Countervailing Power, Integration, and Investment under Incomplete Information*

Simon Loertscher† Leslie M. Marx‡

July 10, 2020

Abstract

We provide an incomplete information model to analyze bargaining, countervailing power, integration, and investments. These have been at the forefront of merger decisions. Countervailing power arguments stipulate that horizontal mergers can be socially desirable because they offset bargaining power on the other side of the market. The tradeoff between profit and social surplus with independent private values gives rise to the possibility of social-surplus-increasing countervailing power and socially harmful vertical integration. Moreover, efficient bargaining implies efficient, noncontractible investments; horizontal mergers can harm rivals; bargaining breaks down on the equilibrium path; and countervailing power can align investments with the first-best.

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

*We thank Matt Backus, Allan Collard-Wexler, Soheil Ghili, Brad Larsen, Joao Montez, Rob Porter, Patrick Rey, Glen Weyl, Steve Williams, and audiences at the 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 the University of Melbourne for valuable discussions and comments. Edwin Chan 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 No. 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

Countervailing power features prominently in antitrust debates. The idea that increasing market power on one side of the market countervails existing market power on the other side is appealing to many and at times embraced as if it had "talismanic power" (Steptoe, 1993). Nonetheless, the concept of countervailing power has been controversial since its inception, with John Galbraith viewing it as a mitigant of economic power of "substantial, and perhaps central, importance" and George Stigler lamenting the lack of any explanation for "why bilateral oligopoly should in general eliminate, and not merely redistribute, monopoly gains."[1] A considerable part of the controversy arises because formalizing countervailing power has proved challenging. As George Stigler's comment makes clear, it requires a model in which bargaining power affects not only the division but also the size of social surplus. This is challenging because, as noted by the *New Palgrave Dictionary*, "it is difficult to model bilateral monopoly or oligopoly, and there exists no single canonical model."[2] Recent merger cases further brought to the forefront the effects of vertical integration, and market structure more broadly, on investments, and on how these effects depend on bargaining and bargaining power.

In this paper, we provide and analyze an incomplete information model that sheds new light on the above. At the heart of our incomplete information model is that whether or not bargaining is efficient is endogenous to the market structure and bargaining weights. In particular, countervailing power arises naturally as bargaining weights change, sometimes to the point that the first-best becomes possible. Vertical integration can be socially harmful because it can render inefficient otherwise efficient bargaining. While in industrial organization there is both something of a consensus and a concern that accounting for incomplete information is important but challenging, this paper shows that, assuming independent private values, incomplete information can be modelled in a tractable way. Indeed, the independent private values model has all of the required properties to generate countervailing power effects, bargaining breakdown on the equilibrium path, and a more nuanced view of vertical integration. It also offers a disciplined approach to the surplus-profit tradeoff that is at the core of industrial organization, regulation, and of much of economics more generally.

We model bargaining as an incentive-compatible, individually-rational mechanism that maximizes the weighted sum of expected buyer and supplier surplus, subject to a no-deficit

[1]Galbraith (1954, p. 1) and Stigler (1954, p. 13), respectively. See also Galbraith (1952).

[2]Snyder (2008, p. 1188). As Carlton and Israel (2011, p. 128) put it: "For changes in bargaining outcomes due to a buyer merger to create a true efficiency, it must be that, post-merger, the parties are better able to arrive at an optimal nonlinear price schedule, perhaps due to lower transactions costs, which moves output closer to the competitive level."

constraint, with the weights in the objective reflecting the relative bargaining powers of the buyer and the suppliers. In general, the bargaining weights affect both the division and the size of expected social surplus as illustrated in Figure 1. The buyer's (seller's) surplus is maximized when the buyer (seller) has all the bargaining power and makes a take-it-or-leave-it offer, resulting in the two extremal points of the bargaining frontier. With randomized take-it-or-leave-it offers, one can achieve any convex combination between these two points. But by a logic that is essentially the same as the argument invoked by Paul Samuelson to show that with constant returns to scale the production possibility frontier is concave, away from the extremes one can, in general, do better by reoptimzing.[3] Thus, an equalization of bargaining weights increases social surplus, which is the key to countervailing power.

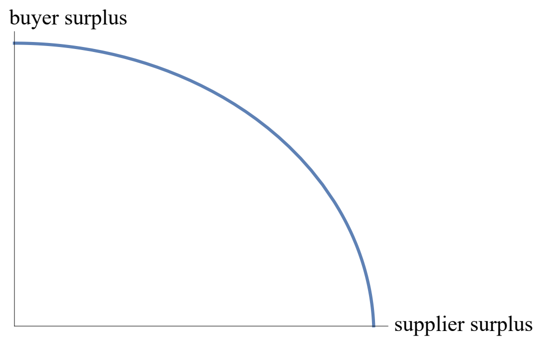

Figure 1: Illustration of a bargaining frontier.

We show that a horizontal merger that "levels the playing field" by equalizing bargaining weights can improve social surplus. Indeed, it can make the first-best possible when, prior to the merger and the change in bargaining weights, it was not achievable because the price-formation process was too strongly tilted towards the buyer. In contrast, a horizontal merger between suppliers that does not affect the bargaining weights never improves social surplus and always makes achieving the first-best more difficult because the merger eliminates a bid and thereby exacerbates the deficit problem of incomplete information bargaining.

It is useful to distinguish between an agent's *productive power* (or strength) and its *bargaining power*. Productively stronger buyers have, or tend to have, higher values, and productively stronger suppliers tend to have lower costs. In contrast, an agent's bargaining power captures its ability to affect bargaining in its favor. While, empirically, productive power and bargaining power may be correlated, conceptually, they are distinct and independent. As a case in point, both business and leisure air travellers are price-takers and hence have the same amount of bargaining power: none. However, business passengers are

[3]See Samuelson (1949). Appendix C provides the derivation of Figure 1 for the bilateral trade problem with uniform distributions.

productively stronger, which is why they are charged higher prices. A merger between two suppliers creates a productively stronger supplier that draws its cost from the minimum of the two pre-merger distributions, without per se increasing the bargaining power of the new firm. Changes in bargaining power are necessary, without being sufficient, for there to be countervailing power because, keeping bargaining weights fixed, the bid elimination associated with a merger always makes the deficit problem for incomplete information bargaining more severe. Countervailing power therefore requires there to be unequal bargaining weights before the merger and more equal bargaining weights post merger. Consequently, a merger defense based on countervailing power arguments needs to demonstrate that the buyer has greater bargaining power pre merger—evidence for which will depend on the price-formation process in the situation of interest[4]—and provide a theory for why this power diminishes without vanishing through the merger.

We show that vertical integration between the buyer and a supplier can create a bilateral trade problem à la Myerson and Satterthwaite (1983) in which the first-best becomes impossible when it was possible prior to integration. This occurs, for example, when vertical integration leaves only one independent supplier in the market and when the buyer's lowest possible value before integration exceeds the suppliers' highest possible cost. In situations like these, vertical integration is thus socially harmful. It is so in ways and for reasons that are absent when the efficiency of the price-formation process is exogenously fixed. Of course, vertical integration also eliminates a bilateral trade problem, namely that within the newly created entity. Therefore, the social surplus effects of vertical integration can go either way. We show that under appropriate assumptions, the likely effects of vertical integration can be estimated using pre-integration data, including the frequency of bargaining breakdown. Although our analysis does not imply that vertical integration is universally bad, it does show that a presumption that vertical integration improves social surplus is not warranted.[5]

Our approach to incomplete information bargaining can also be embedded in a dynamic setting in which agents first make investments and bargain once their values and costs are realized. This extension is motivated in part by the upsurge of interest in investment incentives following the Dow-DuPont merger decision and in part by the prominence of both bargaining weights and non-contractible investments in the theory of the firm à la Grossman-

[4]For example, when prices are formed through procurement auctions, evidence of buyer power includes discriminatory reserve prices, discounts and handicaps in the auction, and random winner selection.

[5]The U.S. Department of Justice and Federal Trade Commission's "Vertical Merger Guidelines" (June 30, 2020) emphasize positive effects of vertical integration: "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" (p. 11). That said, they do also recognize that "While the agencies more often encounter problematic horizontal mergers than problematic vertical mergers, vertical mergers are not invariably innocuous" (p. 2).

Hart-Moore. As does this literature, we assume that investments are not contractible, which implies that the mechanism does not vary with investments off the equilibrium path, and we assume that investments improve agents' type distributions. In this setup, the equilibrium investments are efficient if—and under additional conditions, only if—bargaining is efficient.[6] The contrast to the results obtained in the theory of the firm could hardly be sharper. There, complete information and efficient price formation (for example, Nash bargaining or Shapley value) are imposed by assumption and induce hold up and thereby inefficient investments. In incomplete information models, the privacy of information protects agents against hold up, and efficient bargaining implies efficient investments.

Our model is also amenable to introducing variations 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.[7] 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. 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, meaning 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 by the fact that, 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 harms the supplier 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 harm the agent.

When the model is extended to allow for multi-object demand and supplier-specific preferences by the buyer, bargaining externalities also naturally arise. We provide conditions under which an increase in the buyer's preference for one supplier's product increases the payoffs for all of the suppliers. This occurs when an increase in a buyer's preference for a

[6] As we discuss, a similar result holds for investment by the suppliers in quality.

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

supplier improves the efficiency of incomplete information bargaining to the benefit of all suppliers.

The remainder of the paper is structured as follows. Section 2 introduces the setup. In Section 3, we derive the price-formation process that incorporates bargaining weights and provide a model of incomplete information bargaining. Section 4 derives results pertaining to horizontal and vertical mergers. Section 5 analyzes investment incentives. In Section 6, we extend the model to allow variation in agents' outside options and opportunity costs and to allow multi-object demand with supplier-specific preferences. In addition, we provide an extensive-form implementation. In Section 7, we discuss related literature. Section 8 concludes the paper. Formal mechanism design results and longer proofs are relegated to appendices.

2 Setup

We consider a procurement setup with n suppliers indexed by $i \in \mathcal{N} \equiv \{1, \dots, n\}$, each with the capacity to produce one unit of a good, and one buyer, labeled B, with demand for one unit of the good. The buyer draws its value v independently from a distribution $F(v)$ with support $[\underline{v}, \overline{v}]$ and density $f(v)$ that is positive for all $v \in (\underline{v}, \overline{v})$. Supplier i draws its cost c_i independently from distribution $G_i(c_i)$ with support $[\underline{c}, \overline{c}]$ and density $g_i(c_i)$ that is positive for all $c_i \in (\underline{c}, \overline{c})$. We assume that F and $G_1, \dots, G_n$ are common knowledge, while the realized value v and the realized costs $c_1, \dots, c_n$ are the private information of the buyer and individual suppliers, respectively. To save on notation, we ignore ties among the agents' types.

The buyer and the suppliers have quasilinear preferences. The expected payoff of supplier i with cost c_i when producing the good with probability q_i and receiving the expected monetary transfer m is $\hat{u}_i(c_i; m, q_i) = m - c_i q_i$. The expected payoff of the buyer with value v when receiving the object with probability q and making the expected monetary payment m is $\hat{u}_B(v; m, q) = vq - m$. Under ex post efficiency (and ignoring ties), trade occurs between the buyer and supplier i if and only if $v - c_i > \max_{j \neq i}\{0, v - c_j\}$. The problem is trivial if $\overline{v} \leq \underline{c}$ because then it is never ex post efficient to have trade with any supplier. Therefore, we assume that $\overline{v} > \underline{c}$.

Because we allow both the buyer's value and the suppliers' costs to be 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.[8] Indeed, for $n = 1$, our setup encompasses the classic Myerson-Satterthwaite

[8] To avoid informed-principal problems, we model the mechanism-design problem as one in which a third

(1983) setting, where, as they show, ex post efficient trade is impossible if and only if $\underline{v} < \overline{c}$. 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*. With one supplier and nonoverlapping supports, we obtain ex post efficient trade even under incomplete information.

We denote the buyer's virtual value function by $\Phi(v) \equiv v - \frac{1-F(v)}{f(v)}$ and supplier i's virtual cost function by $\Gamma_i(c) \equiv c + \frac{G_i(c)}{g_i(c)}$.[9] We assume that the virtual value and virtual cost functions are increasing.[10] For $a \in [0, 1]$, we define the a-weighted virtual value function by $\Phi_a(v) \equiv v - (1-a)\frac{1-F(v)}{f(v)}$ and the a-weighted virtual cost function for supplier i by $\Gamma_{i,a}(c) \equiv c + (1-a)\frac{G_i(c)}{g_i(c)}$.[11] Observe that monotonicity of $\Phi(v)$ and $\Gamma_i(c)$ implies that $\Phi_a(v)$ and $\Gamma_{i,a}(c)$ are also monotone. As observed by Mussa and Rosen (1978), $\Phi(v)$ can be interpreted as a marginal revenue function. Analogously, $\Gamma_i(c)$ has the interpretation of being supplier i's marginal cost function.

Although we focus on a procurement setting with one buyer and one or more suppliers, all of our results extend with the appropriate adjustments to a sales auction with one supplier and one or more buyers.

3 Incomplete information bargaining

At the heart of essentially any economic model of exchange 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. As mentioned in the introduction, we stay within this tradition by working with a given price-formation process, but we add to it by introducing a price-formation process with incomplete information that allows for heterogeneous bargaining weights and that 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.[12]

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 the buyer, we still obtain the buyer-optimal mechanism, just as one would if we assumed that the buyer organizes the exchange.

[9]If $f(\overline{v}) = 0$, then define $\Phi(\overline{v}) \equiv \lim_{v \uparrow \overline{v}} \Phi(v)$, and if $g_i(\underline{c}) = 0$, then define $\Gamma_i(\underline{c}) \equiv \lim_{c \downarrow \underline{c}} \Gamma_i(c)$. If $f(\underline{v}) = 0$, then $\Phi(\underline{v}) = -\infty$, and if $g_i(\overline{c}) = 0$, then $\Gamma_i(\overline{c}) = \infty$.

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

[11]This departs from 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.

[12]For evaluating the merits of a countervailing power argument, 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

For the 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.

3.1 Price-formation mechanism

We model incomplete information bargaining as a direct mechanism $\langle \mathbf{Q}, \mathbf{M} \rangle$ operated by the *market*, where the allocation rule $\mathbf{Q} : [\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]^n \to \mathbb{R}_+^{n+1}$ maps the buyer's and suppliers' types to their quantities (the quantity received by the buyer and quantities provided by the suppliers), and the payment rule $\mathbf{M} : [\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]^n \to \mathbb{R}^{n+1}$ maps types to payments (the payment from the buyer and the payments to the suppliers).[13] Feasibility requires that for all type realizations, $Q_i(v, \mathbf{c}) \leq 1$ for all $i \in \mathcal{N}$ and $Q_B(v, \mathbf{c}) = \sum_{i \in \mathcal{N}} Q_i(v, \mathbf{c})$, where the buyer is endowed with free disposal. Of course, excess production will never be optimal, so the sum of the quantities provided by the suppliers will not be greater than the buyer's demand.

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, the value of the outside option—to zero.[14] A direct mechanism has *no deficit* if the expected payment from the buyer is greater than or equal to the sum of the expected payments to the suppliers. For formal definitions, see Appendix A.1.

