ONLINE APPENDIX

Incomplete Information Bargaining with Applications to Mergers, Investment, and Vertical Integration*

Simon Loertscher† and Leslie M. Marx‡

October 20, 2021

*Financial support from the Samuel and June Hordern Endowment, European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement 340903, University of Melbourne Faculty of Business & Economics Eminent Research Scholar Grant, and Australian Research Council Discovery Project Grant DP200103574 is gratefully acknowledged.

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

‡Duke University, 100 Fuqua Drive, Durham, NC 27708, USA, marx@duke.edu.

Contents

Page No.

A. Mechanism design foundations

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

1. Concepts and derivations

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

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

$$\text{(A.1)} \qquad \hat{q}_i^B(z) = \mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}}[Q_i^B(z, \mathbf{v}_{-i}, \mathbf{c})] \quad \text{and} \quad \hat{m}_i^B(z) = \mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}}[M_i^B(z, \mathbf{v}_{-i}, \mathbf{c})].$$

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

A. Key constraints. — The mechanism is *Bayesian incentive compatible* for buyer i if for all $v, z \in [\underline{v}, \overline{v}]$,

$$\text{(A.2)} \qquad \hat{u}_i^B(v) \equiv \hat{q}_i^B(v)v - \hat{m}_i^B(v) \geq \hat{q}_i^B(z)v - \hat{m}_i^B(z),$$

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

$$\text{(A.3)} \qquad \hat{u}_j^S(c) \equiv \hat{m}_j^S(c) - \hat{q}_j^S(c)c \geq \hat{m}_j^S(z) - \hat{q}_j^S(z)c.$$

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

deficit condition if

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

B. INTERIM EXPECTED PAYOFFS. — Standard arguments (see, e.g., Krishna, 2010, Chapter 5.1) proceed as follows:

Focusing on buyer i, incentive compatibility implies that

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

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

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

and for $\varepsilon < 0$ implies

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

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

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

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

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

[1] A function $h : [\underline{v},\overline{v}] \to \mathbb{R}$ is absolutely continuous if for all $\varepsilon > 0$ there exists $\delta > 0$ such that whenever a finite sequence of pairwise disjoint sub-intervals (v_k, v_k') of $[\underline{v},\overline{v}]$ satisfies $\sum_k (v_k' - v_k) < \delta$, then $\sum_k \left|h(v_k') - h(v_k)\right| < \varepsilon$. 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 $x \in [a,b]$.

C. Mechanism budget surplus. — Using the definitions of $\hat{u}_i^B$ and $\hat{u}_j^S$ in (A.2) and (A.3), we can rewrite these as

$$\hat{m}_i^B(v) = \hat{q}_i^B(v)v - \int_{\underline{v}}^{v} \hat{q}_i^B(t)dt - \hat{u}_i^B(\underline{v}) \tag{A.4}$$

and

$$\hat{m}_j^S(c) = \hat{q}_j^S(c)c + \int_{c}^{\overline{c}} \hat{q}_j^S(t)dt + \hat{u}_j^S(\overline{c}). \tag{A.5}$$

The expected payment by buyer i is then

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

where the first equality uses the definition of the expectation, the second uses (A.4), the third switches the order of integration, the fourth integrates, the fifth collects terms, the sixth uses the definition of the virtual value Φ_i, and the last equality uses the definition of the expectation. Similarly, using (A.5), the expected payment to supplier j is

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

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

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

2. *Myerson-Satterthwaite redux*

Consistent with the setup of Myerson and Satterthwaite (1983), this subsection focuses on the case of one seller and one buyer.

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

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

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

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

subject to the no-deficit constraint,

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

and the individual rationality constraints

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

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

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

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

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

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

It follows that for a given ρ, the Lagrangian is maximized with respect to Q pointwise by setting $Q(v,c)=1$ if $\Phi^{1/\rho}(v)\geq\Gamma^{1/\rho}(c)$ and $Q(v,c)=0$ other-

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

wise. Because the virtual types are increasing, this pointwise maximizer, denoted $Q^{\rho}(v, c)$, is increasing in v and decreasing in c for any $\rho \geq 1$. Hence, there is an incentive compatible implementation.

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

B. Proofs

Proof of Lemma 1. We can write (6) as

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

It is then clear that the allocation rule defined in the statement of the lemma maximizes the Lagrangian pointwise subject to the feasibility constraints. Combining this with the optimized Langrange multiplier, $\rho^{\mathbf{w}}$, completes the proof. ■

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

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

ingly $\mathcal{C}$ is also convex. The incomplete information bargaining mechanism solves

$$\max_{\mathbf{u}\in\mathcal{C}} \sum_{j\in\mathcal{N}^S} w_j^S u_j^S + \sum_{i\in\mathcal{N}^B} w_i^B u_i^B. \tag{B.1}$$

