tilde{Q}^S_2(\mathbf{v}, \mathbf{c}) \equiv \begin{cases} \hat{Q}^S_{1,2}(\mathbf{v}, \min\{c_1, c_2\}, \mathbf{c}_{-\{1,2\}}) & \text{if } c_2 < c_1, \\ 0 & \text{otherwise,} \end{cases}$$

which are nonincreasing in the supplier 1's type and supplier 2's type, respectively, and so satisfy incentive compatibility.

We now show that the expected threshold types of the nonmerging agents are the same under $\hat{\mathbf{Q}}$ and $\tilde{\mathbf{Q}}$.

Lemma B.1. *The expected threshold types of trading buyers and trading nonmerging suppliers are the same under $\tilde{\mathbf{Q}}$ and $\hat{\mathbf{Q}}$.*

Proof. Because the merged entity's type is drawn from the distribution of $\min\{c_1, c_2\}$, it suffices to compare threshold types between the two mechanisms for a given pre-merger type vector $(\mathbf{v}, \mathbf{c})$ and the corresponding post-merger type vector $(\mathbf{v}, \min\{c_1, c_2\}, \mathbf{c}_{-\{1,2\}})$. The threshold types for a trading buyer depend at most on the weighted virtual cost of the cutoff supplier (the highest weighted virtual cost supplier that trades) and the weighted virtual values of buyers with units that do not trade. Analogously, the threshold types for a trading supplier depend at most on the weighted virtual value of the cutoff buyer (the lowest weighted virtual value buyer that trades) and the weighted virtual costs of suppliers with units that do not trade.

First, consider the threshold types for trading buyers. It is never the case that both supplier 1 and supplier 2 trade under $\tilde{\mathbf{Q}}$. If supplier 1 or supplier 2 trades, then using $k_1^S = k_2^S = K^B$, the cutoff weighted virtual cost is $\Gamma_{1,2}^{w/\hat{\rho}}(\min\{c_1, c_2\})$. Thus, trading buyers' threshold types depend on the merging suppliers only through $\Gamma_{1,2}^{w/\hat{\rho}}(\min\{c_1, c_2\})$, which is the same as under $\hat{\mathbf{Q}}$. If neither supplier 1 nor supplier 2 trades under $\tilde{\mathbf{Q}}$, then the merged supplier with type equal to $\min\{c_1, c_2\}$ does not trade under $\hat{\mathbf{Q}}$, and again the buyer's threshold payments are the same under $\hat{\mathbf{Q}}$ and $\tilde{\mathbf{Q}}$. This completes the demonstration that trading buyers' expected threshold types are the same under $\tilde{\mathbf{Q}}$ and $\hat{\mathbf{Q}}$.

Second, consider the threshold types for trading nonmerging suppliers. Suppose that nonmerging supplier i trades under $\tilde{\mathbf{Q}}$. Then because $k_1^S = k_2^S = K^B$, it must be that supplier i has a lower weighted virtual cost than $\Gamma_{1,2}^{w/\hat{\rho}}(\min\{c_1, c_2\})$, i.e., ignoring ties between types,

$$\Gamma_i^{w_i^S/\hat{\rho}}(c_i) < \Gamma_{1,2}^{w/\hat{\rho}}(\min\{c_1, c_2\}).$$

If supplier i were to report a type greater than $\hat{x} \equiv \Gamma_i^{w_i^S/\hat{\rho}^{-1}}(\Gamma_{1,2}^{w/\hat{\rho}}(\min\{c_1, c_2\}))$, it would trade zero units, so its threshold types are all less than or equal to $\hat{x}$. This implies that its threshold types only depend on c_1 and c_2 through $\min\{c_1, c_2\}$. Thus, nonmerging suppliers' expected threshold payments are the same under $\tilde{\mathbf{Q}}$ and $\hat{\mathbf{Q}}$. $\square$

Using Lemma B.1 and defining the payments $\tilde{\mathbf{M}}$ for the nonmerging agents to be their threshold payments associated with $\tilde{\mathbf{Q}}$ minus $\hat{\eta}_i^B \hat{\pi}$ for buyer $i \in \mathcal{N}^B$ and plus $\hat{\eta}_j^S \hat{\pi}$ for supplier $j \in \mathcal{N}^S \backslash \{1, 2\}$, it follows that the buyers' and nonmerging suppliers' expected payments are the same under $\tilde{\mathbf{M}}$ as under $\hat{\mathbf{M}}$.

Now consider the payments of the merging suppliers. Suppose supplier 1 trades, which implies that supplier 2 does not trade. Consider supplier 1's threshold payment for the q-th unit under $\tilde{\mathbf{Q}}$. It is the worst type that supplier 1 could report and still trade its q-th unit. The only difference in the calculation of supplier 1's threshold payment under $\tilde{\mathbf{Q}}$ versus the merged entity's threshold payment under $\hat{\mathbf{Q}}$ is that under $\tilde{\mathbf{Q}}$, supplier 1's threshold types are bounded above by c_2 because then any report greater than c_2 results in supplier 1's trading zero units. Thus, when supplier 1 trades (and supplier 2 does not), its expected threshold payment under $\tilde{\mathbf{Q}}$ is strictly less than the expected threshold payment of the merged entity under $\hat{\mathbf{Q}}$. Similarly, when supplier 2 trades, its expected threshold payment under $\tilde{\mathbf{Q}}$ is less than the expected threshold payment of the merged entity under $\hat{\mathbf{Q}}$. Thus, letting $\tilde{\tau}_j^S(\mathbf{v}, \mathbf{c})$ be the threshold payment of supplier j under $\tilde{\mathbf{Q}}$ and $\hat{\tau}^S(\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}})$ be the threshold payment of the merged entity under $\hat{\mathbf{Q}}$, we have

$$0 < \mathbb{E}_{\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}}}\left[\hat{\tau}^S(\mathbf{v}, c_{1,2}, \mathbf{c}_{-\{1,2\}})\right] - \mathbb{E}_{\mathbf{v}, \mathbf{c}}\left[\tilde{\tau}_1^S(\mathbf{v}, \mathbf{c}) + \tilde{\tau}_2^S(\mathbf{v}, \mathbf{c})\right] \equiv \Delta.$$

This implies that under $\tilde{\mathbf{Q}}$, the budget surplus in the pre-merger market not including fixed payments is $\hat{\pi} + \Delta$, where Δ is the amount by which the merging suppliers' combined threshold payments are smaller in the pre-merger market under $\tilde{\mathbf{Q}}$ than in the post-merger market under $\hat{\mathbf{Q}}$.

Let $\tilde{\boldsymbol{\eta}}$ be the pre-merger shares (recall that we assume that the merger does not alter shares, so for nonmerging agents, the shares in $\tilde{\boldsymbol{\eta}}$ are the same as in $\hat{\boldsymbol{\eta}}$, and $\tilde{\eta}_1^S + \tilde{\eta}_2^S = \hat{\eta}_{1,2}^S$). Define for each supplier $j \in \{1,2\}$,

$$\tilde{M}_j^S(\mathbf{v},\mathbf{c}) \equiv \tilde{\tau}_j^S(\mathbf{v},\mathbf{c}) + \tilde{\eta}_j^S \hat{\pi},$$

and note that

$$\begin{aligned}
\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\tilde{M}_1^S(\mathbf{v},\mathbf{c}) + \tilde{M}_2^S(\mathbf{v},\mathbf{c})\right] &= \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\tilde{\tau}_1^S(\mathbf{v},\mathbf{c}) + \tilde{\tau}_2^S(\mathbf{v},\mathbf{c}) + \hat{\eta}_{1,2}^S\hat{\pi}\right] \\
&= \mathbb{E}_{\mathbf{v},c_{1,2},\mathbf{c}_{-\{1,2\}}}\left[\hat{\tau}_i^S(\mathbf{v},c_{1,2},\mathbf{c}_{-\{1,2\}})\right] + \hat{\eta}_{1,2}^S\hat{\pi} - \Delta \\
&= \mathbb{E}_{\mathbf{v},c_{1,2},\mathbf{c}_{-\{1,2\}}}\left[\hat{M}_{1,2}^S(\mathbf{v},c_{1,2},\mathbf{c}_{-\{1,2\}})\right] - \Delta.
\end{aligned}$$

It follows that $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle$ is an incentive compatible, individually rational pre-merger mechanism and satisfies the no-deficit constraint. In addition, there is budget surplus Δ to be allocated to the agents according to $\tilde{\boldsymbol{\eta}}$.

Comparing expected weighted surpluses, we have (dropping the arguments $(v, c_{1,2}, \mathbf{c}_{-\{1,2\}})$ on $\langle\hat{\mathbf{Q}},\hat{\mathbf{M}}\rangle$ and $(\mathbf{v},\mathbf{c})$ on $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle$):

$$\begin{aligned}
&\mathbb{E}_{\mathbf{v},c_{1,2},\mathbf{c}_{-\{1,2\}}}\left[\sum_{i\in\mathcal{N}^B} w_i^B\left(\hat{Q}_i^B v_i - \hat{M}_i^B\right) + \sum_{j\in\mathcal{N}^S\setminus\{1,2\}} w_j^S\left(\hat{M}_j^S - \hat{Q}_j^S c_j\right) + w\left(\hat{M}_{1,2}^S - \hat{Q}_{1,2}^S c_{1,2}\right)\right] \\
= \;&\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{i\in\mathcal{N}^B} w_i^B\left(\tilde{Q}_i^B v_i - \tilde{M}_i^B\right) + \sum_{j\in\mathcal{N}^S\setminus\{1,2\}} w_j^S\left(\tilde{M}_j^S - \tilde{Q}_j^S c_j\right) + \sum_{j\in\{1,2\}} w\left(\tilde{M}_j^S + \Delta - \tilde{Q}_j^S c_j\right)\right] \\
\leq \;&\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{i\in\mathcal{N}^B} w_i^B\left(\tilde{Q}_i^B v_i - \tilde{M}_i^B + \tilde{\eta}_i^B\Delta\right) + \sum_{j\in\mathcal{N}^S\setminus\{1,2\}} w_j^S\left(\tilde{M}_j^S - \tilde{Q}_j^S c_j + \tilde{\eta}_j^S\Delta\right)\right. \\
&\left.+ \sum_{j\in\{1,2\}} w\left(\tilde{M}_j^S - \tilde{Q}_j^S c_j + \tilde{\eta}_j^S\Delta\right)\right],
\end{aligned}$$

where the inequality uses the fact that $\tilde{\eta}_i^x > 0$ only if $w_i^x = \max\mathbf{w}$ and $\sum\tilde{\boldsymbol{\eta}} = 1$, which implies that

$$\sum_{i\in\mathcal{N}^B} w_i^B\tilde{\eta}_i^B + \sum_{j\in\mathcal{N}^S\setminus\{1,2\}} w_j^S\tilde{\eta}_j^S + \sum_{j\in\{1,2\}} w\tilde{\eta}_j^S = \max\mathbf{w} \geq w,$$

and where the inequality is strict if $w < \max \mathbf{w}$.

Thus, $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle$ is a feasible incomplete information bargaining mechanism in the pre-merger market and generates expected weighted surplus under that is weakly greater (strictly if $w < \max \mathbf{w}$) than under $\langle \hat{\mathbf{Q}}, \hat{\mathbf{M}} \rangle$ in the post-merger market. It follows that the optimized incomplete information bargaining mechanism in the pre-merger market has weakly greater weighted social surplus (strictly if $w < \max \mathbf{w}$) than $\langle \hat{\mathbf{Q}}, \hat{\mathbf{M}} \rangle$ does in the post-merger market. Further, if all nonmerging agents have zero bargaining weight, then $\hat{\rho} = \max \mathbf{w} = w$ and the pre-merger mechanism has $\tilde{\rho} = \max \mathbf{w} = w$. It follows that no further optimization of the mechanism is possible, and so expected weighted welfare, and indeed, expected surplus for all agents, is the same before and after the merger. ■

Proof of Proposition 6. A supplier merger of the type considered results in a merged entity with a weighted virtual cost function

$$\Gamma^a_{i,j}(c) \equiv c + (1-a)\frac{1-(1-G_i(c))(1-G_j(c))}{g_i(c)(1-G_j(c)) + g_j(c)(1-G_i(c))},$$

which has $\Gamma^a_{i,j}(\overline{c}) = \infty$ for all $a \in [0,1)$.

Part (i): Suppose that the pre-merger market is efficient. Let $\hat{\rho}^{\mathbf{w}}$ denote the post-merger Lagrange multiplier on the no-deficit constraint and note that $\max \mathbf{w} \leq \hat{\rho}^{\mathbf{w}}$. When $w^S_{i,j} < \max \mathbf{w}$, we have $\Gamma^{w^S_{i,j}/\hat{\rho}^{\mathbf{w}}}_{i,j}(\overline{c}) = \infty$. This implies that the post-merger market does not achieve the first-best because for an open set of types with $c_{i,j}$ sufficiently close to $\overline{c}$ such that $\overline{v} < \Gamma^{w^S_{i,j}/\rho^{\mathbf{w}}}_{i,j}(c_{i,j})$, we have $c_{i,j} < \min_{\ell \in \mathcal{N}^B} v_\ell$ and $c_{i,j} < \min_{\ell \in \mathcal{N}^S \setminus \{i,j\}} c_\ell$, which implies that the merged entity trades under the first-best, but $\max_{\ell \in \mathcal{N}^B} \Phi^{w^B_\ell/\rho^{\mathbf{w}}}_\ell(v_\ell) \leq \overline{v} < \Gamma^{w^S_{i,j}/\rho^{\mathbf{w}}}_{i,j}(c_{i,j})$, which implies that the merged entity does not trade under incomplete information bargaining. Thus, expected social surplus decreases as a result of the merger.