Fix a mechanism $\langle \mathbf{Q}, \mathbf{M} \rangle$ and type realizations $(v, \mathbf{c})$. Then for an allocation rule $\mathbf{Q}$ with $Q_i(v, \mathbf{c}) \in \{0, 1\}$ for all $i \in \{B\} \cup \mathcal{N}$, the buyer's and supplier i's *ex post* surpluses as a function of the type realizations are

$$U_{B;\mathbf{Q},\mathbf{M}}(v, \mathbf{c}) \equiv \hat{u}_B(v; M_B(v, \mathbf{c}), Q_B(v, \mathbf{c})) = vQ_B(v, \mathbf{c}) - M_B(v, \mathbf{c}),$$

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 put the question of transaction costs on center stage in economics, this limits the value of the approach.

[13]Any model of trade involves a mechanism that 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.

[14]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.1 in terms of Bayesian incentive compatibility and interim individual rationality, one could also use the ex post versions of those constraints.

and

$$U_{i;\mathbf{Q},\mathbf{M}}(v,\mathbf{c}) \equiv \hat{u}_i(c_i; M_i(v,\mathbf{c}), Q_i(v,\mathbf{c})) = M_i(v,\mathbf{c}) - c_i Q_i(v,\mathbf{c}),$$

where $\hat{u}_B$ and $\hat{u}_i$ are the functions introduced in Section 2. The budget surplus generated by the mechanism is

$$R_{\mathbf{M}}(v,\mathbf{c}) \equiv M_B(v,\mathbf{c}) - \sum_{i\in\mathcal{N}} M_i(v,\mathbf{c}),$$

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

$$W_{\mathbf{Q}}(v,\mathbf{c}) \equiv \sum_{i\in\mathcal{N}} (v - c_i) Q_i(v,\mathbf{c}).$$

To capture bargaining power, we endow the agents with bargaining weights $\mathbf{w} = (w_i)_{i\in\{B\}\cup\mathcal{N}}$, where $w_i \in [0,1]$ is agent 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

$$W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(v,\mathbf{c}) \equiv w_B U_{B;\mathbf{Q},\mathbf{M}}(v,\mathbf{c}) + \sum_{i\in\mathcal{N}} w_i U_{i;\mathbf{Q},\mathbf{M}}(v,\mathbf{c}), \tag{1}$$

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

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

We contrast the outcome under the market mechanism with what a *planner* does, whose mechanism maximizes $\mathbb{E}_{v,\mathbf{c}}[W_{\mathbf{Q}}(v,\mathbf{c})]$ subject to these same constraints. We denote by $\mathbf{Q}^{\mathbf{w}}$ the allocation rule that solves the market's problem and by $\mathbf{Q}^*$ the allocation rule that solves the planner's problem. We say $\mathbf{Q}^*$ is *first-best* if the constraint (2) is not binding, in which case we denote it $\mathbf{Q}^{FB}$. Otherwise, $\mathbf{Q}^*$ is said to be *second-best*. Accordingly, the first-best and second-best quantities or outcomes are then the quantities or outcomes that arise under the mechanisms with allocation rules $\mathbf{Q}^{FB}$ and $\mathbf{Q}^*$, respectively.

Letting $\mathcal{M}$ be the set of incentive-compatible, individually-rational, no-deficit mechanisms, we define a *incomplete information bargaining* mechanism with bargaining weights $\mathbf{w}$ to be a mechanism in $\mathcal{M}$ that maximizes expected weighted welfare, $\mathbb{E}_{v,\mathbf{c}}[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(v,\mathbf{c})]$. Notice that, because we evaluate outcomes according to expected welfare $\mathbb{E}_{v,\mathbf{c}}[W_{\mathbf{Q}}(v,\mathbf{c})]$, the bargaining weights $\mathbf{w}$ are indeed only bargaining weights, that is, they do not affect how outcomes are evaluated, although they affect the distribution of social surplus and, as we will see, sometimes the size of social surplus.

The Lagrangian associated with maximizing expected weighted welfare (1) subject to the no-deficit constraint (2) can be written as $\mathbb{E}_{v,\mathbf{c}}\left[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(v,\mathbf{c})+\rho R_{\mathbf{M}}(v,\mathbf{c})\right]$, where ρ is the Lagrange multiplier on the no-deficit constraint. Using the mechanism design techniques described in Appendix A.1, the incentive-compatibility constraint implies that the Lagrangian can be rewritten as an expression that has one term involving fixed payments to the worst-off types and another term given by

$$\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}[w_B(\overbrace{v-\Phi(v)}^{\text{buyer surplus}})+w_i(\overbrace{\Gamma_i(c_i)-c_i}^{\text{supplier } i \text{ surplus}})+\rho(\overbrace{\Phi(v)-\Gamma_i(c_i)}^{\text{budget surplus}})]Q_i(v,\mathbf{c})\right]. \tag{3}$$

Because any budget surplus can be reallocated to the agents through fixed payments, the shadow price of budget surplus, ρ, satisfies $\rho \geq \max \mathbf{w}$. In addition, because a positive expected budget surplus is always possible given our assumption that $\overline{v} > \underline{c}$, the shadow price is finite. Denoting by $\mathbf{Q}_\rho$ the pointwise maximizer of (3) for a given ρ, the optimum is characterized by the smallest ρ greater than or equal to $\max \mathbf{w}$ such that the no-deficit constraint is satisfied at $\mathbf{Q}_\rho$.[15] In what follows, it will be convenient to work with the inverse of the multiplier, $\beta \equiv 1/\rho$, which is an element of $(0, \frac{1}{\max \mathbf{w}}]$. We denote the largest $\beta \in (0, \frac{1}{\max \mathbf{w}}]$ such that the no-deficit condition is satisfied by $\beta^{\mathbf{w}}$, which we refer to as the budget parameter.[16]

An immediate implication is that when $\mathbf{w} = \mathbf{1}$, indeed whenever bargaining weights are symmetric, incomplete information bargaining delivers the second-best quantities. As an illustration, consider a bilateral trade problem, i.e., assume $n = 1$. As mentioned, efficient trade is impossible if and only if the supports overlap, i.e., $\underline{v} < \overline{c}$. Because bargaining weights of $\mathbf{w} = (1, 1)$ yield the second-best outcome, this implies that $\beta^{(1,1)} < 1$ holds if and only if the supports overlap. With nonoverlapping supports, i.e., $\underline{v} \geq \overline{c}$, the incentive-compatibility and individual-rationality constraints can be satisfied by charging the buyer $\underline{v}$ and paying the supplier $\overline{c}$, generating a surplus of $\underline{v} - \overline{c} \geq 0$, which can then be shared between the buyer and the supplier. More generally, for $\mathbf{w} = (1, \mathbf{0})$, we obtain the buyer's optimal mechanism. Because the buyer is essentially the residual claimant and so prefers not to trade rather than accepting a deficit, this mechanism does not run a deficit, implying that $\beta^{(1,\mathbf{0})} = 1$. Similarly, for $\mathbf{w} = (0, \mathbf{1})$, we have the suppliers' optimal mechanism,[17] which for analogous

[15] This follows by 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).

[16] While we do not pursue it here, our approach generalizes directly to the requirement that the mechanism needs to generate a budget surplus of $K \in \mathbb{R}$, which is not more than the maximum budget surplus that any incentive-compatible, individually-rational mechanism can generate. The second-best mechanism that generates K, but otherwise maximizes the same objective, has an allocation rule as defined in Lemma 1, but with $\beta^{\mathbf{w}}$ replaced by $\beta^{\mathbf{w}}_K$, which is a decreasing function of K. Interpreted in this way, we have $\beta^{\mathbf{w}} = \beta^{\mathbf{w}}_0$.

[17] Because there are multiple suppliers, the notion of the suppliers' optimal mechanism deserves brief elaboration. It means that the mechanism produces the same outcome as one in which a single supplier with

reasons does not run a deficit either, implying that $\beta^{(0,\mathbf{1})} = 1$.

Ignoring ties, which occur with probability zero, the allocation rule for incomplete information bargaining is as given in the following lemma:

Lemma 1. *The allocation rule of incomplete information bargaining with bargaining weights* $\mathbf{w}$*,* $\mathbf{Q}^{\mathbf{w}}$*, is such that, for* $i \in \mathcal{N}$*,* $Q_i^{\mathbf{w}}(v, \mathbf{c}) \equiv 1$ *if* $\Phi_{w_B \beta^{\mathbf{w}}}(v) \geq \Gamma_{i, w_i \beta^{\mathbf{w}}}(c_i) = \min_{j \in \mathcal{N}} \Gamma_{j, w_j \beta^{\mathbf{w}}}(c_j)$*, and otherwise* $Q_i^{\mathbf{w}}(v, \mathbf{c}) \equiv 0$*.*

Proof. See Appendix B.

An implication of Lemma 1 is that the probability of trade, and hence social surplus, are increasing in $\beta^{\mathbf{w}}$.

We are left to augment the allocation rule of Lemma 1 with a consistent payment rule. By the *payoff equivalence theorem* (see, e.g., Myerson, 1981; Krishna, 2002; Börgers, 2015), the interim 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. Incentive compatibility implies further that $\underline{v}$ and $\overline{c}$ are the *worst-off types* of the buyer and suppliers, respectively.[18] Thus, to complete the definition of incomplete information bargaining, all that remains to be done is to define these constants.

By standard mechanism design arguments (see Appendix A.1), the expected budget surplus for the mechanism with the allocation rule in Lemma 1, not including the constants reflecting payments to worst-off types, can be written in terms of the allocation rule and virtual types as follows:

$$\pi^{\mathbf{w}} \equiv \sum_{i \in \mathcal{N}} \mathbb{E}_{v, \mathbf{c}} \left[(\Phi(v) - \Gamma_i(c_i)) \cdot Q_i^{\mathbf{w}}(v, \mathbf{c}) \right].$$

Of course, if $\beta^{\mathbf{w}} < 1/\max \mathbf{w}$, then it must be that $\pi^{\mathbf{w}} = 0$, in which case the question of how to allocate the budget surplus is moot. In contrast, if $\beta^{\mathbf{w}} = 1/\max \mathbf{w}$, then $\pi^{\mathbf{w}} \geq 0$. In this case, weighted welfare is maximized when $\pi^{\mathbf{w}}$ is allocated among the buyer and suppliers with bargaining weights equal to $\max \mathbf{w}$. If more than one agent has the maximum bargaining weight, then some surplus sharing rule is required. For example, one might apply equal sharing or distribute the surplus according to Nash bargaining weights—as with Nash bargaining weights, the sharing rule has no social surplus effects. We denote the agents' shares of $\pi^{\mathbf{w}}$ by $\boldsymbol{\eta} \in [0,1]^{n+1}$ satisfying $\eta_i = 0$ for all $i \in \{B\} \cup \mathcal{N}$ such that $w_i < \max \mathbf{w}$ and

cost equal to the minimum of the costs of all the suppliers makes a take-it-or-leave-it offer to the buyer.

[18] That is, for any mechanism satisfying incentive compatibility, $\underline{v} \in \arg\min_{v \in [\underline{v}, \overline{v}]} \mathbb{E}_{\mathbf{c}}[U_B(v, \mathbf{c})]$ and $\overline{c} \in \arg\min_{c \in [\underline{c}, \overline{c}]} \mathbb{E}_{v, \mathbf{c}_{-i}}[U_i(v, \mathbf{c})]$.

$\sum_{i \in \{B\} \cup \mathcal{N}} \eta_i = 1$, yielding interim expected payoffs to the agents' worst-off types of

$$\hat{u}_B(\underline{v}; \mathbf{w}, \boldsymbol{\eta}) = \eta_B \pi^{\mathbf{w}} \quad \text{and} \quad \hat{u}_i(\overline{c}; \mathbf{w}, \boldsymbol{\eta}) = \eta_i \pi^{\mathbf{w}}, \tag{4}$$

for all $i \in \mathcal{N}$.

The outcome of incomplete information bargaining with bargaining weights $\mathbf{w}$ and 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 2 and interim expected payoffs to agents' worst-off types given by (4). Following the mechanism design techniques described in Appendix A.1, we can write these payoffs as stated in the following proposition:

Proposition 1. *Incomplete information bargaining with bargaining weights* $\mathbf{w}$ *and shares* $\boldsymbol{\eta}$ *generates expected payoffs*

$$u_B(\mathbf{w}, \eta_B) = \eta_B \pi^{\mathbf{w}} + \mathbb{E}_v \left[\sum_{i \in \mathcal{N}} \int_{\underline{v}}^{v} \mathbb{E}_{\mathbf{c}} \left[Q_i^{\mathbf{w}}(x, \mathbf{c}) \right] dx \right]$$

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

$$u_i(\mathbf{w}, \boldsymbol{\eta}) = \eta_i \pi^{\mathbf{w}} + \mathbb{E}_{c_i} \left[\int_{c_i}^{\overline{c}} \mathbb{E}_{v, \mathbf{c}_{-i}} \left[Q_i^{\mathbf{w}}(v, x, \mathbf{c}_{-i}) \right] dx \right].$$

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, take as 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 $\mathcal{M}$. Then there exist bargaining weights $\mathbf{w} \in [0,1]^{n+1}$ such that $\mathbf{Q}^{\mathbf{w}}$ is equal to $\tilde{\mathbf{Q}}$ (as shown in the proof of Proposition 2, these weights are derived from the Lagrange multipliers on the constraints that agents' payoffs be at least as great as in $\tilde{\mathbf{u}}$). Further, there exist shares $\boldsymbol{\eta}$ such that the agents' expected payoffs are $\tilde{\mathbf{u}}$ in incomplete information bargaining with bargaining weights $\mathbf{w}$ and shares $\boldsymbol{\eta}$. Conversely, because no money is "left on the table," any expected payoffs from incomplete information bargaining are Pareto undominated among payoffs resulting from mechanisms in $\mathcal{M}$.

Proposition 2. *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* $\mathcal{M}$ *if and only if there exist bargaining weights* $\mathbf{w}$ *and shares* $\boldsymbol{\eta}$ *such that* $\mathbf{Q}^{\mathbf{w}} = \tilde{\mathbf{Q}}$ *and* $(u_B(\mathbf{w}, \boldsymbol{\eta}), u_1(\mathbf{w}, \boldsymbol{\eta}), \ldots, u_n(\mathbf{w}, \boldsymbol{\eta})) = \tilde{\mathbf{u}}$.

Proof. See Appendix B.

As we show in Appendix C, incomplete information bargaining includes the k-double auction of Chatterjee and Samuelson (1983) as a special case (when $n = 1$ and agents draw their types from the 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.

3.2 Social-surplus-increasing equalization of bargaining weights and scope for countervailing power

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 $\mathbf{w}$ with $w_i = w$ for all $i \in \{B\} \cup \mathcal{N}$—and scope for countervailing power.

In particular, denoting by $W^* \equiv \mathbb{E}_{v,\mathbf{c}}[W_{\mathbf{Q}^*}(v, \mathbf{c})]$ the value of the planner's objective under the planner's optimal allocation rule and by $W^{\mathbf{w}} \equiv \mathbb{E}_{v,\mathbf{c}}[W_{\mathbf{Q}^{\mathbf{w}}}(v, \mathbf{c})]$ the value of the planner's objective under the allocation rule chosen by the market, 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 $w \in (0, 1]$.[19] Hence, for any $w \in (0, 1]$, we have $W^* = W^{(w,\dots,w)}$.