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

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

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

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

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

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

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

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

Case 1. Suppose that $K^B \leq K^S$ and (9) fails to hold. (Analogous arguments apply if $K^B > K^S$ and (9) fails.) Then for an open set of types, not all of the n^B buyers trade under $\mathbf{Q}^{\mathbf{w}}$. Thus, in order for $\mathbf{Q}^{\mathbf{w}}$ and $\mathbf{Q}^*$ to coincide, they must agree on not only the ranking within buyers and within suppliers, but also the ranking across buyers and suppliers. Consistent with (ii)–(iv), suppose that all buyers have the same bargaining weight w^B, all suppliers have the same bargaining weight w^S, $w^S < w^B$, and all suppliers have the same distribution. (Analogous analysis applies if $w^S > w^B$ and all buyers have the same distribution.) Then the planner and market both rank the buyers the same and rank the suppliers the same, but they evaluate the buyers' virtual values using weight $w^B/\rho^{\mathbf{w}}$ and the suppliers' virtual costs using weight $w^S/\rho^{\mathbf{w}}$, where $w^B/\rho^{\mathbf{w}} > w^S/\rho^{\mathbf{w}}$. Because either $w^B/\rho^{\mathbf{w}} \neq 1/\rho^{\mathbf{1}}$ or $w^S/\rho^{\mathbf{w}} \neq 1/\rho^{\mathbf{1}}$ or both, $\mathbf{Q}^{\mathbf{w}}(\mathbf{v},\mathbf{c}) \neq \mathbf{Q}^*(\mathbf{v},\mathbf{c})$ for all $(\mathbf{v},\mathbf{c})$ in an open subset of $[\underline{v},\overline{v}]^{n^B} \times [\underline{c},\overline{c}]^{n^S}$.

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

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

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

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

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

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

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

the individual rationality constraints

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

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

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

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

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

which we can rewrite as

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

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

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

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

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

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

We are left to map this back to the original problem of maximizing over weights $\mathbf{w}$ rather than mechanisms. Observe that the payoff profiles constructed in the paragraph above are Pareto undominated. Hence, by Proposition 2, there exist weights associated with each of those such that they are the payoffs generated by the incomplete information bargaining mechanism. Thus, $\omega(u)$ is concave.

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

$$\underline{v} \leq \overline{c}. \tag{B.2}$$

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

$$\sum_{j \in \mathcal{N}^S} u_j^S(\Delta_\lambda) = \lambda \sum_{j \in \mathcal{N}^S} u_j^S(\Delta') + (1-\lambda) \sum_{j \in \mathcal{N}^S} u_j^S(\Delta''), \tag{B.3}$$

and

$$\sum_{i \in \mathcal{N}^B} u_i^B(\Delta_\lambda) = \lambda \sum_{i \in \mathcal{N}^B} u_i^B(\Delta') + (1-\lambda) \sum_{i \in \mathcal{N}^B} u_i^B(\Delta''). \tag{B.4}$$

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

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

Define $y \equiv \max\{\underline{c}, \underline{v}\}$. Then $y \in [\underline{v}, \overline{v})$ and, using (B.2), $y \in [\underline{c}, \overline{c}]$, which says that y is in the range of the buyers' weighted virtual values and suppliers' weighted virtual costs. We have for all $i \in \mathcal{N}^B$,

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

Thus, for

$$v_i \in \left(\Phi_i^{1/\rho'^{-1}}(y), \Phi_i^{1/\rho_\lambda^{-1}}(y)\right), \tag{B.5}$$

we have

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

Using the continuity of the weighted virtual cost functions, it follows that for an open set of types $(\mathbf{v}, \mathbf{c})$ such that for all $i \in \mathcal{N}^B$, v_i satisfies (B.5) and for all $j \in \mathcal{N}^S$, c_j is in a sufficiently tight neighborhood of $\Gamma_j^{\frac{1-\Delta'}{\Delta'}/\rho'^{-1}}(y)$, we have for all $i \in \mathcal{N}^B$ and $j \in \mathcal{N}^S$,

$$\Phi_i^{1/\rho''}(v_i) < \Phi_i^{1/\rho_\lambda}(v_i) < \Gamma_j^{\frac{1-\Delta'}{\Delta'}/\rho'}(c_j) < \Phi_i^{1/\rho'}(v_i). \tag{B.6}$$

Recalling from Proposition 1 the incomplete information bargaining allocation rule and letting $\mathbf{Q}^\Delta$ denote the unique incomplete information bargaining allocation rule associated with Δ, (B.6) implies that for an open set of type, $\mathbf{Q}^{\Delta''}$ and $\mathbf{Q}^{\Delta_\lambda}$ specify no trade, while $\mathbf{Q}^{\Delta'}$ has trade. Hence, the allocation rule $\mathbf{Q}^{\Delta_\lambda}$ is *not* a convex combination of the allocation rules $\mathbf{Q}^{\Delta'}$ and $\mathbf{Q}^{\Delta''}$. Because a convex combination of $\mathbf{Q}^{\Delta'}$ and $\mathbf{Q}^{\Delta''}$ implies a convex combination of the payoffs, it follows that $\mathbf{Q}^{\Delta_\lambda}$ does not induce a convex combination of the payoffs, a contradiction. Hence, the Williams frontier must be strictly concave.