Part (ii): Suppose that the pre-merger market is not efficient. Proposition 5 implies that if the merging suppliers have all the bargaining weight, then the merger does not affect social surplus, positively or negatively. In contrast, if the buyers or a subset of buyers have all the bargaining weight, then the merger has no effect on the Lagrange multiplier on the no-deficit constraint because the buyer-optimal mechanism satisfies the no-deficit constraint when the multiplier is equal to its minimum value of $\max \mathbf{w}$. Thus, for a merger that does not alter bargaining weights, the only effect on the allocation rule comes through the effect on the merged entity's virtual cost function. Thus, a merger reduces social surplus if the merged entity's weighted virtual cost function $\Gamma^a_{1,2}$ satisfies $\Gamma^a_{1,2}(\min\{c_1, c_2\}) > \min\{\Gamma^a_1(c_1), \Gamma^a_2(c_2)\}$ for all $c_1, c_2 \in (\underline{c}, \overline{c})$ and $a \in [0,1]$, which holds for symmetric suppliers, because then trade occurs for a strictly smaller set of realized types in the post-merger versus per-merger market.

■

Proof of Proposition 8. We begin by considering the first part of the statement. Choosing $\underline{v}$ such that $\underline{v} \geq \bar{c}$, we have nonoverlapping supports. Hence, with symmetric bargaining weights, the pre-integration market achieves the first-best. Consequently, vertical integration cannot increase social surplus. The second part follows from Theorem 4 (and Table 1) in Williams (1999), which shows that $\underline{v} > \underline{c}$ and n^S sufficiently large is sufficient for first-best to be possible when the suppliers draw their types from identical distributions. ■

Proof of Proposition 9. If $\underline{v} = 0$, then the first-best is not possible without vertical integration; for $n^B = n^S = 1$, see Myerson and Satterthwaite (1983), and for $n^B = 1 < n^S$, see Williams (1999), whose results imply that for a necessary condition for first-best to be possible is that $\underline{c}$ is strictly smaller than the lower bound of the support of the buyer's distribution. Because the suppliers are assumed symmetric, we drop the supplier subscripts on the distribution G and on the unweighted and weighted virtual cost functions Γ and Γ^a. Let $L_{n^S}(c) \equiv 1-(1-G(c))^{n^S}$ with density $l_{n^S}(c)$ denote the distribution of the lowest cost draw of a seller. Because we assume 1 buyer, we drop the buyer subscript on the distribution F and on the unweighted and weighted virtual value functions Φ and Φ^a.

For $\underline{v} \in [0,1)$, define the truncated distribution and density

$$F_{\underline{v}}(v) \equiv \frac{F(v)-F(\underline{v})}{1-F(\underline{v})} \quad \text{and} \quad f_{\underline{v}}(v) \equiv \frac{f(v)}{1-F(\underline{v})}.$$

Observe that for $\underline{v}' > \underline{v}$, we have $F_{\underline{v}'}(v) \leq F_{\underline{v}}(v)$, with a strict inequality for $v \in (\underline{v}, 1)$, that is, $F_{\underline{v}'}$ first-order stochastically dominates $F_{\underline{v}}$. Define the weighted virtual value function associated with the truncated distribution by, for $a \in [0,1]$ and $v \in [\underline{v}, 1]$,

$$\Phi^a_{\underline{v}}(v) \equiv v-(1-a)\frac{1-F_{\underline{v}}(v)}{f_{\underline{v}}(v)} = v-(1-a)\frac{1-F(v)}{f(v)} = \Phi^a(v),$$

which reflects the well-known truncation-invariance result for virtual value functions.

Analogous to how we define $\pi^{\mathbf{w}}$ in (6), given $\underline{v}$ and ρ, define $\pi(\underline{v},\rho)$ to be the budget surplus of the mechanism with the allocation rule that solves (5) for $\mathbf{w} = \mathbf{1}$, not including fixed payments:

$$\pi(\underline{v},\rho) \equiv \begin{cases} \frac{1}{1-F(\underline{v})} \displaystyle\int_0^{\Gamma^{1/\rho^{-1}}(1)} \int_{\Phi^{1/\rho^{-1}}(\Gamma^{1/\rho}(c))}^{1} (\Phi(v)-\Gamma(c))f(v)l_n(c)dvdc & \text{if } \underline{v} \in [0, \Phi^{1/\rho^{-1}}(0)], \\ \displaystyle\int_0^{\Gamma^{1/\rho^{-1}}(1)} \int_{\max\{\underline{v},\Phi^{1/\rho^{-1}}(\Gamma^{1/\rho}(c))\}}^{1} (\Phi(v)-\Gamma(c))f_{\underline{v}}(v)l_n(c)dvdc & \text{if } \underline{v} \in [\Phi^{1/\rho^{-1}}(0), 1). \end{cases}$$

Let $\rho_{\underline{v}}$ be such the smallest $\rho \in [1,\infty)$ such that $\pi(\underline{v},\rho) \geq 0$. Because the first-best is not achieved when $\underline{v}=0$, it follows that $\rho_0 > 1$. For all $\underline{v} \in [0, \Phi^{1/\rho_0^{-1}}(0)]$, the sign of $\pi(\underline{v},\rho)$ does not depend on $\underline{v}$, so we have $\rho_{\underline{v}} = \rho_0$. In contrast, for $\underline{v} \in (\Phi^{1/\rho_0^{-1}}(0), 1)$, $\pi(\underline{v},\rho)$ is increasing in $\underline{v}$ because an increase in $\underline{v}$ induces a first-order stochastic dominance shift in $F_{\underline{v}}$ and because $\Phi(v)$ increases with v. Because $\pi(\underline{v},\rho)$ is increasing in ρ (because $\Gamma^{1/\rho^{-1}}(1)$ is decreasing in ρ and $\Phi^{1/\rho^{-1}}(\Gamma^{1/\rho}(c))$ is increasing in ρ), it follows that $\rho_{\underline{v}}$ is strictly decreasing in $\underline{v}$ for $\underline{v} \in (\Phi^{1/\rho_0^{-1}}(0), 1)$ provided that $\rho_{\underline{v}} > 1$. Thus, we have,

$$\frac{\partial \rho_{\underline{v}}}{\partial \underline{v}} \leq 0, \tag{24}$$

with a strict inequality if and only if $\underline{v} \in (\Phi^{1/\rho_0^{-1}}(0), 1)$ and $\rho_{\underline{v}} > 1$.

First-best social surplus given $\underline{v}$ satisfies

$$W^{FB}(\underline{v}) = \frac{1}{1-F(\underline{v})} \int_{\underline{v}}^{1} \int_0^{v} (v-c)l_n(c)f(v)dcdv,$$

and second-best social surplus given $\underline{v}$ satisfies

$$W^{SB}(\underline{v}) = \begin{cases} \frac{1}{1-F(\underline{v})} \displaystyle\int_{\Phi^{1/\rho_0^{-1}}(0)}^{1} \int_0^{\Gamma^{1/\rho_0^{-1}}(\Phi^{1/\rho_0}(v))} (v-c)l_n(c)f(v)dcdv & \text{if } \underline{v} \in [0, \Phi^{1/\rho_0^{-1}}(0)], \\ \displaystyle\int_{\underline{v}}^{1} \int_0^{\Gamma^{1/\rho_{\underline{v}}^{-1}}(\Phi^{1/\rho_{\underline{v}}}(v))} (v-c)l_n(c)f(v)dcdv & \text{if } \underline{v} \in [\Phi^{1/\rho_0^{-1}}(0), 1). \end{cases}$$

Turning to the ratio $\mathcal{G}(\underline{v}) = \frac{W^{FB}(\underline{v})-W^{SB}(\underline{v})}{W^{FB}(\underline{v})}$, for $\underline{v} \in [0, \Phi^{1/\rho_0^{-1}}(0)]$, we have

$$\mathcal{G}(\underline{v}) = 1 - \frac{\displaystyle\int_{\Phi^{1/\rho_0^{-1}}(0)}^{1} \int_0^{\Gamma^{1/\rho_0^{-1}}(\Phi^{1/\rho_0}(v))} (v-c)l_n(c)f(v)dcdv}{\displaystyle\int_{\underline{v}}^{1} \int_0^{v} (v-c)l_n(c)f(v)dcdv},$$

and so for $\underline{v} \in (0, \Phi^{1/\rho_0}{}^{-1}(0)]$,

$$\mathcal{G}'(\underline{v}) = -\frac{\left(\int_{\Phi^{1/\rho_0}{}^{-1}(0)}^{1} \int_0^{\Gamma^{1/\rho_0}{}^{-1}(\Phi^{1/\rho_0}(v))} (v-c)l_n(c)f(v)dcdv\right)\left(\int_0^{\underline{v}} (\underline{v}-c)f(\underline{v})l_n(c)dc\right)}{\left(\int_{\underline{v}}^1 \int_0^v (v-c)l_n(c)f(v)dcdv\right)^2} < 0,$$

which is the desired result.

Consider now the case with $\underline{v} \in (\Phi^{1/\rho_0}{}^{-1}(0), 1)$. Observe first that because $\Phi^{1/\rho}$ and $\Gamma^{1/\rho^{-1}}$ are both decreasing in ρ, it follows that

$$\frac{\partial \Gamma^{1/\rho^{-1}}(\Phi^{1/\rho}(v))}{\partial \rho} < 0. \tag{25}$$

Continuing, for $\underline{v} \in (\Phi^{1/\rho_0}{}^{-1}(0), 1)$, we have $\mathcal{G}(\underline{v}) = 1 - \frac{h(\underline{v})}{k(\underline{v})}$, where

$$h(\underline{v}) \equiv \int_{\underline{v}}^1 \int_0^{\Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v))} (v-c)l_n(c)f(v)dcdv$$

and

$$k(\underline{v}) \equiv \int_{\underline{v}}^1 \int_0^v (v-c)l_n(c)f(v)dcdv.$$

Thus, for $\underline{v} \in (\Phi^{1/\rho_0}{}^{-1}(0), 1)$, the sign of $\mathcal{G}'(\underline{v})$ is equal to the sign of $h(\underline{v})k'(\underline{v}) - h'(\underline{v})k(\underline{v})$. Using $0 < h(\underline{v}) < k(\underline{v})$, a sufficient condition for $\mathcal{G}'(\underline{v}) < 0$ is that $k'(\underline{v}) < h'(\underline{v})$. To see that this holds, first note that if $\rho_{\underline{v}} > 1$, then for all $v \in [\underline{v}, 1]$,

$$\Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v)) < v, \tag{26}$$

and second notice that if $\rho_{\underline{v}} > 1$, then

$$\begin{aligned} h'(\underline{v}) - k'(\underline{v}) &= \int_0^{\underline{v}} (\underline{v}-c)f(\underline{v})l_n(c)dc - \int_0^{\Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(\underline{v}))} (\underline{v}-c)l_n(c)f(\underline{v})dc \\ &\quad + \int_{\underline{v}}^1 (v - \Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v)))l_n(\Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v)))f(v)\frac{\partial \Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v))}{\partial \rho_{\underline{v}}}\frac{\partial \rho_{\underline{v}}}{\partial \underline{v}}dv \\ &> \int_{\underline{v}}^1 (v - \Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v)))l_n(\Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v)))f(v)\frac{\partial \Gamma^{1/\rho_{\underline{v}}}{}^{-1}(\Phi^{1/\rho_{\underline{v}}}(v))}{\partial \rho_{\underline{v}}}\frac{\partial \rho_{\underline{v}}}{\partial \underline{v}}dv \\ &> 0 \end{aligned}$$

where the first inequality uses $\rho_{\underline{v}} > 1$ and (26), and the second inequality uses $\rho_{\underline{v}} > 1$, (24), (25), and (26), completing the proof that for all $\underline{v} \in (\Phi^{1/\rho_0}{}^{-1}(0), 1)$ such that $\rho_{\underline{v}} > 1$, $\mathcal{G}'(\underline{v}) < 0$. ■

Proof of Proposition 10. We have proved the first part in the text and are thus left to prove the second part.