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, $\mathbf{Q}^{\mathbf{w}}(v, \mathbf{c}) \neq \mathbf{Q}^*(v, \mathbf{c})$ for all $(v, \mathbf{c})$ in an open subset of $[\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]^n$. As stated in the next proposition, specific conditions are required for the planner *not* to prefer an equalization of bargaining weights. Of course, there is no benefit to the planner from equalization if the bargaining weights are already equalized, but even when the weights are not all the same, there is also no benefit to the planner if the market has full trade, i.e.,

$$\min_{i \in \mathcal{N}} \Gamma_{i,w_i\beta^{\mathbf{w}}}(\overline{c}) \leq \Phi_{w_B\beta^{\mathbf{w}}}(\underline{v}), \tag{5}$$

which implies that $\beta^{\mathbf{w}} = 1/\max \mathbf{w}$, and if there is sufficient symmetry among the suppliers that it is always the lowest-cost supplier that trades. Specifically, the suppliers must have equal bargaining weights, $w_1 = \cdots = w_n = w$, and must either have the maximum bargaining weight, $w = \max \mathbf{w}$, so that $\Gamma_{i,w\beta^{\mathbf{w}}}(c_i) = c_i$, or it must be that $G_1 = \cdots = G_n$, so that the supplier with the lowest weighted virtual cost is also the supplier with the lowest actual

[19] To see this, note that $W_{\mathbf{Q},\mathbf{M}}^{(w,\dots,w)}(v, \mathbf{c}) = w(W_{\mathbf{Q}}(v, \mathbf{c}) - R_{\mathbf{M}}(v, \mathbf{c}))$, 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.

cost.[20]

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) (5) *holds;*

(ii) for all $i \in \mathcal{N}$, $w_i = w$;

(iii) either $w_B < w$ *or for all* $i \in \mathcal{N}$, $G_i = G$.

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, for any $n \geq 2$ and independent of any distributional assumptions, there exist asymmetric bargaining weights $\mathbf{w}$ such that the planner benefits from an equalization of bargaining weights, i.e., there exist $\mathbf{w}$ such that $W^{\mathbf{w}} < W^*$. The same is true for $n = 1$, unless $\overline{c} \leq \Phi(\underline{v})$, and $\underline{v} \geq \Gamma_1(\overline{c})$.[21] This shows that, quite generally, equalization of bargaining power increases social surplus. Some of the benefits that the planner obtains from more equal bargaining weights stem from an equalization of bargaining weights among suppliers, which eliminates socially wasteful discrimination among suppliers based on differently weighted virtual costs. 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. Adhering to this terminology, we say that there is *scope for countervailing power* in a market if the planner prefers an equalization of bargaining power across the two sides of the market.

To analyze the scope for countervailing power, we therefore focus on the case in which suppliers have symmetric bargaining weights, that is, $w_1 = \cdots = w_n$. In this case, payoffs are pinned down by the bargaining differential between the buyer side of the market and the supplier side of the market as captured by Δ given by $\Delta \equiv \frac{w_B - w_1}{\max\{w_B, w_1\}} \in [-1, 1]$. This

[20]That ex ante symmetry among suppliers implies that there is no inefficiency in production when the production decision is based on (equally weighted) virtual costs rather than costs hinges on the assumption that the virtual cost functions are increasing. Without that assumption, the weighted virtual cost 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 random tie-breaking.

[21]These distributional assumptions are restrictive in the sense that they are not satisfied if the supports of the buyer's and suppliers' type distributions overlap because $\Phi(v) < v$ for any $v < \overline{v}$ and $\Gamma_i(c) > c$ for any $c > \underline{c}$. Further, the conditions fail in many cases even when there is no overlap—for example, if the supplier draws its cost from the uniform distribution on $[0, 1]$ and $F(v)$ is uniform on $[\underline{v}, \underline{v} + 1]$, then the conditions hold if and only if $\underline{v} \geq 2$. If $\overline{c} > \Phi(\underline{v})$, 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^*.

allows us to write every agent's ex ante expected surplus as a function of Δ only. Denoting by $u_S(\Delta) \equiv \sum_{i\in\mathcal{N}} u_i(\Delta)$ aggregate expected supplier surplus when the bargaining weights of the two sides differ by Δ, we can trace out the frontier of expected buyer and expected aggregate supplier surplus, which we refer to as the *Williams frontier* in honour of Williams (1987), who first studied this frontier in a model with one buyer and one supplier. Formally, the Williams' frontier is defined as $\mathcal{F} \equiv \{(u_S(\Delta), u_B(\Delta)) \mid \Delta \in [-1, 1]\}$, with associated mapping $\omega : [u_S(1), u_S(-1)] \to [u_B(-1), u_B(1)]$ defined by $\omega(u) = \max\{y \mid (u, y) \in \mathcal{F}\}$. For the special case of $n = 1$ and uniformly distributed types, the Williams frontier coincides with the payoff frontier for the k-double auction, which is depicted in Figure 1. The frontier there is concave to the origin, which as shown next is a general property:

Proposition 4. *The Williams frontier is concave to the origin, i.e., ω is strictly decreasing and concave, and strictly concave if there is at most one value of Δ that achieves the first-best.*

Proof. See Appendix B.

Building on this, the concavity of the Williams frontier has the following implication:

Corollary 1. *A change in the buyer-side and supplier-side bargaining weights that moves them closer together, i.e., that moves Δ closer to zero, weakly increases social surplus: if $\Delta' < \Delta \leq 0$ or $0 \leq \Delta < \Delta'$, then $u_B(\Delta') + u_S(\Delta') \leq u_B(\Delta) + u_S(\Delta) \leq u_B(0) + u_S(0)$.*

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.[22] 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 budget surplus of the mechanism, which is zero. What changes as the bargaining weights change is the allocation rule, which

[22] 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.

transitions from, say, the buyer-optimal one via the second-best to the one that is optimal for the suppliers. 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 Merger review

In this section, we analyze horizontal mergers and vertical integration. Throughout this section, we assume that after horizontal or vertical integration, the integrated entity can efficiently solve its internal agency problem, which is a standard assumption.[23] We evaluate outcomes from an ex ante perspective, that is, before firms' types are realized. This includes the profitability of mergers or vertical integration.

4.1 Horizontal mergers

As do Farrell and Shapiro (1990), we model a horizontal merger without cost synergies as allowing merging suppliers to re-optimize by producing using the lower of their two costs.[24] Following this approach, the merged supplier's cost is drawn from the distribution of the minimum of two independent cost draws, one from each of the two merging suppliers' distributions. That is, a merged entity formed from suppliers 1 and 2 draws its cost from the distribution

$$\hat{G}(c) \equiv 1 - (1 - G_1(c))(1 - G_2(c)), \tag{6}$$

with associated weighted virtual type function $\hat{\Gamma}_a(c)$, which we assume to be increasing.

We say that the merged entity's virtual cost dominates the pre-merger virtual costs if for all $a \in [0, 1)$ and c_i and c_j with $\min\{c_i, c_j\} \in (\underline{c}, \overline{c})$, we have

$$\hat{\Gamma}_a(\min\{c_i, c_j\}) > \min\{\Gamma_a(c_i), \Gamma_a(c_j)\}. \tag{7}$$

It is straightforward to show that (7) is satisfied if the merging suppliers have symmetric cost distributions, $G_1 = G_2$.

In principle, a merger need not affect agents' bargaining powers. But it does affect the productive power of the merging suppliers. If suppliers 1 and 2, each drawing a cost from the

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

[24] 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).

distribution G, merge, then the merged entity draws its cost from the distribution $1-(1-G)^2$ and has an associated virtual cost that dominates the pre-merger virtual costs in the sense of (7). This means that, in the absence of any bargaining power changes (i.e., the merged entity's bargaining power is $w = w_1 = w_2$), the merger causes the probability of trade by the merging suppliers to weakly decrease. Further, because the merger tightens the no-deficit constraint, the decrease is strict unless the bargaining processes treats the merging suppliers nondiscriminatorily, i.e., has weighted virtual costs equal to costs, both before and after the merger.[25] Note that a merger that does not change the probabilities of trade also does not affect the buyer's payoff if $\eta_B = 0$. Consequently, we have the following result:

Proposition 5. *A merger between two symmetric suppliers, each with bargaining weight w before merger, which is also the merged entity's bargaining weight, is neutral for social surplus if there is no discrimination against the merged entity. If, in addition, $\eta_B = 0$, then it is also neutral for the buyer. Otherwise, the merger reduces expected buyer surplus and expected social surplus.*

According to Proposition 5, the productive power effect of a merger of symmetric suppliers harms the buyer and society, except in boundary cases. Proposition 5 generalizes the insights from Loertscher and Marx (2019) to a setting in which incomplete information pertains to both sides of the market and bargaining power is not restricted to be with the buyer.

In addition to changing the productive power 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 that somehow "levels the playing field" endows merging parties with *countervailing power* is based on this very conception. It finds support in the empirical literature (Ho and Lee, 2017; Bhattacharyya and Nain, 2011) and, as mentioned, features prominently in concurrent antitrust debates and cases.[26] Nonetheless, a major obstacle to analyzing the effects of countervailing power in existing modeling approaches is that these take the efficiency of the price-formation process as given. This is true for all oligopoly models, in which agents on one side of the market (typically buyers) are assumed to be price-takers and also applies to randomized take-it-or-leave-it offers.

In contrast, as illustrated in Figure 1 and stated in Corollary 1, with incomplete information, a change in bargaining weights has an impact on social surplus because the efficiency

[25] Letting $\hat{\beta}^{\mathbf{w}}$ be the post-merger budget parameter, holding bargaining power fixed, the merger causes the probability of trade by the merging suppliers to decrease if and only if $w\hat{\beta}^{\mathbf{w}} < 1$, in which case the relevant weighted virtual cost function is not the identity function and so is discriminatory.

[26] 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).

of the mechanism varies with bargaining weights. Consequently, a merger that results in buyer-side and supplier-side bargaining power 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 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.[27] Indeed, Figure 2(b) provides an example in which countervailing power restores 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.

(a) Social-surplus-enhancing countervailing power

(b) First-best-achieving countervailing power

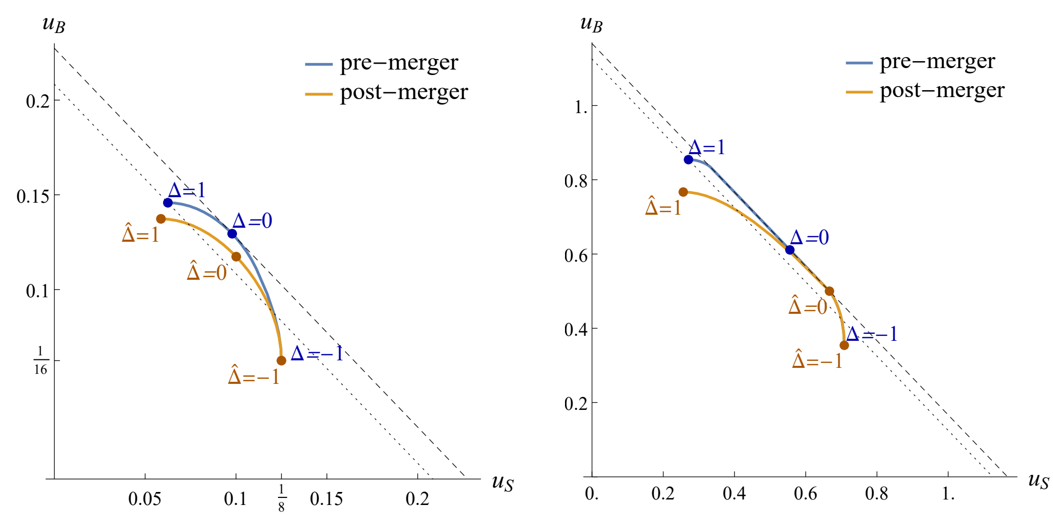

Figure 2: Williams frontiers for the case of 2 symmetric pre-merger suppliers (blue) and 1 post-merger supplier (orange). We assume that in the pre-merger market, $w_1 = w_2$, and that the merged entity also has bargaining weight w_1. The bargaining power differential for the pre-merger market is denoted by $\Delta \equiv (w_B - w_1)/\max\{w_B, w_1\}$ and for the post-merger market by $\hat{\Delta} \equiv (w_B - w_1)/\max\{w_B, w_1\}$. We assume shares $\boldsymbol{\eta}$ that are symmetric among 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 $[1, 2]$. In each case, the merged entity's cost distribution is given by (6).

We summarize with the following result:

Corollary 2. *A merger between two symmetric suppliers that only has productive power effects and reduces social surplus is more harmful than a merger between the same two sup-*

[27] 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.

pliers 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.

Corollary 2 demonstrates that there is the possibility of a countervailing power defense for a merger. A merger that only involves productive power effects reduces social surplus if the first-best is not possible post merger. In contrast, a merger that causes bargaining weights to shift in favor of the suppliers may improve expected social surplus despite the adverse productive power effects. Of course, a merger with countervailing power is bad for the buyer for two reasons: competition among suppliers is reduced *and* the remaining suppliers have increased bargaining power. Thus, merger review based on a buyer-surplus standard would never be moved by a countervailing power defense.[28] In contrast, merger review based on a social-surplus standard may well be.

Our analysis allows us to identify necessary conditions for a countervailing power defense. 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 buyer would need to have greater bargaining power than the suppliers prior to the merger, so that movement towards the equalization of bargaining power is possible. Third, the buyer would need to retain at least some bargaining power following the merger—buyer power would need to diminish, but not vanish—so that society is not simply trading a dominant buyer for dominant suppliers.

This brings to mind 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).[29] Our conclusions are consistent with that view insofar as the buyer must have power before a merger and retain at least some power after a merger in order for a countervailing power defense to make economic sense.

The necessary conditions for a countervailing power defense raise the question of how one would ascertain that a buyer has bargaining power. For example, if a market is characterized as a k-double auction, then evidence of buyer power would be that transactions always oc-

[28] In the alternative, but analogous, setting of a sale auction with one supplier and multiple buyers, a merger of buyers harms the supplier and possibly benefits (all) buyers. This gives rise to the possibility of buyer-surplus-increasing mergers, which might be viewed as procompetitive. To increase social surplus, the merger would have to reduce the supplier's bargaining power without eliminating it.

[29] The EC merger guidelines 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).

cur at the buyer's price.[30] For a procurement auction, evidence consistent with buyer power but inconsistent with the absence of buyer power 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.[31]

4.2 Vertical integration

We now analyze vertical integration between the buyer and one of the suppliers. Following vertical integration between the buyer and supplier i, incomplete information bargaining works as before, with the exception that the vertically integrated firm now acts as a buyer with value $y = \min\{v, c_i\}$, whose distribution we denote by $\hat{F}(y) \equiv F(y) + G_i(y)(1 - F(y))$, which we assume to exhibit increasing virtual value. If there is no trade between the vertically integrated firm and the nonintegrated suppliers, then the integrated firm's payoff is equal to $\max\{0, v - c_i\}$ due to internal sourcing, and the nonintegrated suppliers have payoffs of zero. Observe that for $y \in [\underline{v}, \overline{v}]$, we have $\hat{F}(y) \geq F(y)$. Thus, vertical integration makes the buyer productively less powerful.