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

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

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

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

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

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

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

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

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

and define buyer i's allocation rule by

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

For suppliers 1 and 2, define the allocation rule by

$$\tilde{Q}_1^S(\mathbf{v},\mathbf{c}) \equiv \begin{cases} \hat{Q}_{1,2}^S(\mathbf{v},\min\{c_1,c_2\},\mathbf{c}_{-\{1,2\}}) & \text{if } c_1 \leq c_2, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\tilde{Q}_2^S(\mathbf{v},\mathbf{c}) \equiv \begin{cases} \hat{Q}_{1,2}^S(\mathbf{v},\min\{c_1,c_2\},\mathbf{c}_{-\{1,2\}}) & \text{if } c_2 < c_1, \\ 0 & \text{otherwise,} \end{cases}$$

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

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

LEMMA B.1: *The expected threshold payments of nonmerging firms are the same under $\tilde{\mathbf{Q}}$ and $\hat{\mathbf{Q}}$.*

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

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

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

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

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

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

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

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

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

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

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

and note that

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

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

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

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

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

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

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

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

Proof of Proposition 7. We begin by considering the first part of the statement. Choosing $\underline{v}$ such that $\underline{v} \geq \overline{c}$, we have nonoverlapping supports. Hence, with symmetric bargaining weights, the pre-integration market achieves the first-best. Consequently, vertical integration cannot increase social surplus. 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 $n^S - 1 \geq 1$ suppliers). For this setting, impossibility of first-best trade obtains (see, e.g., Delacrétaz et al., 2019), regardless of bargaining weights. The second part follows from Theorem 4 (and Table 1) in Williams (1999), which shows that $\underline{v} > \underline{c}$ and n^S sufficiently large is sufficient for first-best to be possible when the suppliers draw their types from identical distributions. ■

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

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

$$\text{(B.7)} \qquad \hat{u}^B_{i,\mathbf{Q}}(v_i; \mathbf{e}^B_{-i}, \mathbf{e}^S) = \int_{\underline{v}}^{v_i} q^B_i(x; \mathbf{e}^B_{-i}, \mathbf{e}^S)dx,$$

and, taking expectations with respect to v_i, one obtains

$$\text{(B.8)} \qquad u^B_{i,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{v}}^{\overline{v}} q^B_i(x; \mathbf{e}^B_{-i}, \mathbf{e}^S)(1 - F_i(x; e^B_i))dx$$

up to a constant, and, analogously,

$$\text{(B.9)} \qquad u^S_{j,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{c}}^{\overline{c}} q^S_j(x; \mathbf{e}^B, \mathbf{e}^S_{-j})G_j(x; e^S_j)dx$$

up to a constant.

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

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

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

$$\text{(B.10)} \qquad \overline{e}^B_i \in \arg\max_{e^B_i} u^B_{i,\mathbf{Q}^{FB}}(e^B_i, \overline{\mathbf{e}}^B_{-i}, \overline{\mathbf{e}}^S) - \Psi^B_i(e^B_i)$$

and

$$\overline{e}_j^S \in \arg\max_{e_j^S} u_{j,\mathbf{Q}^{FB}}^S(\overline{\mathbf{e}}^B, e_j^S, \overline{\mathbf{e}}_{-j}^S) - \Psi_j^S(e_j^S). \tag{B.11}$$

Assume that (10)–(11) hold. Let $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ denote the incomplete information bargaining allocation rule given in Lemma 1, but with the virtual types defined in terms of the type distributions associated with investment $\overline{\mathbf{e}}$, and let $\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}$ denote the associated multiplier on the no-deficit constraint. Suppose that first-best investments $\overline{\mathbf{e}}$ are Nash equilibrium investments, which implies that for all $i \in \mathcal{N}^B$ and $j \in \mathcal{N}^S$,

$$\overline{e}_i^B \in \arg\max_{e_i^B} u_{i,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}^B(e_i^B, \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) - \Psi_i^B(e_i^B) \tag{B.12}$$

and

$$\overline{e}_j^S \in \arg\max_{e_j^S} u_{j,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}^S(\overline{\mathbf{e}}^B, e_j^S, \overline{\mathbf{e}}_{-j}^S) - \Psi_j^S(e_j^S). \tag{B.13}$$

Assumptions (10)–(11) ensure that the first-best investments are characterized by their first-order conditions. Thus, using (B.8) and (B.10), we have for all $i \in \mathcal{N}^B$,

$$-\int_{\underline{v}}^{\overline{v}} q_i^{FB,B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx - \Psi_i^{B\prime}(\overline{e}_i^B) = 0. \tag{B.14}$$

Similarly, using (B.8) and (B.12), we have

$$-\int_{\underline{v}}^{\overline{v}} q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx - \Psi_i^{B\prime}(\overline{e}_i^B) = 0. \tag{B.15}$$

Combining (B.14) and (B.15), we have

$$\int_{\underline{v}}^{\overline{v}} (q_i^{FB,B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) - q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S)) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx = 0. \tag{B.16}$$

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

$$\mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}|\overline{\mathbf{e}}_{-i}^B,\overline{\mathbf{e}}^S}\left[\int_{\underline{v}}^{\overline{v}} (Q_i^{FB,B}(x, \mathbf{v}_{-i}, \mathbf{c}) - Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x, \mathbf{v}_{-i}, \mathbf{c})) \tfrac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx\right] = 0. \tag{B.17}$$