Let $\hat{u}^B_{i,\mathbf{Q}}(v_i; \mathbf{e}^B_{-i}, \mathbf{e}^S)$ denote the interim expected payoff of buyer i with type v_i, not including the (constant) interim expected payment to the worst-off type and not including investment costs, when the allocation rule is $\mathbf{Q}$ and other agents investments are $(\mathbf{e}^B_{-i}, \mathbf{e}^S)$. Define $\hat{u}^S_{i,\mathbf{Q}}(c_i; \mathbf{e}^B, \mathbf{e}^S_{-i})$ analogously. Let $u^B_{i,\mathbf{Q}}(\mathbf{e})$ and $u^S_{i,\mathbf{Q}}(\mathbf{e})$ denote the expected payoffs of buyer i and supplier i, respectively, when the allocation rule is $\mathbf{Q}$ and investments are $\mathbf{e}$. For any allocation rule $\mathbf{Q}$, let $q^B_i(v_i; \mathbf{e}^B_{-i}, \mathbf{e}^S) \equiv \mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}|\mathbf{e}^B_{-i},\mathbf{e}^S}\left[Q^B_i(\mathbf{v}, \mathbf{c})\right]$ and $q_i(c_i; \mathbf{e}^B, \mathbf{e}^S_{-i}) \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}_{-i}|\mathbf{e}^B,\mathbf{e}^S_{-i}}\left[Q^S_i(\mathbf{v}, \mathbf{c})\right]$. As discussed in Appendix A.1, by the payoff equivalence theorem, we have, up to a constant,

$$\hat{u}^B_{i,\mathbf{Q}}(v_i; \mathbf{e}^B_{-i}, \mathbf{e}^S) = \int_{\underline{v}}^{v_i} q^B_i(x; \mathbf{e}^B_{-i}, \mathbf{e}^S)dx, \tag{27}$$

and, taking expectations with respect to v_i, one obtains

$$u^B_{i,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{v}}^{\overline{v}} q^B_i(x; \mathbf{e}^B_{-i}, \mathbf{e}^S)(1 - F(x; e^B_i))dx \tag{28}$$

up to a constant, and, analogously,

$$u^S_{j,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{c}}^{\overline{c}} q^S_j(x; \mathbf{e}^B, \mathbf{e}^S_{-j})G_j(x; e^S_j)dx \tag{29}$$

up to a constant.

By the definition of $\overline{\mathbf{e}}$ as the vector of first-best investments, we have

$$\overline{\mathbf{e}} \in \arg\max_{\mathbf{e}} \sum_{i\in\mathcal{N}^B} u^B_{i,\mathbf{Q}^{FB}}(\mathbf{e}) + \sum_{j\in\mathcal{N}^S} u^S_{j,\mathbf{Q}^{FB}}(\mathbf{e}) - \sum_{i\in\mathcal{N}^B} \Psi^B_i(e^B_i) - \sum_{j\in\mathcal{N}^S} \Psi^S_j(e^S_j).$$

which implies that for all $i \in \mathcal{N}^B$ and $j \in \mathcal{N}^S$,

$$\overline{e}^B_i \in \arg\max_{e^B_i} u^B_{i,\mathbf{Q}^{FB}}(e^B_i, \overline{\mathbf{e}}^B_{-i}, \overline{\mathbf{e}}^S) - \Psi^B_i(e^B_i) \tag{30}$$

and

$$\overline{e}_j^S \in \arg\max_{e_j^S} u_{j,\mathbf{Q}^{FB}}^S(\overline{\mathbf{e}}^B, e_j^S, \overline{\mathbf{e}}_{-j}^S) - \Psi_j^S(e_j^S). \tag{31}$$

Assume that (9)–(10) hold. Let $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ denote the incomplete information bargaining allocation rule given in Lemma 1, but with the virtual types defined in terms of the type distributions associated with investment $\overline{\mathbf{e}}$, and let $\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}$ denote the associated multiplier on the no-deficit constraint. Suppose that first-best investments $\overline{\mathbf{e}}$ are Nash equilibrium investments, which implies that for all $i \in \mathcal{N}^B$ and $j \in \mathcal{N}^S$,

$$\overline{e}_i^B \in \arg\max_{e_i^B} u_{i,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}^B(e_i^B, \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) - \Psi_i^B(e_i^B) \tag{32}$$

and

$$\overline{e}_j^S \in \arg\max_{e_j^S} u_{j,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}^S(\overline{\mathbf{e}}^B, e_j^S, \overline{\mathbf{e}}_{-j}^S) - \Psi_j^S(e_j^S). \tag{33}$$

Assumptions (9)–(10) ensure that the first-best investments are characterized by their first-order conditions. Thus, using (28) and (30), we have for all $i \in \mathcal{N}^B$,

$$-\int_{\underline{v}}^{\overline{v}} q_i^{FB,B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx - \Psi_i^{B\prime}(\overline{e}_i^B) = 0. \tag{34}$$

Similarly, using (28) and (32), we have

$$-\int_{\underline{v}}^{\overline{v}} q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx - \Psi_i^{B\prime}(\overline{e}_i^B) = 0. \tag{35}$$

Combining (34) and (35), we have

$$\int_{\underline{v}}^{\overline{v}} (q_i^{FB,B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) - q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S)) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx = 0. \tag{36}$$

Writing this in terms of the ex post allocation rules, we have for all $i \in \mathcal{N}^B$,

$$\mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}|\overline{\mathbf{e}}_{-i}^B,\overline{\mathbf{e}}^S}\left[\int_{\underline{v}}^{\overline{v}} (Q_i^{FB,B}(x, \mathbf{v}_{-i}, \mathbf{c}) - Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x, \mathbf{v}_{-i}, \mathbf{c})) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx\right] = 0. \tag{37}$$

Steps analogous to those leading to (36) imply that for all $j \in \mathcal{N}^S$,

$$\int_{\underline{c}}^{\overline{c}} (q_j^{FB,S}(x; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-j}^S) - q_j^{\mathbf{w},\overline{\mathbf{e}},S}(x; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-j}^S)) \frac{\partial G_j(x; \overline{e}_j^S)}{\partial e} dx = 0. \tag{38}$$

By Lemma 1, we know that the total number of trades induced by $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}(\mathbf{v}, \mathbf{c})$ is the maximum such that the lowest weighted virtual value of any trading buyer is greater than or equal to the highest weighted virtual cost of any trading supplier. Further, the total number of trades induced by $\mathbf{Q}^{FB}(\mathbf{v},\mathbf{c})$ is the maximum such that the lowest value of any trading buyer is greater than or equal to the highest cost of any trading supplier. Because virtual costs are greater than or equal to actual costs and virtual values are less than or equal to actual values, it follows that $\sum_{i\in\mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v},\mathbf{c}) \le \sum_{i\in\mathcal{N}^B} Q_i^{FB,B}(\mathbf{v},\mathbf{c})$ for all $(\mathbf{v},\mathbf{c})$ (and similarly on the supply side). Because we assume that $\frac{\partial F_i(v;e)}{\partial e} < 0$ for all $v \in (\underline{v},\overline{v})$, (37) then implies that

$$\sum_{i\in\mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v},\mathbf{c}) = \sum_{i\in\mathcal{N}^B} Q_i^{FB,B}(\mathbf{v},\mathbf{c}) \equiv \xi(\mathbf{v},\mathbf{c}) \tag{39}$$

for all but a zero-measure set of types. By feasibility, the corresponding total supplier-side quantities are also equal to $\xi(\mathbf{v},\mathbf{c})$ for all but a zero-measure set of types. Thus, it only remains to show that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ always induces the same agents to trade as does $\mathbf{Q}^{FB}$.

We begin by considering the case with overlapping supports and then consider the case in which (ii), (iii), or (iv) holds.

Case 1: $\underline{v} < \overline{c}$. Suppose, contrary to what we want to show, that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ discriminates among agents based on virtual types for an open set of types—we then show that this implies that the number of trades under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ must sometimes differ from the number under the first-best, contradicting (39). That is, suppose that there exist suppliers (an analogous argument applies for buyers), which we denote by 1 and 2, and types $(\hat{\mathbf{v}},\hat{\mathbf{c}})$ with $\hat{c}_1 \neq \hat{c}_2$ such that $Q_1^{FB,S}(\hat{\mathbf{v}},\hat{\mathbf{c}}) > Q_1^{\mathbf{w},\overline{\mathbf{e}},S}(\hat{\mathbf{v}},\hat{\mathbf{c}})$ and $Q_2^{FB,S}(\hat{\mathbf{v}},\hat{\mathbf{c}}) < Q_2^{\mathbf{w},\overline{\mathbf{e}},S}(\hat{\mathbf{v}},\hat{\mathbf{c}})$. Because supplier 1 trades under the first-best when supplier 2 has excess capacity, this implies that $\hat{c}_1 < \hat{c}_2$; and because supplier 2 trades in the Nash equilibrium when supplier 1 has excess capacity, this implies that $\Gamma_2^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_2;\overline{e}_2^S) \le \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_1;\overline{e}_1^S)$. It follows that

$$\hat{c}_1 < \hat{c}_2 \le \Gamma_2^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_2;\overline{e}_2^S) \le \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_1;\overline{e}_1^S).$$

Because $\hat{c}_1 < \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_1;\overline{e}_1^S)$, it follows that $w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}} < 1$ and so for all $c \in (\underline{c},\overline{c})$, $c < \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c;\overline{e}_1^S)$. Thus, letting $\tilde{c}_1 \in (\max\{\underline{c},\underline{v}\},\overline{c})$, $\tilde{v}_1 \in (\tilde{c}_1, \min\{\overline{c}, \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\tilde{c}_1;\overline{e}_1)\})$, for all $i \in \mathcal{N}^S\backslash\{1\}$, $\tilde{c}_i = \overline{c}$, and for all $i \in \mathcal{N}^B\backslash\{1\}$, $\tilde{v}_i = \underline{v}$, we have

$$\tilde{c}_1 < \tilde{v}_1 < \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\tilde{c}_1;\overline{e}_1^S) \quad \text{and} \quad \max_{i\in\mathcal{N}^B\backslash\{1\}} \tilde{v}_i < \tilde{v}_1 < \min_{i\in\mathcal{N}^S\backslash\{1\}} \tilde{c}_i,$$

which implies that no trades occur under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ and only supplier 1 and buyer 1 trade under the first-best. By continuity, for all $(\mathbf{v},\mathbf{c})$ in an open set of types around $(\tilde{\mathbf{v}},\tilde{\mathbf{c}})$, we have $\sum_{i\in\mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v},\mathbf{c}) \neq \sum_{i\in\mathcal{N}^B} Q_i^{FB,B}(\mathbf{v},\mathbf{c})$, which contradicts (39). Thus, we conclude that

$\mathbf{Q}^{\mathbf{w},\bar{\mathbf{e}}}$ does not discriminate among suppliers based on virtual types and so $\mathbf{Q}^{\mathbf{w},\bar{\mathbf{e}}}$ induces the same suppliers to produce as does $\mathbf{Q}^{FB}$. An analogous argument shows that the set of trading buyers is the same under $\mathbf{Q}^{\mathbf{w},\bar{\mathbf{e}}}$ as under $\mathbf{Q}^{FB}$.

Case 2: $\underline{v} \geq \bar{c}$ and either (ii), (iii), or (iv) holds. Note that $\underline{v} \geq \bar{c}$ implies that under the first-best, the number of trades is $\min\{K^B, K^S\}$. If (ii) holds, i.e., $K^B = K^S$, then all agents trade under the first-best and so (39) implies that all agents also trade under $\mathbf{Q}^{\mathbf{w},\bar{\mathbf{e}}}$, which completes the proof. Suppose that (iii) holds, so that $K^B < K^S$ and (11) holds. (Analogous arguments apply to the case with $K^B > K^S$ and (12).) Then all buyers consume their full demands under the first-best. By the argument given in the proof of Lemma 2, because bargaining is efficient, we must have $w_1^S = \cdots = w_{n^S}^S$. Given this, (11) implies that the ranking of suppliers according to $\Gamma_i^{w_i^S/\rho_{\bar{\mathbf{e}}}^{\mathbf{w}}}(c_i; \bar{e}_i^S)$ is the same as the ranking according to c_i. Thus, using (39), $\mathbf{Q}^{\mathbf{w},\bar{\mathbf{e}}}$ induces the same suppliers to produce as does $\mathbf{Q}^{FB}$, and again we are done. ■

C Appendix: Extensions

In Section C.1 we extend the model to allow heterogeneous outside options, and in Section C.2, we extend the model to allow buyers to have heterogeneous preferences over suppliers. Section C.3 provides a generalization of the one-to-many setup that encompasses additional models.

C.1 Heterogeneous outside options

The values of agents' outside options are central for determining the division of social surplus in complete information bargaining models. We now briefly discuss how our model can be augmented or reinterpreted to account for similar features. As we show, there are two types of outside options that can vary across agents: the opportunity cost of participating in the mechanism and the opportunity cost of producing (or buying), which we address in turn. Some of the comparative statics with respect to these costs are the same as with complete information bargaining, while other aspects are novel relative to complete information models.

Fixed costs of participating in the mechanism

For the purposes of this extension, we assume that $n^B = K^B = 1$ and drop the buyer subscripts. In this case, we can also assume, without further loss, that $k_j^S = 1$ for all $j \in \mathcal{N}^S$.