Consider first a bilateral trade setting with overlapping supports before integration (i.e., $n = 1$ and $\underline{v} < \overline{c}$). Because the first-best is impossible when the buyer and supplier are independent entities, it follows immediately from our assumption that the integrated entity can resolve the internal agency problem that vertical integration can increase social surplus:

Proposition 6. *With one supplier and overlapping supports, vertical integration increases social surplus (to the first-best) regardless of bargaining weights.*

By Proposition 6, vertical integration can increase social surplus and enable the first-best by essentially eliminating a Myerson-Satterthwaite problem. However, as we show next, it can also create one.

[30] This property does not hinge on particular distributional assumptions. For $k = 1$, the buyer's and seller's optimal bids are $\Gamma_1^{-1}(v)$ and c, respectively, while for $k = 0$, they are v and $\Phi^{-1}(c_1)$. Hence, for $k = 1$ ($k = 0$) the k-double auction is the mechanism that is optimal for the buyer (seller) for any distributions F and G_1 with positive densities on their supports. (If Φ 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).)

[31] 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 its value, 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. (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." A buyer without power purchases from the lowest-cost supplier.

To this end, assume that there are multiple suppliers and consider the case with nonoverlapping supports. The latter assumption implies that prior to vertical integration, the first-best is possible and, indeed, occurs if the pre-integration bargaining weights are symmetric. Hence, vertical integration cannot possibly increase social surplus. This leaves the question of whether vertical integration could be neutral. The following proposition shows that the answer is negative.

Proposition 7. *With two or more suppliers and nonoverlapping supports, vertical integration decreases social surplus whenever pre-integration bargaining weights are symmetric, regardless of post-integration bargaining weights.*

Proof. See Appendix B.

Proposition 7 provides a clear-cut case in which vertical integration is harmful from the perspective of society. This result, as well as the result in Proposition 6, is robust in that it does not depend on specific assumptions about distributions or beliefs of agents. Indeed, because there is always a dominant strategy implementation of the price-formation 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.[32] Of course, our results do rely, inevitably, on support assumptions.

At the heart of both Propositions 6 and 7 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 literature, however, there is now a new effect, namely that the market with the remaining suppliers becomes less efficient. 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.[33]

[32] 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).

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

Connecting bargaining breakdown with vertical integration

A pervasive feature of real-world bargaining is that negotiations often break down. 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, to, possibly, Brexit.[34] 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.

In incomplete information bargaining, negotiations 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, by the Myerson-Satterthwaite theorem, 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.

Assuming that real-world negotiations are appropriately captured by incomplete information bargaining, one can use observed frequencies of negotiation breakdowns to back out the parameters of the distributions from which the buyer and the supplier draw their types. For purpose of illustration, we assume in the following that the buyer's value and the sellers' cost are drawn from parameterized distributions

$$F(v) = 1 - (1-v)^{1/\kappa} \quad \text{and} \quad G_i(c) = c^{1/\kappa_i}, \tag{8}$$

with support $[0,1]$, where the parameters κ and κ_i are positive real numbers and have the interpretation of "capacities" insofar as larger values of κ and κ_i 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.

Pre-integration market conditions, particularly the probability of breakdown, can also be

[34] 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.

used to gauge the social surplus effects of vertical integration. To see this, assume that $n = 2$ in the market before integration by the buyer with supplier 1, $\mathbf{w}$ and $\boldsymbol{\eta}$ are symmetric, and F and G_i are given by (8). As an identifying assumption, stipulate that $(\kappa_1 + \kappa_2)/2 = 1$, that is, the suppliers' capacities are equal to one on average (alternatively one might use, e.g., margin data for identification). Figure 3(a) shows the results of the calibration of these parameterized distributions given data on supplier market shares and the probability of bargaining breakdown.

(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 social surplus following vertical integration

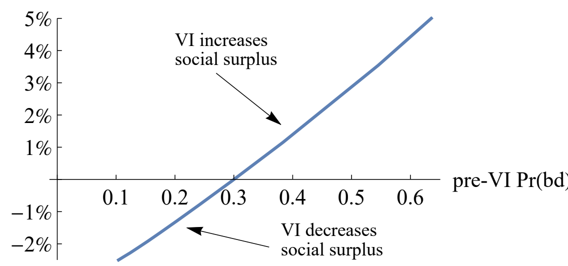

Figure 3: Interaction between the pre-integration breakdown probability and the effect of vertical integration on social surplus. Panel (a): Calibration of distributional parameters based on market shares and breakdown probabilities assuming that $n = 2$, $\mathbf{w}$ and $\boldsymbol{\eta}$ are symmetric, $F(v)$ and $G_i(c)$ are given by (8), and $(\kappa_1+\kappa_2)/2 = 1$. Panel (b): Change in 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).

Now consider the effect on social surplus of vertical integration assuming that there are two pre-integration suppliers with equal market shares and assuming that bargaining weights $\mathbf{w}$ and shares $\boldsymbol{\eta}$ are symmetric both before and after integration. As illustrated in Figure 3(b), in markets where before integration the probability of breakdown is low, the change in social surplus from vertical integration is negative. In that case, the reduced efficiency of price formation with the independent supplier dominates the gain in efficiency associated with internal transactions, and vertical integration reduces social surplus. In contrast, when the probability of breakdown is high prior to integration, then the increased efficiency of internal transactions dominates, and social surplus increases as a result of vertical integration.

5 Investment

Investment incentives feature prominently, and at times controversially, in concurrent policy debates,[35] 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 buyer and each supplier to improve (or more generally change) its type distribution by investing e_j at cost $\Psi_j(e_j)$ for agent $j \in \{B\} \cup \mathcal{N}$. Consistent with G-H-M, we assume that investments are not contractible.[36] 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 actual investments. We assume that the buyer and supplier 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 $\mathbf{Q}^{FB}(v, \mathbf{c})$, where for $i \in \mathcal{N}$, $Q_i^{FB}(v, \mathbf{c}) \equiv 1$ if $v \geq c_i = \min_{j \in \mathcal{N}} c_j$ and $Q_i^{FB}(v, \mathbf{c}) \equiv 0$ otherwise. Then, for a given realization of types, first-best welfare is $W^{FB}(v, \mathbf{c}) \equiv \sum_{i \in \mathcal{N}} (v - c_i) Q_i^{FB}(v, \mathbf{c})$. We let $\overline{\mathbf{e}}$ denote first-best investments, which are a solution to the planner's first-best investment problem, given by $\max_{\mathbf{e}} \mathbb{E}_{v,\mathbf{c}|\mathbf{e}} \left[W^{FB}(v, \mathbf{c}) \right] - \sum_{i \in \{B\} \cup \mathcal{N}} \Psi_i(e_i)$.

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

[35] 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).

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

realization $(v, \mathbf{c})$, supplier i's VCG payoff is $W^{FB}(v, \mathbf{c}) - W^{FB}(v, \overline{c}, \mathbf{c}_{-i})$ plus possibly a constant. Likewise, the buyer's payoff is $W^{FB}(v, \mathbf{c}) - W^{FB}(\underline{v}, \mathbf{c})$, plus possibly a constant.

Taking expectations over $(v, \mathbf{c})$, and noticing that $W^{FB}(v, \overline{c}, \mathbf{c}_{-i})$ is independent of supplier i's type and its distribution, and so independent of e_i it follows that each supplier i's problem at the investment stage, taking as given that the other agents choose investments $\overline{\mathbf{e}}_{-i}$, is $\max_{e_i} \mathbb{E}_{v,\mathbf{c}|e_i,\overline{\mathbf{e}}_{-i}} \left[W^{FB}(v, \mathbf{c})\right] - \Psi_i(e_i)$. An analogous optimization problem applies to the buyer's choice of e_B, noting that $W^{FB}(\underline{v}, \mathbf{c})$ is independent of the buyer's type and its distribution, and so independent of e_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 8 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 investment e, for $i \in \mathcal{N}$, let $G_i(c; e)$ and $F(v; e)$ denote supplier i's and the buyer's type distributions, respectively, with virtual type functions assumed to be monotone. Sufficient conditions for that converse to hold are: for all $i \in \{B\} \cup \mathcal{N}$,

$$\Psi_i'(0) = 0 \;\text{ and for } e > 0,\; \Psi_i'(e) > 0 \;\text{ and }\; \Psi_i''(e) > 0; \tag{9}$$

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

$$\frac{\partial G_i(c; e)}{\partial e} > 0 \;\text{ and }\; \frac{\partial F(v; e)}{\partial e} < 0; \tag{10}$$

and either the type distributions have overlapping supports, $\underline{v} < \overline{c}$, or for all $i \in \mathcal{N}$ and all $c \in [\underline{c}, \overline{c}]$,

$$G_i(c; \overline{e}_i) \equiv G(c). \tag{11}$$

Conditions (9) and (10) imply that the first-best investments $\overline{\mathbf{e}}$ are positive and determined by first-order conditions. With overlapping supports or (11), the efficiency of the buyer's allocation rule implies the efficiency of each supplier's allocation rule.[37]

Proposition 8. *First-best investments are a Nash equilibrium outcome of the simultaneous investment game if incomplete information bargaining is efficient. Conversely, assuming that* (9) *and* (10) *hold and that either supports overlap or* (11) *holds, if first-best investments are a Nash equilibrium outcome, then incomplete information bargaining is efficient.*

[37]Proposition 8 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.

Proof. See Appendix B.

As shown in Proposition 8, when incomplete information bargaining is efficient, the agents' Nash equilibrium investment choices are first-best investments. Because private information protects agents from hold-up,[38] efficient incomplete information bargaining implies efficient investments.[39] Intuitively, given that the allocation rule is efficient, each agent is the 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 Corollary 2 and Proposition 8 allows us to connect investment with countervailing power. By Proposition 8, countervailing power cannot only increase social surplus, holding investments fixed, as stated in Corollary 2, but it can also improve investments to the first-best level. Proposition 8 thus provides an additional channel—investments—through which countervailing power can increase social surplus.

While Proposition 8 focuses on investments that improve agents' own types, the first part of Proposition 8 continues to hold if each supplier can invest in the "quality" of its product, thereby increasing the value of its product to the buyer.[40] 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. Che et al. (2017) consider the latter case and find that the buyer always wants to depart from ex post efficiency in order to boost ex ante investment by suppliers.

Proposition 8 also allows us to analyze the effect of vertical integration on investment. We assume that vertical integration does not affect the cost of investment for the integrated firm, so if the buyer and supplier i integrate and invest $e_B + e_i$, the cost of investment is $\Psi_B(e_B) + \Psi_i(e_i)$. With one supplier in the pre-integration market and overlapping supports, incomplete information bargaining is inefficient, which under conditions (9) and (10), implies

[38] 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.

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

[40] 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.

that equilibrium investments are inefficient. But, by assumption, the allocation is efficient after vertical integration, which by Proposition 8 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 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 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, with* $n = 1$ *and overlapping supports, vertical integration promotes efficient investment; but with nonoverlapping supports and* $n \geq 2$ *suppliers whose distributions satisfy* (11)*, vertical integration disrupts efficient investment if bargaining is efficient prior to vertical integration (which occurs, for example, with symmetric bargaining weights).*

6 Extensions

In this section, we extend the model two ways and then discuss implementation.

6.1 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

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 $x_i \geq 0$ for supplier i. 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 price-formation mechanism with weights $\mathbf{w}$ is the solution to

$$\max_{\langle \mathbf{Q},\mathbf{M}\rangle \in \mathcal{M}} \mathbb{E}_{v,\mathbf{c}}\left[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(v,c)\right] \text{ s.t. } \eta_B \pi^{\mathbf{w}} \geq x_B \text{ and for all } i \in \mathcal{N},\ \eta_i \pi^{\mathbf{w}} \geq x_i.$$

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 where $\beta^{\mathbf{w}}$ is the largest $\beta \in [0, \max \mathbf{w}]$ such that

$$\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}} (\Phi(v) - \Gamma_i(c_i)) \cdot \mathbf{1}_{\Phi_{w_B\beta}(v) \geq \Gamma_{i,w_i\beta}(c_i) = \min_{j\in\mathcal{N}} \Gamma_{j,w_j\beta}(c_j)}\right] \geq x_B + \sum_{i\in\mathcal{N}} x_i, \qquad (12)$$

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

Consider the case of symmetric suppliers in this setup. As the number of suppliers increases, the range of outside options that can be accommodated increases. As the suppliers' outside option increases, the expected social surplus decreases—the need to generate revenue for the suppliers distorts the overall market outcome—and eventually the suppliers' payoffs exceed that of the buyer, even if the buyer has all the bargaining power. Further, if the suppliers' outside option is sufficiently large, then the buyer and society are better off when the number of suppliers is reduced below the maximum number sustainable in the market.

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 $y_i \geq 0$, the support of supplier i's cost distribution is $[\underline{c}_i, \overline{c}_i]$ with $\underline{c}_i = \underline{c} + y_i$ and $\overline{c}_i = \overline{c} + y_i$. If $G_i(c)$ is i's cost distribution without the outside option, then given outside option y_i, its cost distribution is $G_i^o(c) = G_i(c - y_i)$, with density $g_i^o(c) = g_i(c - y_i)$ and support $[\underline{c}_i, \overline{c}_i]$. 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 $y_i > 0$ and $y_B > 0$, we have $G_i^0(c) \leq G_i(c)$ 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 $G_i(c)/g_i(c)$ is increasing in c and $(1-F(v))/f(v)$ is decreasing in v. To see this, let us focus on supplier i. The arguments for the buyer (and of course all other suppliers) are analogous. We denote the weighted virtual cost of supplier i when it has outside option y_i by

$$\Gamma^o_{i,a}(c) \equiv c + (1-a)\frac{G_i(c-y_i)}{g_i(c-y_i)} = \Gamma_{i,a}(c-y) + y < \Gamma_{i,a}(c), \tag{13}$$

where the inequality holds for all $a < 1$ because the monotone hazard rate assumption implies that $\Gamma'_{i,a}(c) > 1$ for all $a < 1$. This in turn has two, somewhat subtle implications. Let z be the threshold for supplier i to trade when its outside option is zero, i.e., keeping z fixed, i trades if and only if $\Gamma_{i,a}(c) \leq z$. (Note that z will be the minimum of the buyer's weighted virtual value and the smallest weighted virtual cost of i's competitors, but this does not matter for the argument that follows.) Assuming that $a < 1$ and $y_i < \overline{c} - \underline{c}$, which implies that $\underline{c}_i < \overline{c}$, it follows that there are costs $c \in [\underline{c}_i, \overline{c}]$ and thresholds z such that supplier i trades when it has the outside option and not without it, that is, $\Gamma^o_{i,a}(c) < z < \Gamma_{i,a}(c)$. 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 i's outside option and hence our argument extends directly once one integrates over z and its density.) Given y_i, supplier i trades if and only if its cost c is below $\tau(y)$ satisfying $\Gamma^o_{i,a}(\tau(y)) = z$. Using (13), this is equivalent to $\Gamma_{i,a}(\tau(y) - y) + y = z$, which in turn is equivalent to $\tau(y) = \Gamma^{-1}_{i,a}(z - y) + y$, whose derivative for $a < 1$ satisfies

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

where the inequalities follow because $\Gamma'_{i,a}(c) > 1$. This implies that, for a fixed z, the probability that i trades decreases in y. To see this, notice that this probability is $G^o_i(\tau(y)) = G_i(\tau(y) - y)$, whose derivative with respect to y is $g_i(\tau(y) - y)(\tau'(y) - 1) < 0$. 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 i 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_i improves supplier i's payoff after its value or cost has been realized, supplier i's ex ante expected payoff decreases in y_i. Moreover, because an increase in y_i worsens supplier i's distribution, the revenue constraint becomes (weakly) tighter, implying a decrease in $\beta^{\mathbf{w}}$, which further reduces supplier i's expected payoff.