Steps analogous to those leading to (B.16) imply that for all $j \in \mathcal{N}^S$,

$$\int_{\underline{c}}^{\overline{c}} (q_j^{FB,S}(x; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-j}^S) - q_j^{\mathbf{w},\overline{\mathbf{e}},S}(x; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-j}^S)) \frac{\partial G_j(x; \overline{e}_j^S)}{\partial e} dx = 0. \tag{B.18}$$

By Lemma 1, we know that the total number of trades induced by $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}(\mathbf{v}, \mathbf{c})$ is the maximum such that the lowest weighted virtual value of any trading buyer is greater than or equal to the highest weighted virtual cost of any trading supplier. Further, the total number of trades induced by $\mathbf{Q}^{FB}(\mathbf{v}, \mathbf{c})$ is the maximum such that the lowest value of any trading buyer is greater than or equal to the highest cost of any trading supplier. Because virtual costs are greater than or equal to actual costs and virtual values are less than or equal to actual values, it follows that $\sum_{i \in \mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v}, \mathbf{c}) \leq \sum_{i \in \mathcal{N}^B} Q_i^{FB,B}(\mathbf{v}, \mathbf{c})$ for all $(\mathbf{v}, \mathbf{c})$ (and similarly on the supply side). Because we assume that $\frac{\partial F_i(v;e)}{\partial e} < 0$ for all $v \in (\underline{v}, \overline{v})$, (B.17) then implies that

$$\sum_{i \in \mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v}, \mathbf{c}) = \sum_{i \in \mathcal{N}^B} Q_i^{FB,B}(\mathbf{v}, \mathbf{c}) \equiv \xi(\mathbf{v}, \mathbf{c}) \tag{B.19}$$

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

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

Case 1: $\underline{v} < \overline{c}$. Suppose, contrary to what we want to show, that $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ discriminates among agents based on virtual types for an open set of types—we then show that this implies that the number of trades under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ must sometimes differ from the number under the first-best, contradicting (B.19). That is, suppose that there exist suppliers (an analogous argument applies for buyers), which we denote by 1 and 2, and types $(\hat{\mathbf{v}}, \hat{\mathbf{c}})$ with $\hat{c}_1 \neq \hat{c}_2$ such that $Q_1^{FB,S}(\hat{\mathbf{v}}, \hat{\mathbf{c}}) > Q_1^{\mathbf{w},\overline{\mathbf{e}},S}(\hat{\mathbf{v}}, \hat{\mathbf{c}})$ and $Q_2^{FB,S}(\hat{\mathbf{v}}, \hat{\mathbf{c}}) < Q_2^{\mathbf{w},\overline{\mathbf{e}},S}(\hat{\mathbf{v}}, \hat{\mathbf{c}})$. Because supplier 1 trades under the first-best when supplier 2 has excess capacity, this implies that $\hat{c}_1 < \hat{c}_2$; and because supplier 2 trades in the Nash equilibrium when supplier 1 has excess capacity, this implies that $\Gamma_2^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_2; \overline{e}_2^S) \leq \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_1; \overline{e}_1^S)$. It follows that

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

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

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

which implies that no trades occur under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ and only supplier 1 and buyer 1 trade under the first-best. By continuity, for all $(\mathbf{v}, \mathbf{c})$ in an open set of types around $(\tilde{\mathbf{v}}, \tilde{\mathbf{c}})$, we have $\sum_{i \in \mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v}, \mathbf{c}) \neq \sum_{i \in \mathcal{N}^B} Q_i^{FB,B}(\mathbf{v}, \mathbf{c})$, which contradicts (B.19). 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 suppliers to produce as does $\mathbf{Q}^{FB}$. An analogous argument shows that the set of trading buyers is the same under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ as under $\mathbf{Q}^{FB}$.

Case 2: $\underline{v} \geq \overline{c}$ and either (ii), (iii), or (iv) holds. Note that $\underline{v} \geq \overline{c}$ implies that under the first-best, the number of trades is $\min\{K^B, K^S\}$. If (ii) holds, i.e., $K^B = K^S$, then all agents trade under the first-best and so (B.19) implies that all agents also trade under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$, which completes the proof. Suppose that (iii) holds, so that $K^B < K^S$ and (12) holds. (Analogous arguments apply to the case with $K^B > K^S$ and (13).) Then all buyers consume their full demands under the first-best. By the argument given in the proof of Lemma 2, because bargaining is efficient, we must have $w_1^S = \cdots = w_{n^S}^S$. Given this, (12) implies that the ranking of suppliers according to $\Gamma_i^{w_i^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c_i; \overline{e}_i^S)$ is the same as the ranking according to c_i. Thus, using (B.19), $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ induces the same suppliers to produce as does $\mathbf{Q}^{FB}$, and again we are done. ■

C. Extensions

In Section C.1, we show how data on bargaining breakdown can be used for estimation. In Section C.2, we extend the model to allow heterogeneous outside options, and in Section C.3, we extend the model to allow buyers to have heterogeneous preferences over suppliers and provide a generalization of the one-to-many setup that encompasses additional models.