We first extend the model to allow the buyer and each supplier to have a positive outside option, denoted by $x_B \geq 0$ for the buyer and $x_j \geq 0$ for supplier j. These outside options are best thought of as fixed costs of participating in the mechanism because they have to be borne regardless of whether an agent trades. In this case, the incomplete information bargaining mechanism with weights $\mathbf{w}$ is the solution to

$$\max_{\langle \mathbf{Q},\mathbf{M}\rangle \in \mathcal{M}} \mathbb{E}_{v,\mathbf{c}}\left[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(v,c)\right] \text{ s.t. } \eta^B \pi^{\mathbf{w}} \geq x_B \text{ and for all } j \in \mathcal{N}^S,\ \eta_j^S \pi^{\mathbf{w}} \geq x_j.$$

Similar to the case in which the value of the outside options was zero for all agents, the allocation rule is as defined in Lemma 1, but now $\rho^{\mathbf{w}}$ is the smallest $\rho \geq \max \mathbf{w}$ such that

$$\mathbb{E}_{v,\mathbf{c}}\left[\sum_{j\in\mathcal{N}^S} (\Phi(v) - \Gamma_j(c_j)) \cdot \mathbf{1}_{\Phi^{w^B/\rho}(v) \geq \Gamma_j^{w_j/\rho}(c_j) = \min_{\ell\in\mathcal{N}^S} \Gamma_\ell^{w_\ell/\rho}(c_\ell)}\right] \geq x_B + \sum_{j\in\mathcal{N}^S} x_j, \tag{40}$$

if such a ρ exists (if no such ρ exists, then the constraints cannot be met).

Production-relevant outside options

Alternatively, one can think of outside options as affecting a supplier's cost of producing or as the buyer's best alternative to procuring the good. Typically, one would expect these to be more sizeable than the costs of participating in the mechanism. To allow for heterogeneity in these production-relevant outside options, we now relax the assumption that all suppliers' cost distributions have the identical support $[\underline{c}, \overline{c}]$ and assume instead that, with a commonly known outside option of value $y_j \geq 0$, the support of supplier j's cost distribution is $[\underline{c}_j, \overline{c}_j]$ with $\underline{c}_j = \underline{c} + y_j$ and $\overline{c}_j = \overline{c} + y_j$. If $G_j(c)$ is j's cost distribution without the outside option, then given outside option y_j, its cost distribution is $G_j^o(c) = G_j(c - y_j)$, with density $g_j^o(c) = g_j(c - y_j)$ and support $[\underline{c}_j, \overline{c}_j]$. In other words, increasing a supplier's outside option shifts its distribution to the right without changing its shape. Likewise, given outside option $y_B \geq 0$, the distribution of the buyer's value v is $F^o(v) = F(v + y_B)$ with density $f^o(v) = f(v + y_B)$ and support $[\underline{v} - y_B, \overline{v} - y_B]$.

Increasing the value of an agent's outside option has two effects. First, it worsens its distribution in the sense that for $y_j > 0$ and $y_B > 0$, we have $G_j^0(c) \leq G_j(c)$ for all c and $F^o(v) \geq F(v)$ for all v. Hence, under the first-best, an agent is less likely to trade the larger is the value of its outside option. While this effect differs from what one would usually obtain in complete information models, it is an immediate implication of the "worsening" of the agent's distribution.

The second effect is less immediate and partly, but not completely, offsets the first under the assumption that hazard rates are monotone, that is, assuming that $G_j(c)/g_j(c)$ is increasing in c and $(1 - F(v))/f(v)$ is decreasing in v. To see this, let us focus on supplier j. The arguments for the buyer (and of course all other suppliers) are analogous. We denote the weighted virtual cost of supplier j when it has outside option y_j by

$$\Gamma_{j,a}^o(c) \equiv c + (1 - a)\frac{G_j(c - y_j)}{g_j(c - y_j)} = \Gamma_{j,a}(c - y) + y < \Gamma_{j,a}(c), \tag{41}$$

where the inequality holds for all $a < 1$ because the monotone hazard rate assumption implies that $\Gamma_{j,a}'(c) > 1$ for all $a < 1$. This in turn has two, somewhat subtle implications. Let z be the threshold for supplier j to trade when its outside option is zero, i.e., keeping z fixed, supplier j trades if and only if $\Gamma_{j,a}(c) \leq z$. (Note that z will be the minimum of the buyer's weighted virtual value and the smallest weighted virtual cost of supplier j's competitors, but this does not matter for the argument that follows.) Assuming that $a < 1$ and $y_j < \overline{c} - \underline{c}$, which implies that $\underline{c}_j < \overline{c}$, it follows that there are costs $c \in [\underline{c}_j, \overline{c}]$ and thresholds z such that supplier j trades when it has the outside option and not without it, that is, $\Gamma_{j,a}^o(c) < z < \Gamma_{j,a}(c)$. This reflects the reasonably well-known result that optimal mechanisms tend to discriminate in favor of weaker agents (McAfee and McMillan, 1987), which in this case is the agent with the positive outside option. It also resonates with intuition from complete information models: keeping costs fixed, the agent with the better outside option is treated more favorably, indeed, it is evaluated according to a smaller weighted virtual cost. However, from an ex ante perspective, the larger is the value of the outside option, the less likely is the agent to trade. To see this, consider a fixed realization of z. (The distribution of these thresholds is not be affected by supplier j's outside option and hence our argument extends directly once one integrates over z and its density.) Given y_j, supplier j trades if and only if its cost c is below $\tau(y)$ satisfying $\Gamma^o_{j,a}(\tau(y)) = z$. Using (41), this is equivalent to $\Gamma_{j,a}(\tau(y) - y) + y = z$, which in turn is equivalent to $\tau(y) = \Gamma^{-1}_{j,a}(z - y) + y$, whose derivative for $a < 1$ satisfies

$$0 < \tau'(y) = -\frac{1}{\Gamma'_{j,a}(\Gamma^{-1}_{j,a}(z - y))} + 1 < 1,$$

where the inequalities follow because $\Gamma'_{j,a}(c) > 1$. This implies that, for a fixed z, the probability that supplier j trades decreases in y. To see this, notice that this probability is $G^o_j(\tau(y)) = G_j(\tau(y) - y)$, whose derivative with respect to y is $g_j(\tau(y) - y)(\tau'(y) - 1) < 0$. In words, although the threshold $\tau(y)$ increases in y, it does so with a slope that is less than 1, which implies that the probability that supplier j trades decreases in y. This effect is not present in complete information models, which in a sense take an ex post perspective by looking at outcomes realization by realization. While improving the outside option y_j improves supplier j's payoff after its value or cost has been realized, supplier j's ex ante expected payoff decreases in y_j. Moreover, because an increase in y_j worsens supplier j's distribution, the revenue constraint becomes (weakly) tighter, implying an increase in $\rho^{\mathbf{w}}$, which further reduces supplier j's expected payoff.

C.2 Preferences over suppliers and bargaining externalities

To allow for and investigate bargaining externalities, we restrict attention to the case of one buyer, $n^B = 1$, with demand for $K^B \geq 1$ units, and $n^S \geq 2$ suppliers, but we generalize the setup to allow the buyer to have heterogeneous preferences over the suppliers. To this end, we let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$ be a commonly known vector of taste parameters of the buyer, with the meaning that the value to the buyer of trade with supplier j when the buyer's type is v is $\theta_j v$. Thus, under (ex post) efficiency, trade should occur between the buyer and supplier j if and only if $\theta_j v - c_j$ is positive and among the K^B highest values of $(\theta_\ell v - c_\ell)_{\ell \in \mathcal{N}^S}$. The problem is trivial if $\max_{j \in \mathcal{N}^S} \theta_j \overline{v} \leq \underline{c}$ because then it is never ex post efficient to have trade with any supplier, so assume that $\max_{i\in\mathcal{N}} \theta_i \overline{v} > \underline{c}$.

This setup encompasses (i) differentiated products by letting the supplier-specific taste parameters differ; (ii) a one-buyer version of the Shapley and Shubik (1972) model by setting $K^B = 1$; and (iii) a version of the Shapley-Shubik model in which the buyer has demand for multiple products of the suppliers by setting $K^B > 1$. For a generalization of the one-to-many setup that encompasses additional models, see Section C.3.

We define the virtual surplus $\Lambda_j^{\mathbf{w},\boldsymbol{\theta}}$ associated with trade between the buyer and supplier j, accounting for the agents' bargaining weights $\mathbf{w}$ and the buyer's preferences $\boldsymbol{\theta}$, with $\rho^{\mathbf{w},\boldsymbol{\theta}}$ defined analogously to before as $\Lambda_j^{\mathbf{w},\boldsymbol{\theta}}(v, c_j) \equiv \theta_j \Phi^{w^B/\rho^{\mathbf{w},\boldsymbol{\theta}}}(v) - \Gamma_j^{w_j^S/\rho^{\mathbf{w},\boldsymbol{\theta}}}(c_j)$. Let $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv (\Lambda_j^{\mathbf{w},\boldsymbol{\theta}}(v, c_j))_{j\in\mathcal{N}^S}$ and denote by $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})_{(K^B)}$ the K^B-highest element of $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})$. As before, in order to save notation, we ignore ties.

Lemma C.2. *Assuming that $n^B = 1$ and $n^S \geq 2$, in the generalized setup with buyer preferences $\boldsymbol{\theta}$, incomplete information bargaining with weights $\mathbf{w}$ has the allocation rule for $j \in \mathcal{N}^S$, $Q_j^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv 1$ if $\Lambda_j^{\mathbf{w},\boldsymbol{\theta}}(v, c_j) \geq \max\{0, \boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})_{(K^B)}\}$, and otherwise $Q_j^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv 0$.*

Proof. The extension to allow supplier specific quality parameters follows by analogous arguments to Lemma 1 noting that the buyer's value for supplier j's good is $\theta_j v$, whose distribution is $\hat{F}(x) \equiv F(x/\theta_j)$ on $[\theta_j \underline{v}, \theta_j \overline{v}]$ with density $\hat{f}(x) = \frac{1}{\theta_j} f(v/\theta_j)$. Thus, the virtual type when the buyer's value is v is

$$\theta_j v - \frac{1 - \hat{F}(\theta_j v)}{\hat{f}(\theta_j v)} = \theta_j v - \theta_j \frac{1 - F(v)}{f(v)} = \theta_j \Phi(v).$$

Thus, the parameter θ_j "factors out" of the virtual type function. The extension to multi-object demand follows by standard mechanism design arguments. ■

We can now use this generalized setup to analyze bargaining externalities between suppliers. If $K^B < n$, then one effect of an increase in θ_i is that agents other than i are less likely to be among the at-most K^B agents that trade. In contrast, if $K^B \geq n$ and $\rho^{\mathbf{w},\boldsymbol{\theta}} > \max \mathbf{w}$, then the probability that supplier j trades, $\Pr(\theta_j \Phi^{w^B/\rho^{\mathbf{w},\boldsymbol{\theta}}}(v) \geq \Gamma_j^{w_j^S/\rho^{\mathbf{w},\boldsymbol{\theta}}}(c_j))$, does not depend on the preference parameters of the other suppliers except through their effect on $\rho^{\mathbf{w},\boldsymbol{\theta}}$. If $\rho^{\mathbf{w},\boldsymbol{\theta}} > \max \mathbf{w}$, then an increase in a rival supplier's preference parameter causes an increase in $\rho^{\mathbf{w},\boldsymbol{\theta}}$, which increases the probability of trade and so benefits the supplier. Thus, we have the following result:

Proposition 11. *Assuming that $n^B = 1$ and $n^S \geq 2$, in the generalized setup with bargaining weights $\mathbf{w}$ and buyer preferences $\boldsymbol{\theta}$, if $K^B \geq n$ and $\rho^{\mathbf{w},\boldsymbol{\theta}} > \max \mathbf{w}$, then an increase in the preference parameter for one supplier increases the payoffs for all suppliers.*

The result of Proposition 11 does not necessarily extend to the case with $K^B < n$, as shown in the following example.

Example with bargaining externalities

In Table 1, we consider the case of one buyer and two suppliers with symmetric bargaining weights. Assuming that F, G_1, and G_2 are the uniform distribution on $[0,1]$ and that $\theta_2 = 1$, we allow the buyer's preference for supplier 1, θ_1, and the buyer's total demand, K^B, to vary.

Table 1: Outcomes for one-to-many price formation for the case of one buyer and two suppliers with $\mathbf{w} = \mathbf{1}$, symmetric $\boldsymbol{\eta}$, types that are uniformly distributed on $[0,1]$, and $\theta_2 = 1$. The values of K^B and θ_1 vary as indicated in the table.

	$K^B = 1$		$K^B = 2$	
θ_1:	1	2	1	2
$1/\rho^{\mathbf{w},\boldsymbol{\theta}}$	0.73	0.76	0.67	0.72
u_B	0.13	0.34	0.14	0.38
u_1	0.05	0.21	0.07	0.22
u_2	0.05	0.01	0.07	0.08

As shown in Table 1, focusing on the case with $K^B = 1$, an increase in the buyer's preference for supplier 1 from $\theta_1 = 1$ to $\theta_1 = 2$ benefits supplier 1 (u_1 increases) but harms supplier 2 (u_2 decreases). The increase in the buyer's preference for supplier 1 means that supplier 2 is less likely to trade. As a result, supplier 2 is harmed by the increase in the buyer's preference for supplier 1. But when $K^B = 2$, the results differ. Supplier 1 again benefits from being preferred by the buyer, but in this case supplier 2 also benefits, albeit less than supplier 1. The increase in the buyer's value from trade with supplier 1 means that the value of $\rho^{\mathbf{w},\boldsymbol{\theta}}$ decreases, so supplier 2 trades more often. As a result of the change from $\theta_1 = 1$ to $\theta_1 = 2$, both u_1 and u_2 increase.

C.3 Generalization

Here we provide a further generalization of the setup with one buyer and multiple suppliers to allow a more general structure for the buyer's preferences over suppliers.