6.2 Bargaining externalities

To allow for and investigate bargaining externalities, we now return to the model without investment and with outside options of zero, but we generalize it to allow the buyer to have preferences over suppliers and to have demand for $D \in \{1, 2, \dots\}$ objects. 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 i when the buyer's type is v is $\theta_i v$. Thus, under (ex post) efficiency, trade should occur between the buyer and supplier i if and only if $\theta_i v - c_i$ is positive and among the D highest values of $(\theta_j v - c_j)_{j \in \mathcal{N}}$. The problem is trivial if $\max_{i \in \mathcal{N}} \theta_i \overline{v} \leq \underline{c}$ 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 $D = 1$; and (iii) a version of the Shapley-Shubik model in which the buyer has demand for multiple products of the suppliers by setting $D > 1$. For an extension of the one-to-many setup that encompasses additional models, see Appendix D.

We define the virtual surplus $\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}$ associated with trade between the buyer and supplier i, accounting for the agents' bargaining weights $\mathbf{w}$ and the buyer's preferences $\boldsymbol{\theta}$, with $\beta^{\mathbf{w},\boldsymbol{\theta}}$ defined analogously to before as $\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}(v, c_i) \equiv \theta_i \Phi_{w_B \beta^{\mathbf{w},\boldsymbol{\theta}}}(v) - \Gamma_{i, w_i \beta^{\mathbf{w},\boldsymbol{\theta}}}(c_i)$. Let $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv (\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}(v, c_i))_{i \in \mathcal{N}}$ and denote by $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})_{(D)}$ the D-th highest element of $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})$. As before, in order to save notation, we ignore ties.

Lemma 2. *In the generalized setup with buyer preferences $\boldsymbol{\theta}$, incomplete information bargaining with weights $\mathbf{w}$ has the allocation rule for $i \in \mathcal{N}$, $Q_i^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv 1$ if $\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}(v, c_i) \geq \max\{0, \boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})_{(D)}\}$, and otherwise $Q_i^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv 0$.*

Proof. See Appendix B.

We can now use this generalized setup to analyze bargaining externalities between suppliers. If $D < n$, then one effect of an increase in θ_i is that agents other than i are less likely to be among the at-most D agents that trade. In contrast, if $D \geq n$ and $\beta^{\mathbf{w},\boldsymbol{\theta}} < 1/\max \mathbf{w}$, then the probability that supplier i trades, $\Pr(\theta_i \Phi_{w_B \beta^{\mathbf{w},\boldsymbol{\theta}}}(v) \geq \Gamma_{i,w_i\beta^{\mathbf{w},\boldsymbol{\theta}}}(c_i))$, does not depend on the preference parameters of the other suppliers except through their effect on $\beta^{\mathbf{w},\boldsymbol{\theta}}$. If $\beta^{\mathbf{w},\boldsymbol{\theta}} < \max \mathbf{w}$, then an increase in a rival supplier's preference parameter causes an increase in $\beta^{\mathbf{w},\boldsymbol{\theta}}$, which increases the probability of trade and so benefits the supplier. Thus, we have the following result:

Proposition 9. *In the generalized setup with bargaining weights* $\mathbf{w}$ *and buyer preferences* $\boldsymbol{\theta}$*, if* $D \geq n$ *and* $\beta^{\mathbf{w},\boldsymbol{\theta}} < 1/\max \mathbf{w}$*, then an increase in the preference parameter for one supplier increases the payoffs for all suppliers.*

The result of Proposition 9 does not necessarily extend to the case with $D < n$, as shown in the example of Appendix E.

6.3 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.[41] As another case in point, to support the model of perfectly competitive markets, one might view price-taking buyers and suppliers as submitting demand and supply schedules to a (fictitious) Walrasian auctioneer who then sets market clearing prices. Similarly, in the Cournot model, one might view firms as submitting quantities to an auctioneer or market maker who sets the market clearing price.[42] Under assumptions on the alternation of offers and taking the limit as the time between offers goes to zero, Rubinstein bargaining delivers Nash bargaining outcomes (Rubinstein, 1982); and under additional assumptions, including conditions on firms' marginal contributions and passive beliefs, the limit of an alternating-offers game approximates Nash-in-Nash outcomes (Collard-Wexler et al., 2019).

In light of this, it is perhaps useful to note that, as mentioned above and discussed in Appendix C, for the case of one supplier and uniformly distributed types, the k-double auction

[41] 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.

[42] Microfoundations of the Cournot model along the lines of Kreps and Scheinkman (1983), while dispensing with the assumption of an auctioneer, maintain the assumption of an exogenously given price-formation process by postulating that firms first choose capacities and then prices.

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 G, 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 now, provide micro-foundations for incomplete information bargaining. Specifically, for general distributions and any number of suppliers, the incomplete information bargaining outcome arises in equilibrium in an extensive-form game involving a buyer, suppliers, and a fee-setting broker.

Building on the model of Loertscher and Niedermayer (2019), we define the *fee-setting extensive-form game* to have one buyer, $n \geq 1$ 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. 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 10. *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. See Appendix B.

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.[43]

This is reminiscent of the role of intermediaries in the wholesale used car market as described by Larsen (2020). 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.

7 Related literature

We use the Myersonian mechanism design machinery (Myerson, 1981) to elicit—search for, as it were—agents' private information and determine prices. Our framework builds on the bilateral trade model of Myerson and Satterthwaite (1983), augmented by bargaining weights and multiple suppliers. Thereby, it combines elements of Myerson and Satterthwaite (1983), Williams (1987), and Gresik and Satterthwaite (1989). Specifically, our procurement model allows for multiple suppliers without imposing restrictions on the supports of

[43]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).

the buyer's value and the suppliers' costs other than assuming that all cost distributions have the same support.[44] We generalize Williams' approach of maximizing an objective that assigns differential weights in a bilateral trade problem by allowing for multiple agents. In light of the quote from *New Palgrave Dictionary* in the introductory paragraph, our paper reinterprets Myerson and Satterthwaite (1983) as a bilateral monopoly problem, extends it to allow for bargaining weights and multiple agents on one side of the market, and shows that it is tractable and has all the required features.[45] In particular, inherent to the independent private values setting is the key economic tradeoff between rent extraction and social surplus.[46]

There has also been a recent upsurge of interest in bargaining (see, for example, Larsen, 2020; Backus et al., 2020, 2019; Zhang et al., 2019; Decarolis and Rovigatti, 2020), and buyer power (see, for example, Snyder, 1996; Nocke and Thanassoulis, 2014; Caprice and Rey, 2015; Loertscher and Marx, 2019). Larsen and Zhang (2018) 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.[47] 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.[48] 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,[49]

[44] While Gresik and Satterthwaite (1989) also allow for multiple buyers, they restrict attention to identical cost distributions. In that regard, our setup thus shares similarities with the optimal auction setting of Myerson (1981), with the important difference that our setup has two-sided private information.

[45] For experimental results consistent with the incomplete information bargaining, see Valley et al. (2002, Fig. 3.A). See Larsen (2020) on the first-best and second-best frontiers for wholesale used cars.

[46] Appendix F provides additional discussion of the properties of the independent private values paradigm.

[47] 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.

[48] While the empirical literature examining multilateral bargaining focuses on fixed quantities or linear tariffs, Rey and Vergé (2019) allow for non-linear 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.

[49] 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."

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.[50] 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 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).[51] As discussed, the predictions could hardly differ more starkly because with incomplete (complete) information efficient bargaining implies efficient (inefficient) investment.[52] 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.

[50]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.

[51]Nocke and Thanassoulis (2014) provide model within the paradigm of efficient, complete information bargaining in which there is scope for countervailing power because bargaining power can mitigate frictions due to credit constraints.

[52]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), 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.

8 Conclusions

We provide an incomplete information bargaining model in which the possibility of countervailing 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 buyer nor the one that is optimal for the suppliers (or a supplier) is efficient in general, which opens the scope for increasing social surplus by making bargaining powers more equal. Social-surplus-increasing countervailing power and socially harmful vertical integration arise naturally in this 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. We show that the effects of outside options can differ relative to complete information setups, and we show that bargaining externalities arise naturally.

Our paper shows that an economic agent's strength or weakness has two dimensions that are, conceptually, independent. The first one, which may be thought of as the agent's productive strength or power, refers to 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 productive power 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 countervailing 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 welfare-increasing countervailing power.

Avenues for future research are many. For example, one could augment the setup to have multiple buyers and multiple suppliers, which may give rise to a raising rivals' costs effect of vertical integration. (The key issue to be resolved for such an extension is how to maintain the assumption of one-dimensional private information after integration.) More fundamentally, developing a better understanding of what determines bargaining power would add considerable value. The distinction between productive strength and bargaining power brought to light in the present paper may prove useful in that regard.

References

Allain, 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*.

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.

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

Carlton, D. W. and M. Israel (2011): “Proper Treatment of Buyer Power in Merger Review,” *Review of Industrial Organization*, 39, 127–136.

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.

Che, Y.-K., E. Iossa, and P. Rey (2017): “Prizes versus Contracts as Incentives for Innovation,” TSE Working Paper.

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

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.

Decarolis, F. and G. Rovigatti (2020): "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising," Working Paper, Bocconi University.

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.

Galbraith, J. K. (1952): *American Capitalism: The Concept of Countervailing Power*, Houghton Mifflin, Boston.

——— (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.

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.

Kreps, D. M. and J. A. Scheinkman (1983): "Quantity Precommitment and Bertrand competition yield Cournot outcomes," *Bell Journal of Economics*, 14, 326–337.

Krishna, V. (2002): *Auction Theory*, Elsevier Science, Academic Press.

Larsen, B. J. (2020): "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," *Review of Economic Studies*.

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.

LOERTSCHER, S. AND L. M. MARX (2019): "Merger Review for Markets with Buyer Power," *Journal of Political Economy*, 127.

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

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

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

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.

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. 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.

RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

SALANT, S. W., S. SWITZER, AND R. J. REYNOLDS (1983): "Optimal Auctions with Asymmetric Financial ExternalitiesLosses 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.

SAMUELSON, P. A. (1949): "International Factor-Price Equalisation Once Again," *Economic Jour-*

nal, 59, 181–197.

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

SHAPLEY, L. AND M. SHUBIK (1972): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.

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.

STEPTOE, M. L. (1993): "The Power-Buyer Defense in Merger Cases," *Antitrust Law Journal*, 61, 493–504.

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

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.

ZHANG, X., P. MANCHANDA, AND J. CHU (2019): "'Meet Me Halfway': The Value of Bargaining," Working Paper, London School of Economics.

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 (Appendix A.2).

A.1 Concepts and derivations

For ease of exposition, in this appendix we assume that $n = 1$. The extension to $n > 1$ is straightforward.

Take as given a direct mechanism $\langle Q, M_B, M_S \rangle$, where $Q : [\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}] \to [0, 1]$ and $M_B, M_S : [\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}] \to \mathbb{R}$. Given reports v and c, $Q(v, c) \in [0, 1]$ is the probability with which the supplier trades with the buyer, $M_B(v, c)$ is the payment from the buyer to the mechanism, and $M_S(v, c)$ is the payment from the mechanism to the supplier. By the Revelation Principle, the focus on direct mechanisms is without loss of generality.

Let $\hat{q}_B(z)$ be the buyer's expected quantity if it reports z and the supplier reports truthfully, and let $\hat{m}_B(z)$ be the buyer's expected payment if it reports z and the supplier reports truthfully:

$$\hat{q}_B(z) = \mathbb{E}_c[Q(z, c)] \quad \text{and} \quad \hat{m}_B(z) = \mathbb{E}_c[M_B(z, c)].$$

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

Key constraints

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

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

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

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

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

$$\mathbb{E}_{v,c}\left[M_B(v,c) - M_S(v,c)\right] \geq 0.$$

Interim expected payoffs

Standard arguments (see, e.g., Krishna, 2002, Chapter 5.1) proceed as follows:

Focusing on the buyers, incentive compatibility implies that

$$\hat{u}_B(v) = \max_{z \in [\underline{v}, \overline{v}]} \left\{\hat{q}_B(z)v - \hat{m}_B(z)\right\},$$

i.e., $\hat{u}_B$ is a maximum of a family of affine functions, which implies that $\hat{u}_B$ is convex and so absolutely continuous and differentiable almost everywhere in the interior of its domain.[53] In addition, incentive compatibility implies that $\hat{u}_B(z) \geq \hat{q}_B(v)z - \hat{m}_B(v) = \hat{u}_B(v) + \hat{q}_B(v)(z-v)$, which for $\varepsilon > 0$ implies

$$\frac{\hat{u}_B(v+\varepsilon) - \hat{u}_B(v)}{\varepsilon} \geq \hat{q}_B(v)$$

and for $\varepsilon < 0$ implies

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

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

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

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

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

[53]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 $\hat{u}_B$ and $\hat{u}_S$ in (14) and (15), we can rewrite these as

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

and

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

The expected payment by the buyer is then

$$\begin{aligned}
\mathbb{E}_v\left[\hat{m}_B(v)\right] &= \int_{\underline{v}}^{\overline{v}} \hat{m}_B(v)f(v)dv \\
&= \int_{\underline{v}}^{\overline{v}} \left(\hat{q}_B(v)v - \int_{\underline{v}}^{v} \hat{q}_B(t)dt\right) f(v)dv - \hat{u}_B(\underline{v}) \\
&= \left(\int_{\underline{v}}^{\overline{v}} \hat{q}_B(v)vf(v)dv - \int_{\underline{v}}^{\overline{v}} \int_{t}^{\overline{v}} \hat{q}_B(t)f(v)dvdt\right) - \hat{u}_B(\underline{v}) \\
&= \left(\int_{\underline{v}}^{\overline{v}} \hat{q}_B(v)vf(v)dv - \int_{\underline{v}}^{\overline{v}} \hat{q}_B(t)\left(1 - F(t)\right)dt\right) - \hat{u}_B(\underline{v}) \\
&= \int_{\underline{v}}^{\overline{v}} \hat{q}_B(v)\left(v - \frac{1-F(v)}{f(v)}\right) f(v)dv - \hat{u}_B(\underline{v}) \\
&= \int_{\underline{v}}^{\overline{v}} \hat{q}_B(v)\Phi(v)f(v)dv - \hat{u}_B(\underline{v}) \\
&= \mathbb{E}_v\left[\hat{q}_B(v)\Phi(v)\right] - \hat{u}_B(\underline{v}),
\end{aligned}$$

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, and the last equality uses the definition of the expectation. Similarly, using (17), the expected payment to the supplier is

$$\mathbb{E}_c\left[\hat{m}_S(c)\right] = \int_{\underline{c}}^{\overline{c}} \hat{m}_S(c)g(c)dc = \mathbb{E}_c\left[\hat{q}_S(c)\Gamma(c)\right] + \hat{u}_S(\overline{c}).$$

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

$$\mathbb{E}_{v,c}\left[\left(\Phi(v) - \Gamma(c)\right)Q(v,c)\right] - \hat{u}_B(\underline{v}) - \hat{u}_S(\overline{c}).$$

As mentioned, it is straightforward to extend these results to the case of $n > 1$.