1. Bargaining breakdown and estimation

A pervasive feature of real-world bargaining is that negotiations often break down.[4] Anecdotal examples range from the U.S. government shut down, to the British coal miners' and the U.S. air traffic controllers' strikes in the 1980s, to failures to form coalition governments in countries with proportional representation. Providing systematic evidence of bargaining breakdown, Backus et al. (2020) analyze a data set covering 25 million observations of bilateral negotiations on eBay and find a breakdown probability of roughly 55 percent. More generally, when firms bargain over essential inputs, such as medical equipment for hospitals or computer chips for manufacturers, bargaining breakdown will typically not mean that the firms stop trading with each other, but rather that the latest, quality-improved version of the input is not traded.

Because with incomplete information, bargaining breakdown occurs on the equilibrium path, one can use observed bargaining breakdown frequencies as a moment to match in empirical research rather than as an error.[5] While in practical applications specifics will, of course, depend on the available data and on the econometric approach, we now provide an illustration of how, with incomplete information bargaining, one can use observed frequencies of negotiation breakdowns to back out the parameters of the distributions from which the agents draw their types.

Consider a market with one buyer and two suppliers with single-unit demand and supply and types drawn from parameterized distributions

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

[4] 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.

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

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

The model can also be used to predict the change in social surplus as result of vertical integration, which is displayed in Figure C.1(b). As shown, when the rate of bargaining breakdown in the pre-integration market is sufficiently low, i.e., the pre-integration market is sufficiently efficient, the change in social surplus from vertical integration is negative. In contrast, when the probability of breakdown is sufficiently high prior to integration, the increased efficiency associated with internal transactions dominates, and vertical integration increases social surplus.

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)

(a) Calibration of distributions to data

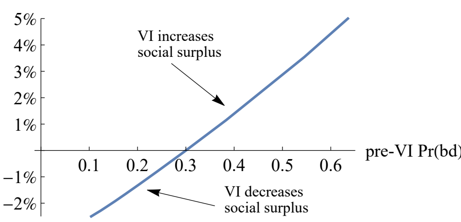

(b) Change in expected social surplus following vertical integration

FIGURE C.1. PRE-INTEGRATION BREAKDOWN AND EFFECTS OF VERTICAL INTEGRATION ON SOCIAL SURPLUS.

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

2. *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.

The comparative statics with respect to increasing an agent's participation cost (see Appendix C.2.A) are intuitive and largely the same as in models with complete information because it increases the agent's share of the surplus that is created; in contrast to complete information models, it may decrease expected social surplus because of distortions to the allocation rule required to cover larger outside options.

The effects of changing an agent's *production-relevant* outside option (see Appendix C.2.B) are more nuanced. For example, as a supplier's outside option improves, the support of its cost distribution shifts upwards by the amount of the improvement, with the result that higher costs become more likely. Hence, the supplier will tend to be less likely to trade. However, under the assumption of monotone hazard rates, this effect is partly (but not completely) offset because, for a given cost realization, the supplier's weighted virtual cost is lower than before the increase in the outside option. This implies that, ex post, given the same cost realization, the supplier is treated more favorably after the outside option increases. This is in line with intuition gleaned from complete information models. But from an ex ante perspective, the increase in the outside option reduces the supplier's expected payoff from incomplete information bargaining because overall it makes the supplier less likely to trade and thereby decreases the supplier's ex ante expected payoff. Moreover, as a supplier's cost distribution worsens, the revenue constraint faced by the mechanism becomes tighter, which further tends to worsen the agent's bargaining outcome.

A. FIXED COSTS OF PARTICIPATING IN THE MECHANISM. — For the purposes of this extension, we assume that $n^B = K^B = 1$ and drop the buyer subscripts. Therefore, we can also assume, without further loss, that $k_j^S = 1$ for all $j \in \mathcal{N}^S$.

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

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

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

$\rho \geq \max \mathbf{w}$ such that

(C.2)

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

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

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

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

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

$$\Gamma_{j,a}^o(c) \equiv c+(1-a)\frac{G_j(c-y_j)}{g_j(c-y_j)} = \Gamma_{j,a}(c-y)+y < \Gamma_{j,a}(c), \tag{C.3}$$