Let $\mathcal{P}$ be the set of subsets of $\mathcal{N}^S$ with no more than K^B elements (including the empty set) and let $\boldsymbol{\theta} = \{\theta_X\}_{X \in \mathcal{P}}$ be a commonly known vector of taste parameters of the buyer satisfying the "size-dependent discounts" condition of Delacrétaz et al. (2019). Specifically, let there be supplier-specific preferences $\{\hat{\theta}_j\}_{j \in \mathcal{N}^S}$ and size-dependent discounts $\{\delta_j\}_{j \in \mathcal{N}^S}$

with $0 = \delta_0 = \delta_1 \leq \delta_2 \leq \cdots \leq \delta_n$ such that for all $X \in \mathcal{P}$, $\theta_X = \sum_{i \in X} \hat{\theta}_i - \delta_{|X|}$. Thus, the buyer's value for purchasing from suppliers in $X \in \mathcal{P}$ when its type is v is $\theta_X v$, which depends on the buyer's value, the buyer's preferences for standalone purchases from the suppliers in X, and a discount that depends on the total number of units purchased. Note that $\theta_\varnothing = 0$, so that the value to the buyer of no trade is zero.

This setup encompasses (i) the homogeneous good model with constant marginal value or decreasing marginal value by setting $\hat{\theta}_j = \theta$ for some common θ and for $j \in \mathcal{N}^S$, δ_j either all zero for constant marginal value or increasing in j for decreasing marginal value; (ii) differentiated products by letting $\hat{\theta}_j$ differ across j and setting all δ_j to zero; (iii) a one-buyer version of the Shapley-Shubik model by setting $K^B = 1$; and (iv) a version of the Shapley-Shubik model in which the buyer has demand for multiple products of the suppliers by setting $K^B > 1$.

Define

$$X^*_\rho(v, \mathbf{c}) \in \arg\max_{X \in \mathcal{P}} \theta_X \Phi^{1/\rho}(v) - \sum_{i \in X} \Gamma_i^{1/\rho}(c_i),$$

i.e., $X^*_\rho(v, \mathbf{c})$ is the set of trading partners for the buyer that maximizes the difference between the weighted virtual value, scaled by $\theta_{X^*_\rho(v,\mathbf{c})}$, and the weighted virtual costs of the trading partners. We then define ρ^* to be the smallest $\rho \geq 1$ such that

$$\mathbb{E}_{v,\mathbf{c}} \left[\theta_{X^*_\rho(v,\mathbf{c})} \Phi(v) - \sum_{i \in X^*_\rho(v,\mathbf{c})} \Gamma_i(c_i) \right] = 0.$$

Given the type realization $(v, \mathbf{c})$, the one-to-many ρ^*-mechanism induces trade between the buyer and suppliers in $X^*_{\rho^*}(v, \mathbf{c})$. The expected payoff of the buyer is

$$\mathbb{E}_v \left[\hat{u}_B(\underline{v}) + \int_{\underline{v}}^{v} \sum_{X \in \mathcal{P}} \theta_X \Pr_{\mathbf{c}} \left(X \in X^*_{\rho^*}(x, \mathbf{c}) \right) dx \right],$$

and the expected payoff of supplier j is

$$\mathbb{E}_{c_j} \left[\hat{u}_j(\overline{c}) + \int_{c_j}^{\overline{c}} \Pr_{v,\mathbf{c}_{-j}} \left(i \in X^*_{\rho^*}(v, x, \mathbf{c}_{-j}) \right) dx \right].$$

D Appendix: Implementation

D.1 k-double auction as a special case

In the k-double auction of Chatterjee and Samuelson (1983), given $k \in [0,1]$, the buyer and supplier in a k-double auction simultaneously submit bids p_B and p_S, and trade occurs at the price $kp_B + (1-k)p_S$ if and only if $p_B \geq p_S$. By construction, the k-double auction never incurs a deficit. If the agents' types are uniformly distributed on $[0,1]$, then the linear Bayes Nash equilibrium of the k-double auction results in trade if and only if $v \geq c\frac{1+k}{2-k} + \frac{1-k}{2}$.[4] As first noted by Myerson and Satterthwaite (1983), for $k = 1/2$ and uniformly distributed types, the k-double auction yields the second-best outcome. Williams (1987) then generalized this insight by showing that, for uniformly distributed types and any $k \in [0,1]$, the k-double auction implements the outcomes of incomplete information bargaining for some bargaining weights. These outcomes are illustrated in Figure D.1.

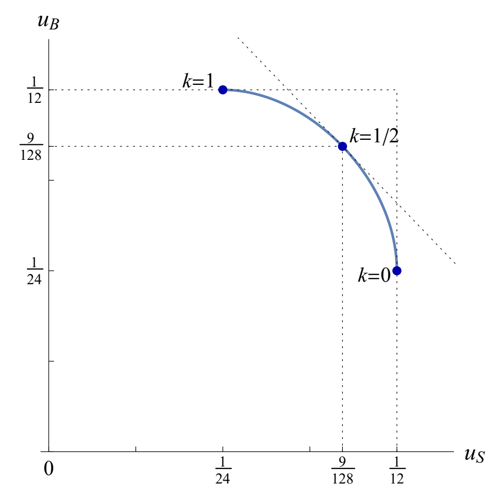

Figure D.1: Payoffs in the k-double auction for all $k \in [0,1]$. Assumes that there is one single-unit supplier and one single-unit buyer that the supplier's cost and the buyer's value are uniformly distributed on $[0,1]$.

To see that incomplete information bargaining encompasses the k-double auction as a special case, note that for the case of one single-unit supplier, one single-unit buyer, and uniformly distributed types, for all $\mathbf{w}$, $\rho^{\mathbf{w}}$ is such that

$$0 = \mathbb{E}_{v,c}\left[(\Phi(v) - \Gamma(c)) \cdot Q^{\mathbf{w}}(v,c)\right] = \int_{\frac{1-w^B/\rho^{\mathbf{w}}}{2-w^B/\rho^{\mathbf{w}}}}^{1} \int_0^{\frac{v-(1-w^B/\rho^{\mathbf{w}})(1-v)}{2-w^S/\rho^{\mathbf{w}}}} (2v - 1 - 2c)\, dc dv,$$

[4] In the linear Bayes Nash equilibrium, a buyer of type v bids $p_B(v) = (1-k)k/(2(1+k)) + v/(1+k)$ and a supplier with cost c bids $p_S(c) = (1-k)/2 + c/(2-k)$. For $k = 1$, $p_B(v) = v/2$ and $p_S(c) = c$, and for $k = 0$, $p_B(v) = v$ and $p_S(c) = (c+1)/2$. Thus, for $k \in \{0,1\}$, the k-double auction reduces to take-it-or-leave-it offers.

where the second equality uses the expression for $Q^{\mathbf{w}}(v,c)$ from Lemma 1 (we write $Q^{\mathbf{w}}$ instead of $Q_1^{\mathbf{w},S}$ because for the case that we consider here, there is only one relevant quantity, and to reduce notation, we drop the agent indices on w^B and w^S). Solving this for $\rho^{\mathbf{w}}$, we get

$$\rho^{\mathbf{w}} = \frac{1}{2}\left(w^B + w^S + \sqrt{w^{B2} - w^B w^S + w^{S2}}\right).$$

Making the substitutions $w^S = 1-\Delta$ and $w^B = \Delta$ and writing $\rho^{\mathbf{w}}$ as a function of Δ, we have

$$\rho^{\Delta} = \frac{1}{2}\left(1 + \sqrt{1 - 3\Delta + 3\Delta^2}\right). \tag{42}$$

It is then straightforward to derive, for a given Δ, the conditions on (v,c) such that there is trade. Equating this condition with the condition for trade in the k-double auction allows one to identify the relation between Δ and k as

$$\Delta_k \equiv \frac{(2-k)k}{1+2k-2k^2}, \tag{43}$$

where Δ_k is increasing in k and varies from 0 to 1 as k varies from 0 to 1.

To see that the price-formation mechanism with bargaining differential Δ_k is equivalent to the k-double auction, substitute the expression for ρ^{Δ} in place of $\rho^{\mathbf{w}}$ into the expression derived from Lemma 1 for $Q^{(1-\Delta,\Delta)}(v,c)$ to get

$$Q^{(1-\Delta,\Delta)}(v,c) \equiv \begin{cases} 1 & \text{if } v \geq \frac{1-2\Delta(1-c)+(1+2c)\sqrt{1-3\Delta+3\Delta^2}}{2(1-\Delta)+2\sqrt{1-3\Delta+3\Delta^2}}, \\ 0 & \text{otherwise.} \end{cases}$$

Using (43), it then follows that

$$Q^{(1-\Delta_k,\Delta_k)}(v,c) = \begin{cases} 1 & \text{if } v \geq c\frac{1+k}{2-k} + \frac{1-k}{2}, \\ 0 & \text{otherwise,} \end{cases}$$

which is the same allocation rule as for the k-double auction.

To conclude, note that we can replicate the incomplete information bargaining outcome with bargaining weights $\mathbf{w}$ by using bargaining weights $(1-\Delta,\Delta)$ with $\Delta = \frac{w^B}{w^B+w^S}$. Thus, for any bargaining weights $\mathbf{w}$, there exists $k \in [0,1]$, namely k such that $\Delta_k = \frac{w^B}{w^B+w^S}$, such that the outcome of the k-double auction is the same as the outcome of incomplete information bargaining with weights $\mathbf{w}$. Conversely, for any $k \in [0,1]$, there exist bargaining weights $\mathbf{w}$, namely $\mathbf{w} = (1-\Delta_k,\Delta_k)$, such that incomplete information bargaining with weights $\mathbf{w}$ yields the same outcome as the k-double auction.

D.2 Axiomatic approach

In this appendix, we provide axiomatic foundations for incomplete information bargaining. Just as the Nash bargaining solution (and cooperative game theory more generally) abstracts away from specific bargaining protocols, our mechanism design based approach does the same. Nash bargaining maps primitives to a bargaining solution that specifies agents' payoffs, and our approach maps primitives (type distributions of the agents) to agents' expected payoffs via the unique (or essentially unique) mechanism that satisfies the axioms presented here.

We take a setup with incomplete information involving independent private types as given and impose axioms on the mechanism that defines incomplete information bargaining. This differs from the existing literature, which imposes axioms on outcomes. In light of the stringent discipline that the incomplete information paradigm imposes, this point of departure is necessary. As Ausubel et al. (2002) note, asking for efficient outcomes in bargaining is "fruitless," given the impossibility theorem of Myerson and Satterthwaite (1983).

As we now show, axioms of incentive compatibility, individual rationality, and no deficit identify a set of feasible mechanisms. Additional axioms of constrained efficiency and symmetry pin down a unique mechanism. Generalizing the efficiency and symmetry axioms allows differential weights on agents' welfare, analogous to generalized Nash bargaining.

Observe that the payoff equivalence theorem is *distribution free* (or detail free) insofar as it holds for any distributions $F_1, \dots, F_{n^B}$ and $G_1, \dots, G_{n^S}$ that have compact supports and positive densities on $(\underline{v}, \overline{v})$ and $(\underline{c}, \overline{c})$, respectively. In formulating our axioms, we are therefore guided by the principle that the axioms should make no reference to distributional assumptions and should make no presumptions beyond these foundational assumptions on the setup. That said, in the body of the paper we assumed regularity (i.e., that virtual value and cost functions are increasing) in order to avoid the technicalities of ironing. We do the same here, although all results continue to hold without regularity assumptions when the weighted and unweighted virtual value and cost functions are replaced by their ironed counterparts.

The first three axioms ensure that the incomplete information bargaining mechanism is *feasible*, which means that beyond satisfying resource constraints, the mechanism satisfies incentive compatibility, individual rationality, and does not run a deficit.[5]

[5]To simplify the exposition, we only require that the mechanism does not run a deficit in expectation, allowing for the possibility that ex post the mechanism may run a budget deficit for some realizations. At some overhead cost, ideas along the lines of Arrow (1979) and d'Aspremont and Gérard-Varet (1979) or, alternatively, Crémer and Riordan (1985) can be used to avoid deficits for all type realizations.

Axiom 1: Incentive compatibility: The mechanism is incentive compatible.

Axiom 2: Individual rationality: The mechanism is individually rational.

Axiom 3: No deficit: The mechanism does not run a deficit.

Axioms 1–3 are, obviously, consistent with incomplete information bargaining with any weights $\mathbf{w}$. Axioms 1–3 constrain incomplete information bargaining, but they also hold, in a sense, in the Nash bargaining framework (Nash, 1950). In that complete information setup, incentive compatibility is trivially satisfied because the "mechanism" already knows the agents' types, and participation in Nash bargaining is individually rational because the bargaining outcome gives each agent a payoff of at least its disagreement payoff. In addition, there is no scope for running a deficit. Thus, there is a sense in which Axioms 1–3 are implied by the other aspects and axioms in the Nash bargaining setup.

Our fourth and fifth axioms ensure that social surplus is maximized, conditional on the constraints imposed by the other axioms, and that when that maximizer is not unique, the solution is one that treats the buyers and suppliers symmetrically.

Axiom 4: Efficiency: The mechanism maximizes expected social surplus subject to the conditions of Axioms 1–3.

Axiom 5: Symmetry: Whenever positive surplus is available to be distributed to agents while still respecting Axioms 1–4, it is distributed equally among the agents.