A.2 Myerson-Satterthwaite 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 $\underline{v} < \bar{c}$, 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 $\underline{v} < \bar{c}$ 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 $p = (\underline{v} + \bar{c})/2$ to the supplier achieves this.

By now, the proof of this result can be provided in a couple of lines (see, e.g, Krishna, 2002). Consider the dominant strategy implementation in which the buyer pays $p_B = \max\{c, \underline{v}\}$ and the supplier receives $p_S = \min\{v, \bar{c}\}$ whenever there is trade, and no payments are made otherwise. Notice that $\hat{u}_B(\underline{v}) = 0 = \hat{u}_S(\bar{c})$. 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 $\underline{v} < \bar{c}$, assume to the contrary that $\underline{v} \geq \bar{c}$. Then the mechanism described above continues to satisfy incentive compatibility and individual rationality, but for all type realizations $p_B = \underline{v} \geq \bar{c} = p_S$, which implies that the mechanism does not run a deficit.

B Appendix: Proofs

Proof of Lemma 1. As mentioned in the text, $\mathbf{Q}_\rho$ maximizes (3) pointwise given ρ. Given that $\rho > 0$, we can write (3) as

$$\begin{aligned}
&\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}\left[w_B(v-\Phi(v))+w_i(\Gamma_i(c_i)-c_i)+\rho\left(\Phi(v)-\Gamma_i(c_i)\right)\right]Q_i(v,\boldsymbol{c})\right]\\
= \;&\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}\left[w_Bv+(\rho-w_B)\Phi(v)-w_ic_i-(\rho-w_i)\Gamma_i(c_i)\right]Q_i(v,\boldsymbol{c})\right]\\
= \;&\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}\left[\rho v-(\rho-w_B)\frac{1-F(v)}{f(v)}-\rho c_i-(\rho-w_i)\frac{G_i(c_i)}{g_i(c_i)}\right]Q_i(v,\boldsymbol{c})\right]\\
= \;&\rho\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}\left[v-\frac{\rho-w_B}{\rho}\frac{1-F(v)}{f(v)}-c_i-\frac{\rho-w_i}{\rho}\frac{G_i(c_i)}{g_i(c_i)}\right]Q_i(v,\boldsymbol{c})\right]\\
= \;&\rho\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}\left[\Phi_{w_B/\rho}(v)-\Gamma_{i,w_i/\rho}(c_i)\right]Q_i(v,\boldsymbol{c})\right],
\end{aligned}$$

which is maximized pointwise by having $Q_{i,\rho}(v,\mathbf{c}) = 1$ if and only if $\Phi_{w_B/\rho}(v) \geq \Gamma_{i,w_i/\rho}(c_i) = \min_{j\in\mathcal{N}}\Gamma_{j,w_j/\rho}(c_j)$, and $Q_{i,\rho}(v,\mathbf{c}) = 0$ otherwise. The result then follows by making the substitution $\beta^{\mathbf{w}} = 1/\rho$. ■

Proof of Proposition 2. Let $\langle\mathbf{Q}^{\mathbf{w}},\mathbf{M}^{\mathbf{w},\boldsymbol{\eta}}\rangle \in \mathcal{M}$ denote the incomplete information bargaining mechanism with weights $\mathbf{w}$ and shares $\boldsymbol{\eta}$, with associated expected payoff vector $\mathbf{u}(\mathbf{w},\boldsymbol{\eta})$. Thus, the maximized value of expected weighted welfare over all mechanisms in $\mathcal{M}$ is $\sum_{i\in\{B\}\cup\mathcal{N}} w_iu_i(\mathbf{w},\boldsymbol{\eta})$.

We first show that the expected payoffs from $\langle\mathbf{Q}^{\mathbf{w}},\mathbf{M}^{\mathbf{w},\boldsymbol{\eta}}\rangle$ are Pareto undominated among the expected payoffs for any mechanism in $\mathcal{M}$. Proceeding by contradiction, suppose that $\mathbf{u}(\mathbf{w},\boldsymbol{\eta})$ is Pareto dominated by expected payoff vector $\tilde{\mathbf{u}}$ associated with a mechanism $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle \in \mathcal{M}$, i.e., $\tilde{u}_i \geq u_i(\mathbf{w},\boldsymbol{\eta})$ for all $i \in \{B\}\cup\mathcal{N}$ and there exists $\ell \in \{B\}\cup\mathcal{N}$ such that $\tilde{u}_\ell > u_\ell(\mathbf{w},\boldsymbol{\eta})$. If there exists $i \in \{B\}\cup\mathcal{N}$ such that $\tilde{u}_i > u_i(\mathbf{w},\boldsymbol{\eta})$ and $w_i > 0$, then $\sum_{i\in\{B\}\cup\mathcal{N}} w_i\tilde{u}_i > \sum_{i\in\{B\}\cup\mathcal{N}} w_iu_i(\mathbf{w},\boldsymbol{\eta})$, implying that $\langle\mathbf{Q}^{\mathbf{w}},\mathbf{M}^{\mathbf{w},\boldsymbol{\eta}}\rangle$ does not maximize expected weighted welfare over mechanisms in $\mathcal{M}$, which is a contradiction. So, for all $i \in \{B\}\cup\mathcal{N}$ such that $\tilde{u}_i > u_i(\mathbf{w},\boldsymbol{\eta})$, we have $w_i = 0$. It follows that $\sum_{i\in\{B\}\cup\mathcal{N}} w_i\tilde{u}_i = \sum_{i\in\{B\}\cup\mathcal{N}} w_iu_i(\mathbf{w},\boldsymbol{\eta})$, which says that $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle$ maximizes expected weighted welfare. By the uniqueness of the allocation rule identified in Lemma 1, we have $\tilde{\mathbf{Q}} = \mathbf{Q}^{\mathbf{w}}$. Thus, using the payoff equivalence theorem, the difference $\tilde{u}_\ell - u_\ell(\mathbf{w},\boldsymbol{\eta})$ reflects an increase in the interim expected payment to agent ℓ's worst-off type. It follows that there exists a mechanism in $\mathcal{M}$

that has allocation rule $\mathbf{Q}^{\mathbf{w}}$ and payment rule based on $\tilde{\mathbf{M}}$, but that redirects agent ℓ's fixed payment to an agent with positive bargaining weight, that has greater expected weighted welfare than $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$, which is a contradiction. This concludes the first part of the proof.

We now turn to the proof that any Pareto undominated payoff vector can be achieved using $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$ with appropriately chosen $\mathbf{w}$ and $\boldsymbol{\eta}$. Let $\tilde{\mathbf{u}}$ be a Pareto undominated payoff profile associated with $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle \in \mathcal{M}$. By the assumption of Pareto undominatedness, $\tilde{u}_B$ solves

$$\max_{\langle \mathbf{Q}, \mathbf{M} \rangle \in \mathcal{M}} u_B \text{ s.t. } u_1 \geq \tilde{u}_1, \ldots, u_n \geq \tilde{u}_n.$$

Using incentive compatibility and individual rationality, the above problem can be recast as choosing $\mathbf{Q}$, $\hat{u}_B(\underline{v}) \geq 0$, and $\hat{u}_i(\overline{c}) \geq 0$ for all $i \in \mathcal{N}$, subject to incentive compatibility, no deficit, and the suppliers' payoff constraints, which has associated Lagrangian

$$\begin{aligned} \mathcal{L} \;=\; & \mathbb{E}_{v,\mathbf{c}} \left[\sum_{i \in \mathcal{N}} \left[v - \Phi(v) + \rho(\Phi(v) - \Gamma_i(c_i)) + \mu_i(\Gamma_i(c_i) - c_i) \right] Q_i(v, \mathbf{c}) \right] \\ & + (1-\rho)\hat{u}_B(\underline{v}) + \sum_{i \in \mathcal{N}} (\mu_i - \rho)\hat{u}_i(\overline{c}) - \sum_{i \in \mathcal{N}} \mu_i \tilde{u}_i, \end{aligned}$$

where $\rho \geq 0$ is the multiplier on the no-deficit constraint and $\mu_i \geq 0$ is the multiplier on the constraint that $u_i \geq \tilde{u}_i$. If $\rho < 1$, then we maximize $\mathcal{L}$ by increasing $\hat{u}_B(\underline{v})$ unboundedly, in violation of the no-deficit constraint, so $\rho \geq 1$.

Maximizing $\mathcal{L}$ pointwise, $\tilde{Q}_i$ is defined by

$$\tilde{Q}_i(v, \mathbf{c}) = \begin{cases} 1 & \text{if } \Phi_{\frac{1}{\rho}}(v) \geq \Gamma_{i,\frac{\mu_i}{\rho}}(c_i) = \min_{j \in \mathcal{N}} \Gamma_{j,\frac{\mu_j}{\rho}}(c_j), \\ 0 & \text{otherwise,} \end{cases}$$

where ρ is the smallest value greater than or equal to 1 such that the no-deficit constraint is satisfied under allocation rule $\tilde{Q}_i(v, \mathbf{c})$. It then follows that

$$\hat{u}_i(\overline{c}) = \tilde{u}_i - \mathbb{E}_{v,\mathbf{c}} \left[(\Gamma_i(c_i) - c_i)\tilde{Q}_i(v, \mathbf{c}) \right]$$

and

$$\hat{u}_B(\underline{v}) = \mathbb{E}_{v,\mathbf{c}} \left[\sum_{i \in \mathcal{N}} (\Phi(v) - \Gamma_i(c_i))\tilde{Q}_i(v, \mathbf{c}) \right] - \sum_{i \in \mathcal{N}} \hat{u}_i(\overline{c}).$$

Case 1: Suppose that $\mathbb{E}_{v,\mathbf{c}} \left[\sum_{i \in \mathcal{N}} (v - \Phi(v)) \tilde{Q}_i(v, \mathbf{c}) \right] < \tilde{u}_B$, which implies that $\hat{u}_B(\underline{v}) > 0$ and $\rho = 1$. If for any $i \in \mathcal{N}$ we have $\mu_i > 1$, then $\mu_i - \rho > 1 - \rho$ and $\mathcal{L}$ is maximized by decreasing $\hat{u}_B(\underline{v})$ to zero and increasing $\hat{u}_i(\overline{c})$, contradicting $\hat{u}_B(\underline{v}) > 0$, so we have $\mu_i \leq 1$. Further, if $\hat{u}_i(\overline{c}) > 0$, then it must be that $\mu_i = \rho = 1$. So, letting $\mathbf{w} = (1, \mu_1, \ldots, \mu_n)$, all agents whose worst-off types have a positive interim expected payoff have the maximum bargaining weight of 1. Thus, the payoffs $\tilde{\mathbf{u}}$ are generated by $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$ with $\mathbf{w} = (1, \mu_1, \ldots, \mu_n)$ and for $i \in \mathcal{N}$,

$$\eta_i = \begin{cases} \frac{\tilde{u}_i - \mathbb{E}_{v,\mathbf{c}}\left[(\Gamma_i(c_i)-c_i)\tilde{Q}_i(v,\mathbf{c})\right]}{\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}(\Phi(v)-\Gamma_i(c_i))\tilde{Q}_i(v,\mathbf{c})\right]} & \text{if } \mu_i = 1, \\ 0 & \text{otherwise,} \end{cases} \tag{18}$$

and $\eta_B = 1 - \sum_{i\in\mathcal{N}} \eta_i$.

Case 2: Suppose that $\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}} (v - \Phi(v))\, \tilde{Q}_i(v,\mathbf{c})\right] = \tilde{u}_B$ and that for some $j \in \mathcal{N}$, $\mathbb{E}_{v,\mathbf{c}}\left[(\Gamma_j(c_j) - c_j)\, \tilde{Q}_j(v,\mathbf{c})\right] < \tilde{u}_j$. Then we can repeat the analysis in Case 1, but centered on supplier j rather than on the buyer, noting that Pareto undominatedness implies that $\tilde{u}_j$ solves

$$\max_{\langle \mathbf{Q},\mathbf{M}\rangle\in\mathcal{M}} u_j \text{ s.t. } u_B \geq \tilde{u}_B \text{ and for } i \in N\backslash\{j\},\ \tilde{u}_i \geq \tilde{u}_i.$$

Then we obtain the analogous result that the payoffs $\tilde{\mathbf{u}}$ are replicated by $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$ with $\mathbf{w} = (\mu_B, \mu_1, \ldots, \mu_{j-1}, 1, \mu_{j+1}, \ldots, \mu_n)$ and

$$\eta_B = \begin{cases} \frac{\tilde{u}_B - \mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}(v-\Phi(v))\tilde{Q}_i(v,\mathbf{c})\right]}{\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}}(\Phi(v)-\Gamma_i(c_i))\tilde{Q}_i(v,\mathbf{c})\right]} & \text{if } \mu_B = 1, \\ 0 & \text{otherwise,} \end{cases}$$

for $i \in \mathcal{N}\backslash\{j\}$, η_i as given in (18), and $\eta_j = 1 - \eta_B - \sum_{i\in\mathcal{N}\backslash\{j\}} \eta_i$.

Case 3: Suppose that $\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}} (v - \Phi(v))\, \tilde{Q}_i(v,\mathbf{c})\right] = \tilde{u}_B$ and that for all $i \in \mathcal{N}$, $\mathbb{E}_{v,\mathbf{c}}\left[(\Gamma_i(c_i) - c_i)\, \tilde{Q}_i(v,\mathbf{c})\right] = \tilde{u}_i$. Then $\hat{u}_B(\underline{v}) = 0$ and for all $i \in \mathcal{N}$, $\hat{u}_i(\overline{c}) = 0$. It follows that $\rho \geq 1$ and that for all $i \in \mathcal{N}$, $\mu_i \leq \rho$.

Case 3a: Suppose that for all $i \in \mathcal{N}$, μ_i is finite. Letting $\overline{w} \equiv \max\{1, \mu_1, \ldots, \mu_n\}$, $\tilde{Q}_i$ is the same allocation rule as $Q_i^{\mathbf{w}}$, where $\mathbf{w} = (\frac{1}{\overline{w}}, \frac{\mu_1}{\overline{w}}, \ldots, \frac{\mu_n}{\overline{w}})$, and the payoffs $\tilde{\mathbf{u}}$ are replicated by $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$ with any specification of $\boldsymbol{\eta}$.

Case 3b: Suppose that for some nonempty $\hat{\mathcal{N}} \subseteq \mathcal{N}$, for all $i \in \hat{\mathcal{N}}$, μ_i is infinite, which implies that ρ is infinite. It follows that the allocation rule $\tilde{Q}_i(v,\mathbf{c})$ maximizes $\mathbb{E}_{v,\mathbf{c}}[\sum_{i\in\hat{\mathcal{N}}}(\Phi(v) - c_i)Q_i(v,\mathbf{c})]$. In this case, $\tilde{Q}_i$ is the same allocation rule as $Q_i^{\mathbf{w}}$, where $w_B = 0$, for all $i \in \hat{\mathcal{N}}$, $w_i = 1$, and for all $i \in \mathcal{N}\backslash\hat{\mathcal{N}}$, $w_i = 0$, and the payoffs $\tilde{\mathbf{u}}$ are replicated by $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$ with any specification of $\boldsymbol{\eta}$. ■

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), (ii), and (iii) 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 (5) fails, $w_1 = \cdots = w_n = w$, and $G_1 = \cdots = G_n$. Then the planner and market both rank the suppliers the same, but they evaluate either the virtual values or the virtual costs (or both) using different weights because either $w_B\beta^{\mathbf{w}} \neq \beta^{\mathbf{1}}$ or $w_i\beta^{\mathbf{w}} \neq \beta^{\mathbf{1}}$ or both. It follows that $\mathbf{Q}^{\mathbf{w}}(v, \mathbf{c}) \neq \mathbf{Q}^*(v, \mathbf{c})$ for all $(v, \mathbf{c})$ in an open subset of $[\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]^n$.