where the inequality holds for all $a < 1$ because the monotone hazard rate assumption implies that $\Gamma_{j,a}'(c) > 1$ for all $a < 1$. This in turn has two, somewhat subtle implications. Let z be the threshold for supplier j to trade when its outside option is zero, i.e., keeping z fixed, supplier j trades if and only if $\Gamma_{j,a}(c) \leq z$. (Note that z will be the minimum of the buyer's weighted virtual value and the smallest weighted virtual cost of supplier j's competitors, but this does not matter for the argument that follows.) Assuming that $a < 1$ and $y_j < \bar{c} - \underline{c}$, which implies that $\underline{c}_j < \bar{c}$, it follows that there are costs $c \in [\underline{c}_j, \bar{c}]$ and thresholds z such that supplier j trades when it has the outside option and not without it, that is, $\Gamma^o_{j,a}(c) < z < \Gamma_{j,a}(c)$. This reflects the reasonably well-known result that optimal mechanisms tend to discriminate in favor of weaker agents (McAfee and McMillan, 1987), which in this case is the agent with the positive outside option. It also resonates with intuition from complete information models: keeping costs fixed, the agent with the better outside option is treated more favorably, indeed, it is evaluated according to a smaller weighted virtual cost. However, from an ex ante perspective, the larger is the value of the outside option, the less likely is the agent to trade. To see this, consider a fixed realization of z. (The distribution of these thresholds is not affected by supplier j's outside option and hence our argument extends directly once one integrates over z and its density.) Given y_j, supplier j trades if and only if its cost c is below $\tau(y)$ satisfying $\Gamma^o_{j,a}(\tau(y)) = z$. Using (C.3), this is equivalent to $\Gamma_{j,a}(\tau(y)-y)+y = z$, which in turn is equivalent to $\tau(y) = \Gamma^{-1}_{j,a}(z - y) + y$, whose derivative for $a < 1$ satisfies

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

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

3. *Buyer preferences over suppliers and bargaining externalities*

To allow for and investigate bargaining externalities, we restrict attention to the case of one buyer, $n^B = 1$, with demand for $K^B \geq 1$ units, and $n^S \geq 2$ suppliers, but we generalize the setup to allow the buyer to have heterogeneous preferences over the suppliers. To this end, we let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_{n^S})$ be a commonly known vector of taste parameters of the buyer, with the meaning that the value to the buyer of trade with supplier j when the buyer's type is v is $\theta_j v$. Thus, under (ex post) efficiency, trade should occur between the buyer and supplier j if and only if $\theta_j v - c_j$ is positive and among the K^B highest values of $(\theta_\ell v - c_\ell)_{\ell \in \mathcal{N}^S}$. The problem is trivial if $\max_{j \in \mathcal{N}^S} \theta_j \overline{v} \leq \underline{c}$ because then it is never ex post efficient to have trade with any supplier, so assume that $\max_{i \in \mathcal{N}} \theta_i \overline{v} > \underline{c}$. This setup encompasses (i) differentiated products by letting the supplier-specific taste parameters differ; (ii) a one-buyer version of the Shapley and Shubik (1972) model by setting $K^B = 1$; and (iii) a version of the Shapley-Shubik model in which the buyer has demand for multiple products of the suppliers by setting $K^B > 1$. For a generalization of the one-to-many setup that encompasses additional models, see Section 3.B.

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

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

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

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

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

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

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

The result of Proposition C.1 does not necessarily extend to the case with $K^B < n$, as shown in the following example.

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

TABLE C.1—OUTCOMES FOR ONE-TO-MANY PRICE FORMATION.

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

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

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

B. GENERALIZATION OF BUYER PREFERENCES OVER SUPPLIERS. — Here we provide a further generalization of the setup with one buyer and multiple suppliers to allow a more general structure for the buyer's preferences over suppliers.

Let $\mathcal{P}$ be the set of subsets of $\mathcal{N}^S$ with no more than K^B elements (including the empty set) and let $\boldsymbol{\theta} = \{\theta_X\}_{X \in \mathcal{P}}$ be a commonly known vector of taste parameters of the buyer satisfying the "size-dependent discounts" condition of Delacrétaz et al. (2019). Specifically, let there be supplier-specific preferences $\{\hat{\theta}_j\}_{j\in\mathcal{N}^S}$ and size-dependent discounts $\{\delta_j\}_{j\in\mathcal{N}^S}$ with $0 = \delta_0 = \delta_1 \leq \delta_2 \leq \cdots \leq \delta_n$ such that for all $X \in \mathcal{P}$, $\theta_X = \sum_{i\in X} \hat{\theta}_i - \delta_{|X|}$. Thus, the buyer's value for purchasing from suppliers in $X \in \mathcal{P}$ when its type is v is $\theta_X v$, which depends on the buyer's value, the buyer's preferences for standalone purchases from the suppliers in X, and a discount that depends on the total number of units purchased. Note that $\theta_\varnothing = 0$, so that the value to the buyer of no trade is zero.

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

Define

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

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

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

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

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

and the expected payoff of supplier j is

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

D. 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.[6]

In light of this, it is perhaps useful to note that, as mentioned in the body of the paper and discussed in Appendix D.1, for the case of one supplier, one buyer, and uniformly distributed types, the k-double auction of Chatterjee and Samuelson (1983) provides an extensive-form game that delivers the same outcomes as incomplete information bargaining. In addition, as we show in Appendix D.2, our approach has axiomatic foundations analogous to those that underpin Nash bargaining. Further, intermediaries like eBay, Amazon, and Alibaba play a prominent trade in organizing markets and, as we show in Appendix D.3, provide micro-foundations for incomplete information bargaining. Specifically, building on the model of Loertscher and Niedermayer (2019), for general distributions, one buyer, and any number of suppliers, the incomplete information bargaining outcome arises in equilibrium in an extensive-form game involving a buyer, suppliers, and a fee-setting broker. This is reminiscent of the role of intermediaries in the wholesale used car market as described by Larsen (2021). There, auction houses run auctions, facilitate further bargaining in the substantial number of cases in which the auction does not result in trade, and collect fees from traders.