Axioms 4 and 5 identify a unique mechanism within the class of direct mechanisms that maximize expected social surplus subject to incentive compatibility, individually rationality, and no deficit, namely incomplete information bargaining with symmetric bargaining weights $\mathbf{w}$ and symmetric $\boldsymbol{\eta}$.

Axioms 4 and 5 have clear counterparts in the "efficiency" and "symmetry" axioms that underlie the Nash bargaining solution. The efficiency axiom in Nash bargaining requires efficiency for any realization of types, whereas Axiom 4 requires efficiency subject to feasibility constraints. Axiom 5 requires that the outcome treat the buyer and suppliers symmetrically whenever that can be done within the context of the other axioms, which is similar to Nash's requirement of symmetry.

If for symmetric $\mathbf{w}$, we have $\pi^{\mathbf{w}} = 0$, as is the case when $\rho^{\mathbf{w}} > \max \mathbf{w}$, then Axioms 1–4 imply that $\hat{u}_1^B(\underline{v}) = \cdots = \hat{u}_{n^B}^B(\underline{v}) = \hat{u}_1^S(\overline{c}) = \cdots = \hat{u}_{n^S}^S(\overline{c}) = 0$, and so the symmetry axiom has no additional bite beyond the other axioms. But if $\pi^{\mathbf{w}} > 0$, then the symmetry axiom requires that this surplus be allocated symmetrically among the agents, resulting in expected interim payoffs to the worst-off types that are positive and equal.

In this case when $n^B = n^S = 1$ and $\underline{v} > \overline{c}$, all five axioms are satisfied using the posted-price mechanism with $p = (\underline{v} + \overline{c})/2$. Notice the similarity to Nash bargaining here—the posted price is the same price at which a buyer with value $\underline{v}$ and a supplier with cost $\overline{c}$ would trade under Nash's axioms and assumptions.

Finally, Nash bargaining specifies, in addition to efficiency and symmetry, axioms of invariance to affine transformations of the utility functions and independence to irrelevant alternatives. In incomplete information bargaining, the assumption of risk neutrality (and the associated quasilinear preferences) means that invariance to affine transformations of the utility functions is maintained. And a restriction that certain allocations or transfer payments are not permitted does not affect the outcome of incomplete information bargaining as long as the optimal allocation and transfers remain available. Thus, the incomplete information bargaining mechanism satisfies the additional axioms of Nash.

We now state our characterization result.

Theorem 1. *The incomplete information bargaining mechanism with symmetric* $\mathbf{w}$ *and* $\boldsymbol{\eta}$*, is the unique direct mechanism satisfying Axioms* 1–5.

Proof of Theorem 1. When $\mathbf{w}$ is symmetric, then by definition, the incomplete information bargaining mechanism maximizes welfare subject to incentive compatibility, individual rationality, and no deficit. Further, because the allocation pins down the agents' interim expected payoffs up to a constant, the mechanism is unique up to the payoffs of the worst-off types, $\hat{u}_1^B(\underline{v}), \ldots, \hat{u}_{n^B}^B(\underline{v})$ and $\hat{u}_1^S(\overline{c}), \ldots, \hat{u}_{n^S}^S(\overline{c})$, but these are uniquely pinned down by the assumption of symmetric $\boldsymbol{\eta}$. ■

We extend our efficiency and symmetry axioms to allow for different bargaining weights for the buyer and suppliers, with at least one of the weights being positive, as follows:

Axiom $4'(\mathbf{w})$**: Generalized efficiency with weights** $\mathbf{w}$: The mechanism maximizes expected weighted welfare, $\mathbb{E}_{\mathbf{v},\mathbf{c}}[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(\mathbf{v},\mathbf{c})]$, subject to the conditions of Axioms 1–3.

Axiom $5'(\mathbf{w})$**: Generalized symmetry with weights** $\mathbf{w}$: Whenever positive surplus is available to be distributed to agents while still respecting Axioms 1–3 and $4'(\mathbf{w})$, it is distributed among the agent(s) with the maximum bargaining weight.

This leads us to the result that incomplete information bargaining is essentially uniquely defined by the axioms and criteria described above, where the "essentially" relates to the possibility of different tie-breaking rules when more than one agent has the maximum bargaining weight. The proof is similar to that of Theorem 1, but with adjustments for the buyers' and suppliers' bargaining weights, and so is omitted.

Theorem 2. *The incomplete information bargaining mechanism with weights* $\mathbf{w}$ *is the essentially unique direct mechanism satisfying Axioms* 1–3, $4'(\mathbf{w})$, *and* $5'(\mathbf{w})$.

D.3 Extensive-form approach

Building on the model of Loertscher and Niedermayer (2019), we define the *fee-setting extensive-form game* to have one buyer with single-unit demand, $n^S = n \geq 1$ suppliers, and an intermediary that facilitates the buyer's procurement of inputs from the suppliers and that charges the buyer a fee for its service. Let $\mathcal{N} = \mathcal{N}^S$ denote the set of suppliers. The buyer's value and the suppliers' costs are not known by the intermediary, although the intermediary does know the distributions F and $G_1, \ldots, G_n$ from which those types are independently drawn. The timing is as follows: 1. the intermediary announces (and commits to) a *discriminatory clock auction*, which we define below, and fee schedule $(\sigma_1, \ldots, \sigma_n)$, where σ_i maps the price p paid by the buyer to supplier i to the fee $\sigma_i(p)$ paid by the buyer to the intermediary, should the buyer purchase from supplier i; 2. the buyer sets a reserve r for the auction; 3. the intermediary holds the auction with reserve r, which determines the winning supplier, if any, and the payment to that supplier; 4. given winner i and payment p, supplier i provides the good to the buyer, and the buyer pays p to supplier i and $\sigma_i(p)$ to the intermediary. If no supplier bids below the reserve, then there is no trade and no payments are made, including no payment to the intermediary.

As just mentioned, the intermediary uses a discriminatory clock auction with reserve r. Because this is a procurement, it is a *descending* clock auction, with the clock price starting at the reserve r and descending from there. As in any standard clock auction, participants choose when to exit, and when they exit, they become inactive and remain so. The clock stops when only one active bidder remains, with ties broken by randomization. A *discriminatory* clock auction specifies supplier-specific discounts off the final clock price $(\delta_1, \ldots, \delta_n)$, where δ_i maps the clock price to supplier i's discount—activity by supplier i at a clock price of $\hat{p}$ obligates supplier i to supply the product at the price $\hat{p} - \delta_i(\hat{p})$. By the usual clock auction logic, in the essentially unique equilibrium in non-weakly-dominated strategies, supplier i with cost c_i remains active in the auction until the clock price reaches $\hat{p}$ such that $\hat{p} - \delta_i(\hat{p}) = c_i$, and then supplier i exits. We assume that the suppliers follow these strategies.

Turning to the incentives of the buyer and intermediary, the buyer chooses the reserve to maximize its expected payoff, and the intermediary chooses the auction discounts and the fee structure to maximize the expected value of its objective. To allow for the possibility that the intermediary has an interest in promoting the surplus of the agents, we assume that the intermediary's objective is to maximize expected weighted welfare subject to no deficit, with surplus distributed according to shares $\boldsymbol{\eta}$, where we refer to $\mathbf{w}$ in this context as intermediary preference weights and $\boldsymbol{\eta}$ as profit shares.

As we show in the following proposition, the outcome of incomplete information bargaining arises as a Bayes Nash equilibrium of this game:

Proposition 12. *The outcome of incomplete information bargaining with bargaining weights* $\mathbf{w}$ *and shares* $\boldsymbol{\eta}$ *is a Bayes Nash equilibrium outcome of the fee-setting extensive-form game with intermediary preference weights* $\mathbf{w}$ *and profit shares* $\boldsymbol{\eta}$*.*

Proof. Consider the Bayes Nash equilibrium of the fee-setting game with intermediary preference weights $\mathbf{w}$. To begin, we assume that $\pi^{\mathbf{w}} \equiv \mathbb{E}_{v,\mathbf{c}}[\sum_{i\in\mathcal{N}}(\Phi(v) - \Gamma_i(c_i)) \cdot Q_i^{\mathbf{w}}(v, \mathbf{c})] = 0$, and then we address the required adjustments for the case with $\pi^{\mathbf{w}} > 0$ at the end.

Suppose that the intermediary sets auction discounts relative to the clock price $\hat{p}$ of $\delta_i(\hat{p}) \equiv \hat{p} - \Gamma_i^{w_i/\rho^{\mathbf{w}}\,-1}(\hat{p})$ and a fee schedule given by, for all $i \in \mathcal{N}$,

$$\sigma_i(p) \equiv \Phi^{w^B/\rho^{\mathbf{w}}\,-1}(\Gamma_i^{w_i/\rho^{\mathbf{w}}}(\Gamma_i^{-1}(p))) - p,$$

and suppose that the buyer sets a reserve of $\Phi^{w^B/\rho^{\mathbf{w}}}(v)$. Then, given our assumption that each supplier i follows its weakly dominant strategy of remaining active until a clock price $\hat{p}$ such that $\hat{p} - \delta_i(\hat{p}) = c_i$, supplier i remains active until a price of $\Gamma_i^{w_i/\rho^{\mathbf{w}}}(c_i)$, and so supplier i wins if and only if

$$\Gamma_i^{w_i/\rho^{\mathbf{w}}}(c_i) = \min_{j\in\mathcal{N}} \Gamma_j^{w_j/\rho^{\mathbf{w}}}(c_j) \leq \Phi^{w^B/\rho^{\mathbf{w}}}(v),$$

which, by Lemma 1, corresponds to the intermediary's optimal allocation rule, $\mathbf{Q}^{\mathbf{w}}$. In equilibrium, if supplier i wins the auction, then the auction ends with a clock price of

$$\hat{p} \equiv \min_{j\in\mathcal{N}\setminus\{i\}} \{\Phi^{w^B/\rho^{\mathbf{w}}}(v), \Gamma_j^{w_j/\rho^{\mathbf{w}}}(c_j)\},$$

and the buyer makes a payment $p = \hat{p} - \delta_i(\hat{p})$ to supplier i and a payment of $\sigma_i(p)$ to the intermediary.

To summarize, given the suppliers' optimal bidding strategies and a reserve set by the buyer of $\Phi^{w^B/\rho^{\mathbf{w}}}(v)$, the intermediary's choice of auction format and fee schedule are optimal because they result in the allocation rule that maximizes the weighted objective subject to no deficit and because the allocation rule pins down the payoffs up to nonnegative constants that are zero under our assumption that $\pi^{\mathbf{w}} = 0$. It remains to show that the best response to the intermediary's auction format and fee schedule for a buyer with value v is to choose a reserve of $\Phi^{w^B/\rho^{\mathbf{w}}}(v)$.

To reduce notation, let $x_B \equiv w^B/\rho^{\mathbf{w}}$ and $x_i \equiv w_i/\rho^{\mathbf{w}}$. Define the distribution of supplier i's weighted virtual type $\Gamma_i^{x_i}(c_i)$ by $\tilde{G}_i^{x_i}(z) = G_i(\Gamma_i^{x_i-1}(z))$, and, letting $\mathbf{x} \equiv (x_1, \ldots, x_n)$, define the distribution of the minimum of the weighted virtual types of suppliers other than i by

$$\tilde{G}_{-i}^{\mathbf{x}}(z) = 1 - \prod_{j \in \mathcal{N}\setminus\{i\}} (1 - \tilde{G}_j^{x_j}(z)).$$