Case 2. Suppose that (5) holds, $w_1 \neq w_2$ so that (ii) fails, and $G_1 = \cdots = G_n$. Then the planner and market rank the suppliers 1 and 2 differently for (c_1, c_2) in an open subset of $[\underline{c}, \overline{c}]^2$, implying that $\mathbf{Q}^{\mathbf{w}}(v, \mathbf{c}) \neq \mathbf{Q}^*(v, \mathbf{c})$ for all $(v, \mathbf{c})$ in an open subset of $[\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]^n$.

Case 3. Suppose that (5) holds, $w_1 = \cdots = w_n = w$, $w_B > w$, and $G_1 \neq G_2$, so that (iii) fails. Then the planner and market rank the suppliers 1 and 2 differently for (c_1, c_2) in an open subset of $[\underline{c}, \overline{c}]^2$, implying that $\mathbf{Q}^{\mathbf{w}}(v, \mathbf{c}) \neq \mathbf{Q}^*(v, \mathbf{c})$ for all $(v, \mathbf{c})$ in an open subset of $[\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]^n$. ■

Proof of Proposition 4. To begin, note that under the assumption that suppliers have symmetric bargaining weights, the allocation rule $\mathbf{Q}^{\mathbf{w}}$ depends only on the bargaining differential $\Delta \equiv \frac{w_B - w_1}{\max\{w_B, w_1\}}$. Given Δ, we denote the associated point on the frontier as $(u_S(\Delta), u_B(\Delta))$, where u_S is the sum of all suppliers' expected payoffs.

We first show that ω is strictly decreasing. If it is not strictly decreasing, then there are two points on the frontier, indexed by Δ and Δ', where expected social surplus is strictly greater at the point indexed by Δ' and expected surplus for both the buyer and the suppliers is weakly greater at the point indexed by Δ'. But then weighted welfare must not have been maximized at the point indexed by Δ because total surplus could be increased while still satisfying all the constraints and some of that additional surplus could be allocated to one or more of the agents with a positive bargaining weight (e.g., the buyer if $\Delta \geq 0$ and otherwise a supplier). This completes the proof that ω is strictly decreasing.

Now turn to the question of concavity. As illustrated in the figure below, suppose that the Williams frontier is not concave.

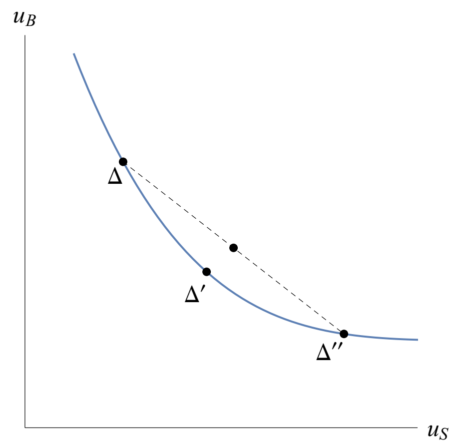

Then there exist points on the frontier, which we denote by their associated bargaining differentials Δ, Δ', and Δ'', with $\Delta > \Delta' > \Delta''$, such that $(u_S(\Delta) + u_S(\Delta''))/2 > u_S(\Delta')$ and $(u_B(\Delta) + u_B(\Delta''))/2 > u_B(\Delta')$. Let $\mu(\Delta)$ denote the incomplete information bargaining mechanism for bargaining differential Δ. Let w'_B and w' be bargaining weights consistent with Δ'. Expected weighted welfare with weights $\mathbf{w}'$ under mechanism $\mu(\Delta')$ satisfies

$$\begin{aligned} w'_B u_B(\Delta') + \sum_{i \in \mathcal{N}} w' u_i(\Delta') &= w'_B u_B(\Delta') + w' u_S(\Delta') \\ &< w'_B \frac{u_B(\Delta) + u_B(\Delta'')}{2} + w' \frac{u_S(\Delta) + u_S(\Delta'')}{2} \\ &= w'_B \frac{u_B(\Delta) + u_B(\Delta'')}{2} + \sum_{i \in \mathcal{N}} w' \frac{u_i(\Delta) + u_i(\Delta'')}{2}, \end{aligned}$$

where the right side is expected weighted welfare with weights $\mathbf{w}'$ under the mechanism that is a 50-50 mixture of $\mu(\Delta)$ and $\mu(\Delta'')$, which since the no-deficit condition is satisfied for this mixture mechanism, contradicts the assumption that $\mu(\Delta')$ is the incomplete information bargaining mechanism with weights $\mathbf{w}'$, thereby completing the proof. ■

Proof of Proposition 7. With nonoverlapping supports and symmetric bargaining weights, the pre-integration market achieves the first-best. After integration between the buyer and supplier i, the buyer's willingness to pay is the cost realization of the integrated supplier, that is, c_i, whose support is $[\underline{c}, \overline{c}]$. Thus, we have a generalized Myerson-Satterthwaite problem (generalized insofar as there is one buyer but $n - 1 \geq 1$ suppliers). For this setting, impossibility of first-best trade obtains (see, e.g., Delacrétaz et al., 2019), regardless of bargaining weights. ■

Proof of Proposition 8. We have proved the first part in the text and are thus left to prove the second part.

We begin with a brief preamble to establish some notation and useful relations. Let $\hat{u}_{B,\mathbf{Q}}(v;\mathbf{e}_{-B})$ denote the interim expected payoff of a buyer with type v, not including the (constant) interim expected payment to the worst-off type, when the allocation rule is $\mathbf{Q}$ and suppliers have investments $\mathbf{e}_{-B}$. For supplier $i \in \mathcal{N}$, define $\hat{u}_{i,\mathbf{Q}}(c_i;\mathbf{e}_{-i})$ analogously. For $i \in \{B\} \cup \mathcal{N}$, let $u_{i,\mathbf{Q}}(\mathbf{e})$ denote the expected payoff of agent i when the allocation rule is $\mathbf{Q}$ and investments are $\mathbf{e}$. For any allocation rule $\mathbf{Q}$, let $q_B(v;\mathbf{e}_{-B}) \equiv \mathbb{E}_{\mathbf{c}|\mathbf{e}_{-B}}\left[Q_B(v,\mathbf{c})\right]$, and for all $i \in \mathcal{N}$, let $q_i(c_i;\mathbf{e}_{-i}) \equiv \mathbb{E}_{v,\mathbf{c}_{-i}|\mathbf{e}_{-i}}\left[Q_i(v,c_i,\mathbf{c}_{-i})\right]$. As discussed in Appendix A.1, by the payoff equivalence theorem, we have, up to a constant,

$$\hat{u}_{B,\mathbf{Q}}(v;\mathbf{e}_{-B}) = \int_{\underline{v}}^{v} q_B(x;\mathbf{e}_{-B})dx, \tag{19}$$

and, taking expectations with respect to v, one obtains

$$u_{B,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{v}}^{\overline{v}} q_B(x;\mathbf{e}_{-B})(1-F(x;e_B))dx \tag{20}$$

up to a constant, and, analogously, for all $i \in \mathcal{N}$,

$$u_{i,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{c}}^{\overline{c}} q_i(x;\mathbf{e}_{-i})G_i(x;e_i)dx \tag{21}$$

up to a constant.

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

$$\overline{\mathbf{e}} \in \arg\max_{\mathbf{e}} \sum_{i\in\{B\}\cup\mathcal{N}} u_{i,\mathbf{Q}^{FB}}(\mathbf{e}) - \sum_{i\in\{B\}\cup\mathcal{N}} \Psi_i(e_i).$$

which implies that for all $i \in \{B\} \cup \mathcal{N}$,

$$\overline{e}_i \in \arg\max_{e_i} u_{i,\mathbf{Q}^{FB}}(e_i,\overline{\mathbf{e}}_{-i}) - \Psi_i(e_i). \tag{22}$$

Assume that (9) and (10) hold and that either supports are overlapping or (11) holds. 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 $\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}$ denote the associated budget parameter. Suppose that first-best investments $\overline{\mathbf{e}}$ are Nash equilibrium investments, which implies that for all $i \in \{B\} \cup \mathcal{N}$,

$$\overline{e}_i \in \arg\max_{e_i} u_{i,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}(e_i,\overline{\mathbf{e}}_{-i}) - \Psi_i(e_i). \tag{23}$$

Assumptions (9) and (10) ensure that the first-best investments in (22) and (23) are characterized by their first-order conditions. Thus, using (20) and (22), we have

$$-\int_{\underline{v}}^{\overline{v}} q_B^{FB}(x;\overline{\mathbf{e}}_{-B})\frac{\partial F(x;\overline{e}_B)}{\partial e}dx - \Psi_B'(\overline{e}_B) = 0. \tag{24}$$

Similarly, using (20) and (23), we have

$$-\int_{\underline{v}}^{\overline{v}} q_B^{\mathbf{w},\overline{\mathbf{e}}}(x;\overline{\mathbf{e}}_{-B})\frac{\partial F(x;\overline{e}_B)}{\partial e}dx - \Psi_B'(\overline{e}_B) = 0. \tag{25}$$

Combining (24) and (25), we have

$$\int_{\underline{v}}^{\overline{v}} (q_B^{FB}(x;\overline{\mathbf{e}}_{-B}) - q_B^{\mathbf{w},\overline{\mathbf{e}}}(x;\overline{\mathbf{e}}_{-B}))\frac{\partial F(x;\overline{e}_B)}{\partial e}dx = 0. \tag{26}$$

Writing this in terms of the ex post allocation rules, we have

$$\mathbb{E}_{\mathbf{c}|\overline{\mathbf{e}}_{-B}}\left[\int_{\underline{v}}^{\overline{v}} (Q_B^{FB}(x,\mathbf{c}) - Q_B^{\mathbf{w},\overline{\mathbf{e}}}(x,\mathbf{c}))\frac{\partial F(x;\overline{e}_B)}{\partial e}dx\right] = 0. \tag{27}$$

Steps analogous to those leading to (26) imply that for all $i \in \mathcal{N}$,

$$\int_{\underline{c}}^{\overline{c}} (q_i^{FB}(x;\overline{\mathbf{e}}_{-i}) - q_i^{\mathbf{w},\overline{\mathbf{e}}}(x;\overline{\mathbf{e}}_{-i}))\frac{\partial G_i(x;\overline{e}_i)}{\partial e}dx = 0. \tag{28}$$

By Lemma 1, we know that $Q_B^{\mathbf{w},\overline{\mathbf{e}}}(v,\mathbf{c}) = 1$ implies that $\min_{i\in\mathcal{N}} c_i \leq \min_{i\in\mathcal{N}} \Gamma_{i,w_i\beta_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c_i;\overline{e}_i) \leq \Phi_{w_B\beta_{\overline{\mathbf{e}}}^{\mathbf{w}}}(v;\overline{e}_B) \leq v$. Thus, $Q_B^{\mathbf{w},\overline{\mathbf{e}}}(v,\mathbf{c}) \leq Q_B^{FB}(v,\mathbf{c})$ for all $(v,\mathbf{c})$. Because we assume that $\frac{\partial F(v;e)}{\partial e} < 0$ for all $v \in (\underline{v},\overline{v})$, (27) then implies that

$$Q_B^{\mathbf{w},\overline{\mathbf{e}}}(v,\mathbf{c}) = Q_B^{FB}(v,\mathbf{c}) \tag{29}$$

for all but a zero-measure set of types. Condition (29) implies that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ is such that $Q_B^{\mathbf{w},\overline{\mathbf{e}}} = Q_B^{FB}$, that is, $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ and $\mathbf{Q}^{FB}$ induce trade in the same instances. It remains to show that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ always induces the same supplier to produce as does $\mathbf{Q}^{FB}$.

We begin by considering the case with overlapping supports and then consider the case in which (11) 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 suppliers based on virtual types for an open set of types. That is, suppose that there exist agents, 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 supplier 1 trades in the first-best while supplier 2 trades under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$: $Q_1^{FB}(\hat{v},\hat{\mathbf{c}})=1$ and $Q_2^{\mathbf{w},\overline{\mathbf{e}}}(\hat{v},\hat{\mathbf{c}})=1$. This implies that

$$\hat{c}_1 < \hat{c}_2 \leq \Gamma_{2,w_2\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\hat{c}_2;\overline{e}_2) \leq \Gamma_{1,w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\hat{c}_1;\overline{e}_1) \text{ and } \Gamma_{2,w_2\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\hat{c}_2;\overline{e}_2) \leq \Phi_{w_B\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(v;\overline{e}_B)$$

and that suppliers other than 1 and 2 have types greater than or equal to $\hat{c}_1$ and virtual types greater than or equal to $\Gamma_{2,w_2\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\hat{c}_2;\overline{e}_2)$. Because $\hat{c}_1 < \Gamma_{1,w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\hat{c}_1;\overline{e}_1)$, it follows that $w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}} < 1$ and so for all $c \in (\underline{c},\overline{c})$, $c < \Gamma_{1,w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(c;\overline{e}_1)$. Thus, letting $\tilde{c}_1 \in (\max\{\underline{c},\underline{v}\},\overline{c})$, $\tilde{v} \in (\tilde{c}_1,\min\{\overline{c},\Gamma_{1,w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\tilde{c}_1;\overline{e}_1)\})$, and for all $i \in \mathcal{N}\backslash\{1\}$, $\tilde{c}_i=\overline{c}$, we have

$$\tilde{c}_1 < \tilde{v} < \Gamma_{1,w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\tilde{c}_1;\overline{e}_1) \text{ and } \tilde{v} < \min_{i\in\mathcal{N}\backslash\{1\}} \tilde{c}_i,$$

which implies that

$$\Phi_{w_B\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\tilde{v};\overline{e}_B) < \min_{i\in\mathcal{N}} \Gamma_{i,w_i\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(\tilde{c}_i;\overline{e}_i),$$

and so $Q_B^{FB}(\tilde{v},\tilde{\mathbf{c}})=1$ and $Q_B^{\mathbf{w},\overline{\mathbf{e}}}(\tilde{v},\tilde{\mathbf{c}})=0$. By continuity, for all $(v,\mathbf{c})$ in an open set of types around $(\tilde{v},\tilde{\mathbf{c}})$, $Q_B^{\mathbf{w},\overline{\mathbf{e}}}(v,\mathbf{c}) \neq Q_B^{FB}(v,\mathbf{c})$, which contradicts (29). Thus, we conclude that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ does not discriminate among suppliers based on virtual types and so $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ induces the same supplier to produce as does $\mathbf{Q}^{FB}$.