1. k-double auction as a special case

In the k-double auction of Chatterjee and Samuelson (1983), given $k \in [0, 1]$, the buyer and supplier in a k-double auction simultaneously submit bids p_B and p_S, and trade occurs at the price $kp_B + (1-k)p_S$ if and only if $p_B \geq p_S$. By construction, the k-double auction never incurs a deficit. If the agents' types are uniformly distributed on $[0, 1]$, then the linear Bayes Nash equilibrium of the k-double auction results in trade if and only if $v \geq c\frac{1+k}{2-k} + \frac{1-k}{2}$.[7] 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

[6]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.

[7]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.

$k \in [0,1]$, the k-double auction implements the outcomes of incomplete information bargaining for some bargaining weights. These outcomes are illustrated in Figure D.1.

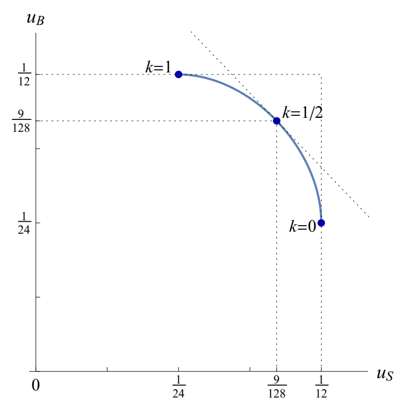

FIGURE D.1. PAYOFFS IN THE k-DOUBLE AUCTION.

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

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

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

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

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

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

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

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

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

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

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

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

Using (D.2), it then follows that

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

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

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

2. *Axiomatic approach*

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Finally, Nash bargaining specifies, in addition to efficiency and symmetry, axioms of invariance to affine transformations of the utility functions and independence of 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 D.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 D.1. When $\mathbf{w}$ is symmetric, then by definition, the incomplete information bargaining mechanism maximizes welfare subject to incentive compatibility, individual rationality, and no deficit. Further, because the allocation pins down the agents' interim expected payoffs up to a constant, the mechanism is unique up to the payoffs of the worst-off types, $\hat{u}_1^B(\underline{v}), \dots, \hat{u}_{n^B}^B(\underline{v})$ and $\hat{u}_1^S(\overline{c}), \dots, \hat{u}_{n^S}^S(\overline{c})$, but these are uniquely pinned down by the assumption of symmetric $\boldsymbol{\eta}$. ■

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

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

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

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

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

3. *Extensive-form approach*

Building on the model of Loertscher and Niedermayer (2019), we define the *fee-setting extensive-form game* to have one buyer with single-unit demand, $n^S = n \geq 1$ suppliers, and an intermediary that facilitates the buyer's procurement of inputs from the suppliers and that charges the buyer a fee for its service. Let $\mathcal{N} = \mathcal{N}^S$ denote the set of suppliers. The buyer's value and the suppliers' costs are not known by the intermediary, although the intermediary does know the distributions F and $G_1, \dots, G_n$ 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 $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$, 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.

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 D.1: *The outcome of incomplete information bargaining with bargaining weights* $\mathbf{w}$ *and shares* $\boldsymbol{\eta}$ *is a Bayes Nash equilibrium outcome of the fee-setting extensive-form game with intermediary preference weights* $\mathbf{w}$ *and profit shares* $\boldsymbol{\eta}$*.*

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

whose first-order condition when $\underline{c} < \Gamma_i(\Gamma_i^{x_i-1}(r))$ is

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

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 $r = \Phi^{x_B}(v) = \Phi^{w^B/\rho^{\mathbf{w}}}(v)$, giving the buyer nonnegative interim expected payoff, which completes the proof for the case with $\pi^{\mathbf{w}} = 0$. If $\pi^{\mathbf{w}} > 0$, then this "excess profit" must be distributed via fixed payments between the agents and the intermediary so that the worst-off type of each agent $i \in \{B\} \cup \mathcal{N}$ has interim expected payoff $\eta_i \pi^{\mathbf{w}}$. ■

Thus, the fee-setting extensive-form game, in which a fee-setting intermediary procures an input for the buyer from competing suppliers, provides a microfoundation for the price-formation mechanism. Reminiscent of Crémer and Riordan

(1985), the sequential nature of the game allows an equilibrium that is Bayesian incentive compatible for one agent, the buyer, and dominant-strategy incentive compatible for the other agents, the suppliers. The equilibrium of the fee-setting game satisfies ex post individual rationality for both the buyer and suppliers, but only balances the intermediary's budget in expectation. In contrast, in Crémer and Riordan (1985), the budget is balanced ex post, but individual rationality is no longer satisfied ex post for all agents.[8]

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

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

E. Details for investment comparative statics

In this appendix, we provide details underlying the comparative statics analysis in Section V, which examines how equilibrium investments are affected by bargaining power and by the extent to which the supports of the value and cost distributions overlap.