The expected payment by the buyer to the suppliers given the reserve r can be written as

$$\begin{aligned}
&\sum_{i\in\mathcal{N}} \mathbb{E}\left[\Gamma_i(c_i) \cdot \mathbf{1}_{\Gamma_i^{x_i}(c_i) \leq \min_{j\neq i}\left\{r, \Gamma_j^{x_j}(c_j)\right\}}\right] \\
= &\sum_{i\in\mathcal{N}} \int_{\underline{c}}^{\max\{\underline{c}, \Gamma_i^{x_i-1}(r)\}} \int_{\Gamma_i^{x_i}(c_i)}^{\infty} \Gamma_i(c_i) d\tilde{G}_{-i}^{\mathbf{x}}(z) dG_i(c_i) \\
= &\sum_{i\in\mathcal{N}} \int_{\underline{c}}^{\max\{\underline{c}, \Gamma_i^{x_i-1}(r)\}} \Gamma_i(c_i)(1 - \tilde{G}_{-i}^{\mathbf{x}}(\Gamma_i^{x_i}(c_i))) dG_i(c_i) \\
= &\sum_{i\in\mathcal{N}} \int_{\underline{c}}^{\max\{\underline{c}, \Gamma_i(\Gamma_i^{x_i-1}(r))\}} y \frac{\left[1 - \tilde{G}_{-i}^{\mathbf{x}}(\Gamma_i^{x_i}(\Gamma_i^{-1}(y)))\right] g_i(\Gamma_i^{-1}(y))}{\Gamma_i'(\Gamma_i^{-1}(y))} dy,
\end{aligned}$$

where the final equality uses the change of variables $y = \Gamma_i(c_i)$. Thus, the buyer with value v maximizes its interim expected payoff by choosing r to solve

$$\max_r \sum_{i\in\mathcal{N}} \left(\int_{\underline{c}}^{\max\{\underline{c}, \Gamma_i(\Gamma_i^{x_i-1}(r))\}} (v - y - \sigma_i(y)) \frac{\left[1 - \tilde{G}_{-i}^{\mathbf{x}}(\Gamma_i^{x_i}(\Gamma_i^{-1}(y)))\right] g_i(\Gamma_i^{-1}(y))}{\Gamma_i'(\Gamma_i^{-1}(y))} dy \right),$$

which, when $\underline{c} < \Gamma_i(\Gamma_i^{x_i-1}(r))$, has first-order condition

$$\begin{aligned}
0 &= \sum_{i\in\mathcal{N}} \Gamma_i'(\Gamma_i^{x_i-1}(r))\Gamma_i^{x_i-1\prime}(r) \left((v - \Gamma_i(\Gamma_i^{x_i-1}(r)) - \sigma_i(\Gamma_i(\Gamma_i^{x_i-1}(r))))\right) \frac{(1 - \tilde{G}_{-i}^{\mathbf{x}}(r))g_i(\Gamma_i^{x_i-1}(r))}{\Gamma_i'(\Gamma_i^{x_i-1}(r))} \\
&= \sum_{i\in\mathcal{N}} \Gamma_i'(\Gamma_i^{x_i-1}(r))\Gamma_i^{x_i-1\prime}(r) \left(v - \Phi^{e_B-1}(r)\right) \frac{(1 - \tilde{G}_{-i}^{\mathbf{x}}(r))g_i(\Gamma_i^{x_i-1}(r))}{\Gamma_i'(\Gamma_i^{x_i-1}(r))},
\end{aligned}$$

where the second equality uses the definition of the fee schedule $\boldsymbol{\sigma}$. Given our assumptions, the second-order condition is satisfied when the first-order condition is, and so the buyer's problem is solved by $r = \Phi^{e_B}(v) = \Phi^{w^B/\rho^{\mathbf{w}}}(v)$, giving the buyer nonnegative interim expected payoff, which completes the proof for the case with $\pi^{\mathbf{w}} = 0$. If $\pi^{\mathbf{w}} > 0$, then this "excess profit" must be distributed via fixed payments between the agents and the intermediary so that the worst-off type of each agent $i \in \{B\} \cup \mathcal{N}$ has interim expected payoff $\eta_i \pi^{\mathbf{w}}$. ■

Thus, the fee-setting extensive-form game, in which a fee-setting intermediary procures an input for the buyer from competing suppliers, provides a microfoundation for the price-formation mechanism. Reminiscent of Crémer and Riordan (1985), the sequential nature of the game allows an equilibrium that is Bayesian incentive compatible for one agent, the buyer, and dominant-strategy incentive compatible for the other agents, the suppliers. The equilibrium of the fee-setting game satisfies ex post individual rationality for both the buyer and suppliers, but only balances the intermediary's budget in expectation. In contrast, in Crémer and Riordan (1985), the budget is balanced ex post, but individual rationality is no longer satisfied ex post for all agents.[6]

The fee-setting extensive-form game is, for example, a reasonable description of the wholesale used car market analyzed by Larsen (2021). There, an intermediary runs auctions, facilitates further bargaining in the substantial number of cases in which the auction does not result in trade, and collects fees from traders.

[6]In the model of Crémer and Riordan (1985), individual rationality is satisfied ex post for the agent that moves first (the buyer in our case) and only ex ante for the agents that move second (suppliers in our case).

E Details

In this appendix, we provide the details underlying the comparative statics results for vertical integration (Section E.1) and for investment (Section E.2) and for estimation with bargaining breakdown (Section 7.2).

E.1 Details for vertical integration comparative statics

Here we provide the details underlying Figure 3. For the purposes of the comparative statics illustrated there, we assume that there is one buyer with single-unit demand, i.e., $n^B = 1$ and $K^B = 1$, $[\underline{c}, \overline{c}] = [0, 1]$, and $G_1 = \cdots = G_{n^S} \equiv G$. We denote the buyer's distribution by F, dropping the buyer subscript since there is only one buyer. We assume that all agents, including a vertically integrated firm, have bargaining weight equal to one. Because we focus on social surplus effects, the tie-breaking shares are not relevant.

Denote by $L_n(c) = 1 - (1 - G(c))^n$ the distribution of the lowest of n independent draws from G and by $l_n(c)$ the associated density. We assume that $\overline{v} \geq 1$. Specifically, for Figure 3(a), we assume that $\underline{v} = 0$ and $\overline{v} = 1$, and for Figure 3(b), we assume that $\overline{v} = 1.2$ and that $\underline{v}$ varies as indicated in the figure. Below, we let n denote the number of nonintegrated suppliers, which means that if there is vertical integration, the total number of suppliers is $n + 1$. Given n and $\rho \geq 1$, we denote by $R_\rho(n)$ the revenue of the mechanism absent vertical integration. We have

$$R_\rho(n) \equiv \int_{\Phi^{1/\rho^{-1}}(0)}^{\overline{v}} \int_0^{\Gamma^{1/\rho^{-1}}(\Phi^{1/\rho}(v))} (\Phi(v) - \Gamma(c)) f(v) l_n(c) dc dv.$$

Because $L_n(c) \leq L_{n+1}(c)$, if $\Gamma(c)$ is increasing, which we assume, this implies that

$$R_\rho(n) < R_\rho(n+1).$$

Because the second-best mechanism given n is characterized by the unique $\rho_n^* \geq 1$ such that

$$R_{\rho_n^*}(n) = 0,$$

and because $R_\rho(n)$ increases in ρ, it follows that

$$\rho_{n+1}^* < \rho_n^*.$$

This means that the more competition there is, the more efficient is the second-best mechanism, that is, the closer is the allocation rule under the second-best mechanism to the one under first-best. Social surplus given n without integration is

$$W(n) \equiv \int_{\Phi^{1/\rho_n^*}{}^{-1}(0)}^{\overline{v}} \int_0^{\Gamma^{1/\rho_n^*}{}^{-1}(\Phi^{1/\rho_n^*}(v))} (v-c)f(v)l_n(c)dcdv.$$

This suggests that eliminating the agency (or double information rent) problem by vertically integrating the buyer with one supplier has less of an impact the larger is n. But there is more to vertical integration than eliminating an independent supplier, which all else equal makes the outside market less efficient because vertical integration also changes the virtual valuation function of the buyer from $\Phi(v) = v - \frac{1-F(v)}{f(v)}$ to

$$\hat{\Phi}(x) \equiv x - \frac{1-H(x)}{h(x)},$$

where $H(x) = 1-(1-F(x))(1-G(x))$ and $h(x) = H'(x)$, with support $[\underline{c}, \overline{c}] = [0,1]$. We assume that F and G are such that $\hat{\Phi}$ is increasing. Denote the corresponding weighted virtual value function by $\hat{\Phi}^a(x) \equiv x-(1-a)\frac{1-H(x)}{h(x)}$. It follows that if $F = G$, as is the case in Figure 3(a), then we have

$$\hat{\Phi}(x) = x - \frac{1}{2}\frac{1-F(x)}{f(x)} = \Phi^{1/2}(x) \geq \Phi(x),$$

where the inequality is strict for all $x < 1$.

Given n independent suppliers and $\rho \geq 1$, revenue from the mechanism under vertical integration, denoted $\hat{R}_\rho(n)$, is

$$\hat{R}_\rho(n) = \int_{\hat{\Phi}^{1/\rho-1}(0)}^{1} \int_0^{\Gamma^{1/\rho^{-1}}(\hat{\Phi}^{1/\rho}(x))} (\hat{\Phi}(x) - \Gamma(c))h(x)l_n(c)dcdx.$$

The second-best mechanism is characterized by $\hat{\rho}_n^*$ such that

$$\hat{R}_{\hat{\rho}_n^*}(n) = 0.$$

One might be inclined to think that the outside market becomes less efficient with vertical integration in the sense that $\hat{\rho}_n^* > \hat{\rho}_{n+1}^*$, where it will be recalled that if there are n independent suppliers with vertical integration there were $n+1$ independent suppliers without it. But it is neither clear whether $\hat{\rho}_n^* > \hat{\rho}_{n+1}^*$ is the case nor what it precisely means: even if ρ were the same, the allocation rules with and without vertical integration change because the virtual valuation changes. While it seems natural to think that $\hat{R}_\rho(n) < R_\rho(n+1)$ because of the decrease in the number of independent suppliers because of vertical integration, there is an additional revenue effect via the change in the virtual valuation function. As noted, for $F = G$, we have $\hat{\Phi}(x) > \Phi(x)$ for all $x < 1$, which has a revenue increasing effect. In addition, the density of x is different from the density of v. For example, if $F = G$, then we have $h(x) = 2f(x)(1 - F(x))$. All of this goes to show that it will, in general, be tricky to say much about which of the various effects dominates.

Social surplus under vertical integration given n independent suppliers, denoted $\hat{W}(n)$, consists of the social surplus from internal production by the vertically integrated firm, which is $\mathbb{E}_{v,c}[\max\{v - c, 0\}]$, plus the value from procuring from the outside market, that is

$$\hat{W}(n) \equiv \int_{\underline{v}}^{\overline{v}} \int_0^v (v - c) f(v) g(c) dc dv + \int_{\hat{\Phi}^{\hat{\rho}_n^*}{}^{-1}(0)}^{1} \int_0^{\Gamma^{\hat{\rho}_n^*}{}^{-1}(\hat{\Phi}^{1/\hat{\rho}_n^*}(x))} (x - c) h(x) l_n(c) dc dx.$$

To see how this is derived, note that if $v < c$, then there is no internal production and the buyer's willingness to pay for an independent supplier is $x = v$; if $v > c$, then the integrated firm's willingness to pay is $x = c$, meaning that if it procures from the external market it does so to replace production by its own supply unit.

Under the assumption that F and G are uniform, we can compute ρ_n^* and $\hat{\rho}_n^*$ and hence $W(n)$ and $\hat{W}(n)$. Figure 3(a) plots the change in social surplus from vertical integration, $\hat{W}(n) - W(n+1)$, for $n \in \{1, \ldots, 10\}$, and Figure 3(b) graphs $\hat{W}(n) - W(n+1)$ as a function of $\underline{v}$ for the cases of $n = 1$ and $n = 2$. In line with the intuition provided above, the social benefits are positive when n is small and negative when n becomes larger. For the uniform example, once $\hat{W}(n) - W(n + 1)$ is negative for some value of n, say n', it remains negative for all $n > n'$, while asymptotically approaching 0 (from below). The nonmonotone behavior of $\hat{W}(n) - W(n+1)$ in n is as expected because when n is large, vertical integration has little effect on internal sourcing (which is unlikely to occur) and small effects on the efficiency of the outside market because, in that case, the market is close to first-best before and after vertical integration. Computations also show that for $n = 0$,

$$\hat{W}(n) - W(n + 1) = 0.0260417 = \frac{1}{6} - \frac{9}{64},$$

where 1/6 is first-best welfare in the Myerson-Satterthwaite problem for uniformly distributed types and 9/64 is second-best welfare when F and G are uniform.

E.2 Details for investment comparative statics

Here we provide details underlying Section 6.3, which examines how equilibrium investments are affected by bargaining power and by the extent to which the supports of the value and cost distributions overlap.

As described in the body of the paper, 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. Specifically, 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 agent's investment e has cost $e^2/2$.

In this setup, the first-best investment e^{FB} is the same for the buyer and supplier and satisfies

$$(e^{FB}, e^{FB}) \in \arg\max_{e_S, e_B} \int_{\underline{v}}^{\underline{v}+1} \int_0^1 (v - c) \cdot \mathbf{1}_{c \leq v} \cdot dG_{e_S}(c) dF_{e_B}(v) - e_B^2/2 - e_S^2/2.$$

For example, if $X = 1.25$ and $\underline{v} = 1$, then the first-best investment is $e_S^{FB} = e_B^{FB} = 0.25$, implying that under the first-best investments, types are uniformly distributed for both the supplier and the buyer.

Second-best investment is also the same for the buyer and suppler and satisfies:[7]

$$(e^{SB}, e^{SB}) \in \arg\max_{e_S, e_B} \int_{\underline{v}}^{\underline{v}+1} \int_0^1 (v - c) \cdot \mathbf{1}_{\Gamma^{1/\rho^{SB}}(c; e_S) \leq \Phi^{1/\rho^{SB}}(v; e_B)} \cdot dG_{e_S}(c) dF_{e_B}(v) - e_B^2/2 - e_S^2/2,$$

where ρ^{SB} is the smallest $\rho \geq 1$ such that $\pi^{(1,1)}(e_B, e_S; \rho) \geq 0$, where

$$\pi^{\mathbf{w}}(e_B, e_S; \rho) \equiv \int_{\underline{v}}^{\underline{v}+1} \int_0^1 (\Phi(v; e_B) - \Gamma(c; e_S)) \cdot \mathbf{1}_{\Gamma^{w_S/\rho}(c; e_S) \leq \Phi^{w_B/\rho}(v; e_B)} \cdot dG_{e_S}(c) dF_{e_B}(v).$$

Now consider the Nash equilibrium investments. Our assumption that investments are not observed implies that given Nash equilibrium investments (e_S^{NE}, e_B^{NE}), trade occurs if and only if $\Gamma^{w_S/\rho^{NE}}(c; e_S^{NE}) \leq \Phi^{w_B/\rho^{NE}}(v; e_B^{NE})$. Further, fixed payments are determined by the agents' shares (η_S, η_B) and the Nash equilibrium budget surplus $\pi^{NE} \equiv \pi(e_B, e_S; \rho^{NE})$.