Case 2: (11) holds and $\underline{v} \geq \overline{c}$. If $w_1=\cdots=w_n$, then (11) and (29) imply that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ induces the same supplier to produce as does $\mathbf{Q}^{FB}$, and we are done. So, let $w_1=\min_{i\in\mathcal{N}} w_i < w_2$. Using (11), for all $c \in (\underline{c},\overline{c})$,

$$\Gamma_{2,w_2\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(c;\overline{e}_2) < \max_{i\in\mathcal{N}} \Gamma_{i,w_i\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(c;\overline{e}_i) \leq \Gamma_{1,w_1\beta^{\mathbf{w}}_{\overline{\mathbf{e}}}}(c;\overline{e}_1), \tag{30}$$

and, dropping investment as an argument in the suppliers' weighted virtual cost functions, for all $c \in (\underline{c},\overline{c})$,

$$\begin{aligned}
q_1^{\mathbf{w},\overline{\mathbf{e}}}(c;\overline{\mathbf{e}}_{-1}) &= \Pr_{v,\mathbf{c}_{-1}|\overline{\mathbf{e}}_{-1}}\left(\Gamma_{1,w_1\beta^{\mathbf{w}}}(c) \leq \min_{i\in\mathcal{N}\backslash\{1\}}\{\Phi_{w_B\beta^{\mathbf{w}}}(v;\overline{e}_B),\Gamma_{i,w_i\beta^{\mathbf{w}}}(c_i)\}\right)\\
&\leq \Pr_{\mathbf{c}_{-1}|\overline{\mathbf{e}}_{-1}}\left(\Gamma_{1,w_1\beta^{\mathbf{w}}}(c) \leq \min_{i\in\mathcal{N}\backslash\{1\}}\Gamma_{i,w_i\beta^{\mathbf{w}}}(c_i)\right)\\
&= \Pr_{\mathbf{c}_{-1}|\overline{\mathbf{e}}_{-1}}\left(c \leq \Gamma^{-1}_{1,w_1\beta^{\mathbf{w}}}(\min_{i\in\mathcal{N}\backslash\{1\}}\Gamma_{i,w_i\beta^{\mathbf{w}}}(c_i))\right)\\
&< \Pr_{\mathbf{c}_{-1}|\overline{\mathbf{e}}_{-1}}\left(c \leq \min_{i\in\mathcal{N}\backslash\{1\}} c_i\right)\\
&= q_1^{FB}(c;\overline{\mathbf{e}}_{-1}),
\end{aligned}$$

where the first inequality follows from the monotonicity of the min operator, the second inequality follows from (30), and the final equality uses $\underline{v} \geq \bar{c}$. Thus, for all $c \in (\underline{c}, \bar{c})$,

$$q_1^{\mathbf{w},\bar{\mathbf{e}}}(c; \bar{\mathbf{e}}_{-1}) < q_1^{FB}(c; \bar{\mathbf{e}}_{-1}), \tag{31}$$

implying that (28) is violated for $i = 1$. This proves that discrimination is not compatible with (28). Hence, $\mathbf{Q}^{\mathbf{w},\bar{\mathbf{e}}} = \mathbf{Q}^{FB}$ follows, which completes the proof. ■

Proof of Lemma 2. The extension to allow supplier specific quality parameters follows by analogous arguments to Lemma 1 noting that the buyer's value for supplier i's good is $\theta_i v$, which has distribution $\hat{F}(x) \equiv F(x/\theta_i)$ on $[\theta_i \underline{v}, \theta_i \bar{v}]$ with density $\hat{f}(x) = \frac{1}{\theta_i} f(v/\theta_i)$. Thus, the virtual type when the buyer's value is v is

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

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

Proof of Proposition 10. 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 $\delta_i(\hat{p}) \equiv \hat{p} - \Gamma^{-1}_{i,w_i\beta^{\mathbf{w}}}(\hat{p})$ and a fee schedule given by, for all $i \in \mathcal{N}$,

$$\sigma_i(p) \equiv \Phi^{-1}_{w_B\beta^{\mathbf{w}}}(\Gamma_{i,w_i\beta^{\mathbf{w}}}(\Gamma_i^{-1}(p))) - p,$$

and suppose that the buyer sets a reserve of $\Phi_{w_B\beta^{\mathbf{w}}}(v)$. 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 $\Gamma_{i,w_i\beta^{\mathbf{w}}}(c_i)$, and so supplier i wins if and only if

$$\Gamma_{i,w_i\beta^{\mathbf{w}}}(c_i) = \min_{j \in \mathcal{N}} \Gamma_{j,w_j\beta^{\mathbf{w}}}(c_j) \leq \Phi_{w_B\beta^{\mathbf{w}}}(v),$$

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

$$\hat{p} \equiv \min_{j \in \mathcal{N} \setminus \{i\}} \{\Phi_{w_B\beta^{\mathbf{w}}}(v), \Gamma_{j,w_j\beta^{\mathbf{w}}}(c_j)\},$$

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 $\Phi_{w_B\beta^{\mathbf{w}}}(v)$, 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 $\Phi_{w_B\beta^{\mathbf{w}}}(v)$.

To reduce notation, let $x_B \equiv w_B\beta^{\mathbf{w}}$ and $x_i \equiv w_i\beta^{\mathbf{w}}$. 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

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

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 $r = \Phi_{x_B}(v) = \Phi_{w_B\beta^{\mathbf{w}}}(v)$, 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 $\eta_i \pi^{\mathbf{w}}$. ■

C Appendix: *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 $v \geq c\frac{1+k}{2-k} + \frac{1-k}{2}$.[54] 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.[55] These outcomes are illustrated in Figure C.1.

To see that incomplete information bargaining encompasses the k-double auction as a special case, note that for the case of one supplier, the allocation $Q_i^{\mathbf{w}}(v, \mathbf{c})$ is the same for

[54] 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.

[55] There is an interesting relation between the effects of equalizing bargaining weights, which is at the center of our analysis, and the effects of equalizing ownership shares in a partnership model à la Cramton et al. (1987). Both effects increase equally weighted social surplus, but, as discussed after Corollary 1, for fundamentally different reasons.

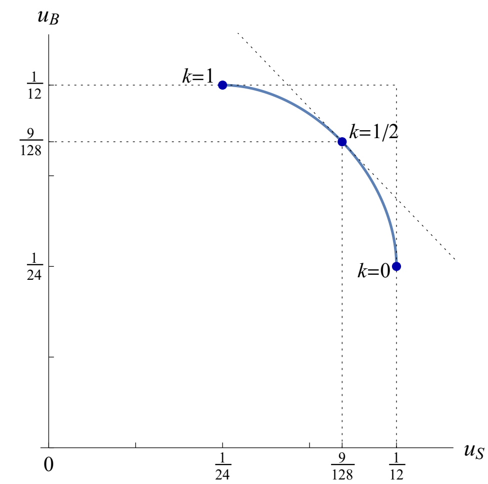

Figure C.1: Payoffs in the k-double auction for all $k \in [0,1]$. Assumes that there is one supplier and that the buyer's value and the supplier's cost are uniformly distributed on $[0,1]$.

all $\mathbf{w}$ with the same bargaining differential Δ defined by

$$\Delta \equiv \frac{w_B - w_1}{\max\{w_B, w_1\}} \in [-1,1].$$

Further, we can span $\Delta \in [-1,1]$ with bargaining weights (w_B, w_1) satisfying $\max\{w_B, w_1\} = 1$, so we restrict attention to such (w_B, w_1) in what follows. Under this restriction, there is a one-to-one mapping between (w_B, w_1) and Δ.

Under the assumption of one supplier and uniformly distributed types, for all $\mathbf{w}$, $\beta^{\mathbf{w}}$ is such that

$$\begin{aligned} 0 &= \mathbb{E}_{v,\mathbf{c}}\left[(\Phi(v) - \Gamma(c)) \cdot Q_1^{\mathbf{w}}(v,c)\right] \\ &= \int_{\frac{1-w_B\beta^{\mathbf{w}}}{2-w_B\beta^{\mathbf{w}}}}^{1} \int_0^{\frac{v-(1-w_B\beta^{\mathbf{w}})(1-v)}{2-w_S\beta^{\mathbf{w}}}} (2v-1-2c)\,dc dv, \end{aligned}$$

where the second equality uses the expression for $Q_1^{\mathbf{w}}(v,c)$ from Lemma 1. Solving this for $\beta^{\mathbf{w}}$ and using $\max\{w_B, w_S\} = 1$, we get

$$\beta^{\mathbf{w}} = \begin{cases} \frac{2w_B + 2w_S - 2\sqrt{w_B^2 - w_B w_S + w_S^2}}{3 w_B w_S} & \text{if } 0 < \min\{w_B, w_S\}, \\ 1 & \text{if } 0 = \min\{w_B, w_S\}. \end{cases}$$

Making the substitution $\Delta = w_B - w_S$ and writing $\beta^{\mathbf{w}}$ as a function of Δ, we have $\beta^{\Delta} = 1$

for $\Delta \in \{-1, 1\}$ and otherwise β^Δ is given by:

$$\beta^\Delta = \frac{4 - 2|\Delta| - 2\sqrt{1 - |\Delta| + \Delta^2}}{3(1 - |\Delta|)}. \tag{32}$$

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

$$\Delta_k \equiv \frac{1 - 2k}{k^2 - \max\{1, 2k\}}. \tag{33}$$

To see that the price-formation mechanism with bargaining differential Δ_k is equivalent to the k-double auction, substitute the expression for β^Δ in place of $\beta^{\mathbf{w}}$ into the expression derived from Lemma 1 for $Q_1^{(1+\Delta,1)}(v, c)$ if $\Delta \in [-1, 0]$ and for $Q_1^{(1,1-\Delta)}(v, c)$ if $\Delta \in (0, 1]$ to get

$$Q_1^\Delta(v, c) \equiv \begin{cases} 1 & \text{if } \Delta \in (-1, 0] \text{ and } v \geq \frac{2c\left(\sqrt{\Delta^2+\Delta+1}+2\Delta+1\right)+\left(2\sqrt{\Delta^2+\Delta+1}-2\Delta-1\right)(\Delta+1)}{2(\Delta+1)\left(\sqrt{\Delta^2+\Delta+1}-\Delta+1\right)}, \\ & \text{or if } \Delta \in [0, 1) \text{ and } v \geq \frac{2c\left(\sqrt{\Delta^2-\Delta+1}+\Delta+1\right)(1-\Delta)+2\sqrt{\Delta^2-\Delta+1}-\Delta-1}{2\left(\sqrt{\Delta^2-\Delta+1}-2\Delta+1\right)}, \\ & \text{or if } \Delta = 1 \text{ and } v \geq 2c, \\ & \text{or if } \Delta = -1 \text{ and } v \geq \frac{c+1}{2}, \\ 0 & \text{otherwise.} \end{cases}$$

It then follows that $Q_1^{\Delta_k}(v, c)$ is the same allocation rule as for the k-double auction, i.e., there is trade if and only if $v \geq c\frac{1+k}{2-k} + \frac{1-k}{2}$.

Assuming one supplier and uniformly distributed types on $[0, 1]$, for any bargaining weights $\mathbf{w}$, there exists $k \in [0, 1]$ such that the outcome of the k-double auction is the same as the outcome of incomplete information bargaining with weights $\mathbf{w}$, and conversely, for any $k \in [0, 1]$, there exist bargaining weights $\mathbf{w}$ such that incomplete information bargaining with weights $\mathbf{w}$ yields the same outcome as the k-double auction.

D Appendix: Generalization

Let $\mathcal{P}$ be the set of subsets of $\mathcal{N}$ with no more than D 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 $\{\hat{\theta}_i\}_{i \in \mathcal{N}}$ and size-dependent discounts $\{\delta_i\}_{i \in \mathcal{N}}$ 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 if 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 $\hat{\theta}_i = \theta$ for some common θ and for $i \in \mathcal{N}$, δ_i either all zero for constant marginal value or increasing in i for decreasing marginal value; (ii) differentiated products by letting $\hat{\theta}_i$ differ by i and setting all δ_i to zero; (iii) a one-buyer version of the Shapley-Shubik model by setting $D = 1$; and (iv) a version of the Shapley-Shubik model in which the buyer has demand for multiple products of the suppliers by setting $D > 1$.

Define

$$X^*_\beta(v, \mathbf{c}) \in \arg\max_{X \in \mathcal{P}} \theta_X \Phi_\beta(v) - \sum_{i \in X} \Gamma_{i,\beta}(c_i),$$

i.e., $X^*_\beta(v, \mathbf{c})$ is the set of trading partners for the buyer that maximizes the difference between the ironed β-weighted virtual value, scaled by $\theta_{X^*_\beta(v,\mathbf{c})}$, and the ironed β-weighted virtual costs of the trading partners. We then define β^* to be the largest $\beta \in [0, 1]$ such that

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

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

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

and the expected payoff of supplier i is

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

E Appendix: Bargaining externalities

In this appendix, we adopt the generalized setup of Section 6.2 with multi-unit demand and buyer preferences over suppliers to illustrate 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 assuming that $\theta_2 = 1$, we allow the buyer's preference for supplier 1, θ_1, and the buyer's total demand, D, 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 D and θ_1 vary as indicated in the table.

	$D=1$		$D=2$	
θ_1:	1	2	1	2
$\beta^{\mathbf{w},\boldsymbol{\theta}}$	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 $D = 1$, 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 $D = 2$, 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 $\beta^{\mathbf{w},\boldsymbol{\theta}}$ increases, 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.

F Appendix: (Unique) properties of the IPV paradigm

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, 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 countervailing power emerges.

Privacy of information endows economic agents with information rents and thereby protects them from hold up, as discussed in our analysis of investments. Even without investment, this protection implies, for example, that first-degree price discrimination is not possible. Rather than being an assumption, the impossibility of first-degree price discrimination is an implication in this setup. Likewise, setting a uniform market clearing price is the optimal mechanism for a monopoly with constant marginal costs facing a continuum of buyers, so under these conditions uniform pricing is a conclusion rather than an assumption.[56] Moreover, the aforementioned assumptions are essentially the *only* assumptions that permit a tractable approach that maintain the basic tradeoff between profit and social surplus. 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.

G 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

[56] In contrast, if the monopoly has increasing marginal costs and the revenue function that it faces is not concave, then setting non-market-clearing prices may be optimal (see Loertscher and Muir, 2020).

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 a case in point, Ausubel et al. (2002) note that 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 and $G_1, \ldots, G_n$ 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 assume 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.[57]

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

[57] 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.

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 buyer 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 $\beta^{\mathbf{w}} < 1/\max \mathbf{w}$, then Axioms 1–4 imply that $\hat{u}_B(\underline{v}) = \hat{u}_1(\overline{c}) = \cdots = \hat{u}_n(\overline{c}) = 0$, 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 $n = 1$ and $\underline{v} > \overline{c}$, all five axioms are satisfied using the posted-price mechanism introduced above 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, $\hat{u}_B(\underline{v})$ and $\hat{u}_1(\overline{c}), \ldots, \hat{u}_n(\overline{c})$, 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, $\mathbb{E}_{v,\mathbf{c}}[W^{\mathbf{w}}_{\mathbf{Q},\mathbf{M}}(v,\mathbf{c})]$, 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 buyer's 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* $5'(\mathbf{w})$*.*

References for the 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.

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

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

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.

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

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

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. (2002): *Auction Theory*, Elsevier Science, Academic Press.

Loertscher, S. and E. V. Muir (2020): "Monopoly pricing, optimal randomization, and resale," Working Paper, University of Melbourne.

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

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

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

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

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.

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

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