[7] The linear virtual type functions are given by

$$\Phi^{\beta}(v; e_B) \equiv v \frac{1 - \beta + X - e_B}{X - e_B} - \frac{(1 + \underline{v})(1 - \beta)}{X - e_B} \quad \text{and} \quad \Gamma^{\beta}(c; e_S) \equiv c \frac{1 - \beta + X - e_S}{X - e_S}.$$

The buyer's Nash equilibrium investment solves

$$e_B^{NE} \in \quad \arg\max_x \quad \int_{\underline{v}}^{\underline{v}+1} \int_0^1 (v - \Phi(v;x)) \cdot \mathbf{1}_{\Gamma^{w_S/\rho^{NE}}(c;e_S^{NE}) \leq \Phi^{w_B/\rho^{NE}}(v;e_B^{NE})} \cdot dG_{e_S^{NE}}(c) dF_e(v) - e^2/2 + \eta_B \pi^{NE},$$

the first-order condition of which is

$$e_B^{NE} = - \int_{\underline{v}}^{\underline{v}+1} \int_0^1 \left. \frac{\partial F_e(v)}{\partial e} \right|_{e=e_B^{NE}} \cdot \mathbf{1}_{\Gamma^{w_S/\rho^{NE}}(c;e_S^{NE}) \leq \Phi^{w_B/\rho^{NE}}(v;e_B^{NE})} \cdot g_{e_S^{NE}}(c) dc dv. \tag{44}$$

Analogously, the supplier's Nash equilibrium investment solves

$$e_S^{NE} \in \quad \arg\max_x \quad \int_{\underline{v}}^{\underline{v}+1} \int_0^1 (\Gamma(c;x) - c) \cdot \mathbf{1}_{\Gamma^{w_S/\rho^{NE}}(c;e_S^{NE}) \leq \Phi^{w_B/\rho^{NE}}(v;e_B^{NE})} \cdot dG_e(c) dF_{e_B^{NE}}(v) - e^2/2 + \eta_S \pi^{NE},$$

whose first-order condition is

$$e_S^{NE} = \int_{\underline{v}}^{\underline{v}+1} \int_0^1 \left. \frac{\partial G_e(c)}{\partial e} \right|_{e=e_S^{NE}} \cdot \mathbf{1}_{\Gamma^{w_S/\rho^{NE}}(c;e_S^{NE}) \leq \Phi^{w_B/\rho^{NE}}(v;e_B^{NE})} \cdot f_{e_B^{NE}}(v) dc dv. \tag{45}$$

Solving for $(e_S^{NE}, e_B^{NE}) \in [0, X]$ and $\rho^{NE} \geq \max\{w_S, w_B\}$ that satisfy (44), (45),

$$\pi^{\mathbf{w}}(e_B^{NE}, e_S^{NE}; \rho^{NE}) \geq 0, \quad \text{and} \quad (\rho^{NE} - \max\{w_S, w_B\}) \pi^{\mathbf{w}}(e_B^{NE}, e_S^{NE}; \rho^{NE}) = 0,$$

we obtain the Nash equilibrium investments and Lagrange multiplier on the no-deficit constraint.

We illustrate the effects of bargaining power and the distributional supports on equilibrium investment in Figure E.2, which expands upon Figure 4 in the body of the paper.

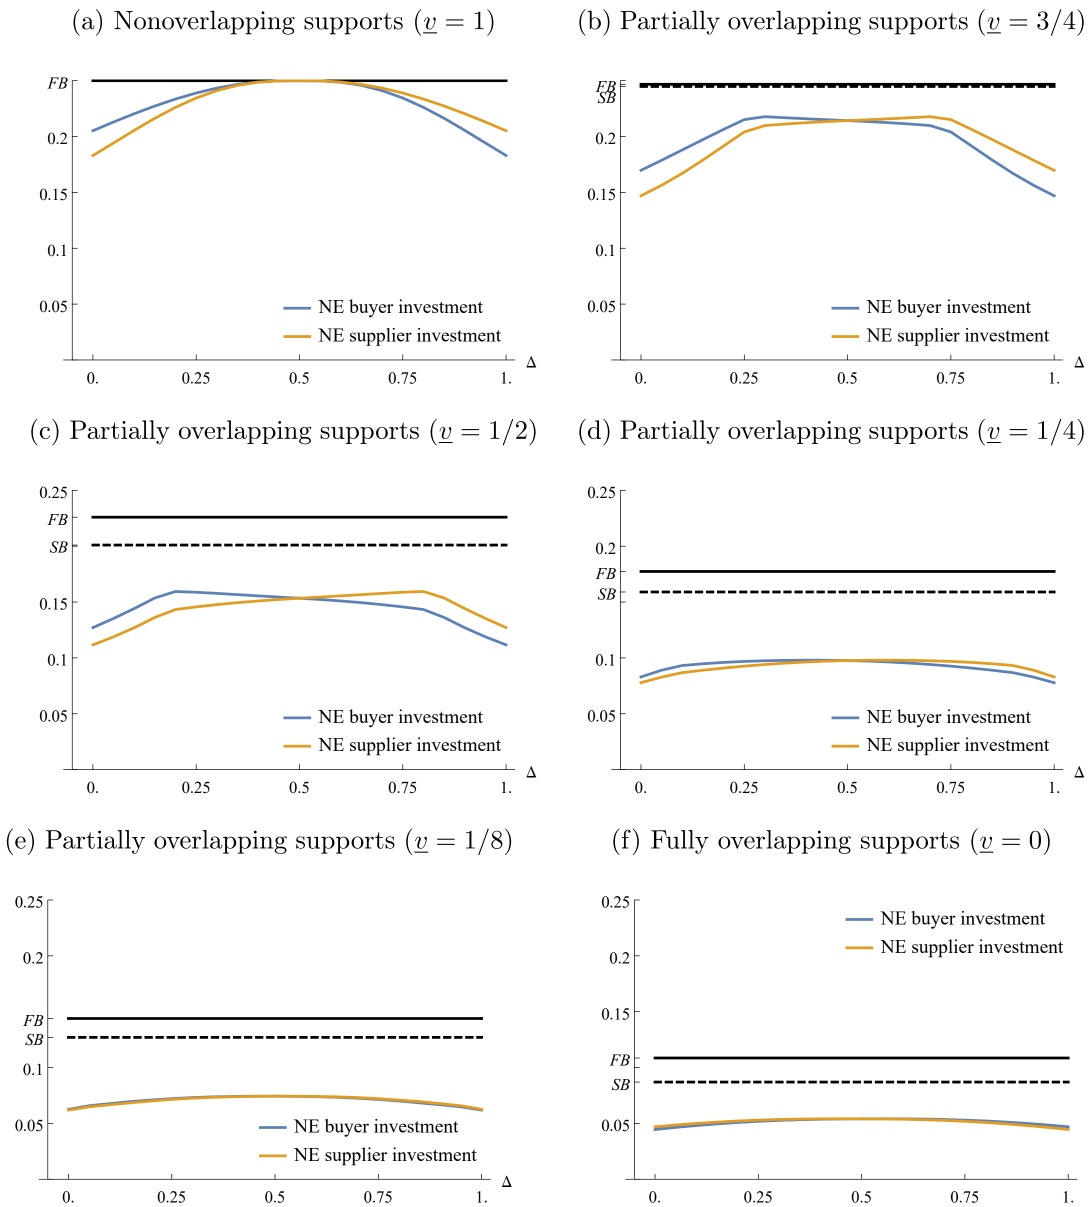

Figure E.2: Nash equilibrium investments with bargaining weights $(w^S, w^B) = (1 - \Delta, \Delta)$ for buyer distributions with varying supports. Assumes the linear virtual type setup for bilateral trade with $F(v) = 1-(1+\underline{v}-v)^{1.25-e_B}$, where $e_B \in [0, 1.25)$ is the buyer's investment, 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. For $\underline{v} \in \{1/4, 1/8, 0\}$, $\rho^{NE} > \max\{w_S, w_B\}$ for all bargaining weights. For $\underline{v} \in \{1/2, 3/4\}$, $\rho^{NE} = \max\{w_S, w_B\}$ for sufficiently asymmetric bargaining weights and $\rho^{NE} > \max\{w_S, w_B\}$ otherwise.

E.3 Details for bargaining breakdown and estimation

To illustrate how, with incomplete information bargaining, one can use observed frequencies of negotiation breakdowns to back out the parameters of the distributions from which the agents draw their types, consider a market with one buyer and two suppliers with single-unit demand and supply and types drawn from parameterized distributions

$$F(v) = 1 - (1 - v)^{1/\kappa} \quad \text{and} \quad G_j(c) = c^{1/\kappa_j}, \tag{46}$$

with support $[0, 1]$, where the parameters κ and κ_j are positive real numbers and have the interpretation of "capacities" insofar as larger values of κ and κ_j imply better distributions in the sense of first-order stochastic shifts. These distributions are analytically convenient because they imply linear virtual type functions. Rather than treating negotiation breakdowns as measurement error, which is difficult to justify if breakdown occurs more than fifty percent of the time in 25 million observations, the frequency of those breakdowns is valuable information that can be used for estimation in the incomplete information framework. Figure E.3(a) provides an example of how the probability of bargaining breakdown can be used to calibrate the model with parameterized distributions.

Then we use the model to predict the change in social surplus as result of vertical integration, which is displayed in Figure E.3(b). As shown, when the rate of bargaining breakdown in the pre-integration market is sufficiently low, i.e., the pre-integration market is sufficiently efficient, the change in social surplus from vertical integration is negative. In contrast, when the probability of breakdown is sufficiently high prior to integration, the increased efficiency associated with internal transactions dominates, and vertical integration increases social surplus.

(a) Calibration of distributions to data

mkt shares	Pr(breakdown)	$(\kappa_1,\ \kappa_2,\ \kappa)$
50-50	10%	(1, 1, 11)
50-50	30%	(1, 1, 3)
50-50	55%	(1, 1, 1)

(b) Change in expected social surplus following vertical integration

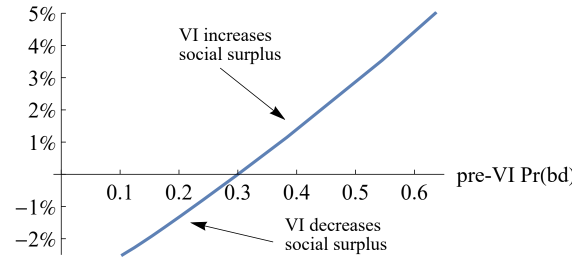

Figure E.3: Interaction between the pre-integration breakdown probability and the effect of vertical integration on social surplus. Panel (a): Calibration of distributional parameters based on market shares and breakdown probabilities assuming that $n^B = 1$ and $n^S = 2$ and that $\mathbf{w}$ and $\boldsymbol{\eta}$ are symmetric, $F(v)$ and $G_j(c)$ are given by (46), and $(\kappa_1 + \kappa_2)/2 = 1$. Panel (b): Change in expected social surplus due to vertical integration as the probability of breakdown in the pre-integration market, "pre-VI Pr(bd)," varies, based on the calibration of Panel (a).

References for the online appendix

Arrow, K. J. (1979): "The Properties Rights Doctrine and Demand Revelation under Incomplete Information," in *Economics and Human Welfare*, ed. by M. Boskin, New York, NY: Academic Press, 23–39.

Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002): "Bargaining with Incomplete Information," in *Handbook of Game Theory*, ed. by R. Aumann and S. Hart, Elsevier Science B.V., vol. 3, 1897–1945.

Boyd, S. and L. Vandenberghe (2004): *Convex Optimization*, Cambridge University Press.

Chatterjee, K. and W. Samuelson (1983): "Bargaining under Incomplete Information," *Operations Research*, 31, 835–851.

Crémer, J. and M. H. Riordan (1985): "A Sequential Solution to the Public Goods Problem," *Econometrica*, 53, 77–84.

Delacrétaz, D., S. Loertscher, L. Marx, and T. Wilkening (2019): "Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade," *Journal of Economic Theory*, 416–454.

d'Aspremont, C. and L.-A. Gérard-Varet (1979): "Incentives and Incomplete Information," *Journal of Public Economics*, 11, 25–45.

Krishna, V. (2010): *Auction Theory*, Elsevier, second ed.

Larsen, B. J. (2021): "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," *Review of Economic Studies*, 88, 851–882.

Loertscher, S. and A. Niedermayer (2019): "An Optimal Pricing Theory of Transaction Fees in Thin Markets," Working Paper, University of Melbourne.

McAfee, R. P. and J. McMillan (1987): "Auctions and Bidding." *Journal of Economic Literature*, 25, 699–738.

Myerson, R. and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

Nash, J. F. (1950): "The Bargaining Problem," *Econometrica*, 18, 155–162.

Shapley, L. and M. Shubik (1972): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.

Williams, S. R. (1987): "Efficient Performance in Two Agent Bargaining," *Journal of Economic Theory*, 41, 154–172.

——— (1999): "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14, 155–180.