As described in the body of the paper, we consider a bilateral trade setup with linear virtual types. We hold fixed the support of the supplier's distribution at $[0, 1]$ and let the support of the buyer's distribution be $[\underline{v}, \underline{v}+1]$, where we vary $\underline{v}$ from 0 to 1. Specifically, we fix $X > 0$ and consider a supplier type distribution of $G_{e_S}(c) \equiv c^{X-e_S}$ with support $[0, 1]$, where $e_S \in [0, X)$ is the supplier's investment, and a buyer type distribution of $F_{e_B}(v) \equiv 1 - (1 + \underline{v} - v)^{X-e_B}$ with support $[\underline{v}, \underline{v} + 1]$, where $e_B \in [0, X)$ is the buyer's investment. We assume that each agent's investment e has cost $e^2/2$.

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

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

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

Second-best investment is also the same for the buyer and suppler and satisfies:[9]

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

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

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

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

[9] The linear virtual type functions are given by

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

solves

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

the first-order condition of which is

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

Analogously, the supplier's Nash equilibrium investment solves

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

whose first-order condition is

$$\text{(E.2)} \quad e_S^{NE} = \int_{\underline{v}}^{\underline{v}+1} \int_0^1 \left. \frac{\partial G_e(c)}{\partial e} \right|_{e=e_S^{NE}} \cdot \mathbf{1}_{\Gamma^{w_S/\rho^{NE}}(c;e_S^{NE}) \leq \Phi^{w_B/\rho^{NE}}(v;e_B^{NE})} \cdot f_{e_B^{NE}}(v) dc dv.$$

Solving for $(e_S^{NE}, e_B^{NE}) \in [0, X]^2$ and $\rho^{NE} \geq \max\{w_S, w_B\}$ that satisfy (E.1), (E.2),

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

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

We illustrate the effects of bargaining power and the distributional supports on equilibrium investment in Figure E.1, which expands upon Figure 3 in the body of the paper.

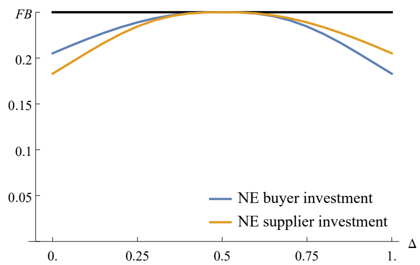

(a) Nonoverlapping supports ($\underline{v} = 1$)

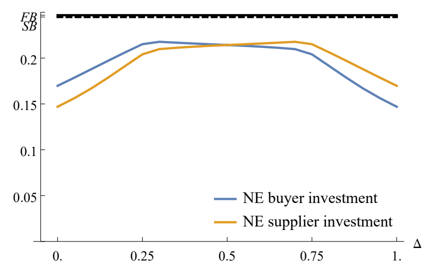

(b) Partially overlapping supports ($\underline{v} = 3/4$)

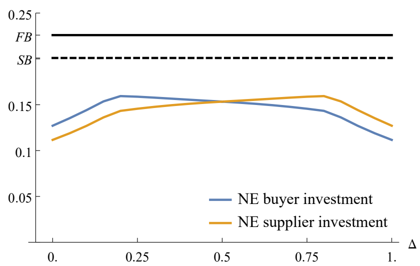

(c) Partially overlapping supports ($\underline{v} = 1/2$)

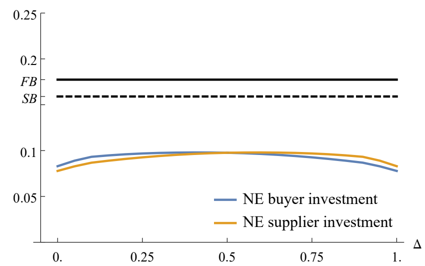

(d) Partially overlapping supports ($\underline{v} = 1/4$)

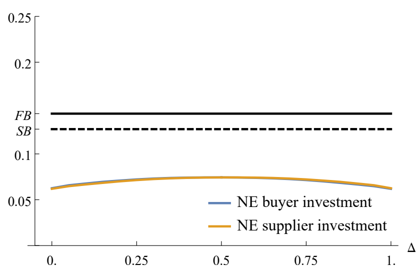

(e) Partially overlapping supports ($\underline{v} = 1/8$)

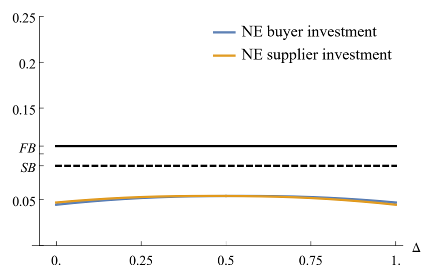

(f) Fully overlapping supports ($\underline{v} = 0$)

FIGURE E.1. NASH EQUILIBRIUM INVESTMENTS.

Notes: Nash equilibrium investments with bargaining weights $(w^S, w^B) = (1 - \Delta, \Delta)$ for buyer distributions with varying supports. Assumes the linear virtual type setup for bilateral trade with $F(v) = 1 - (1 + \underline{v} - v)^{1.25 - e_B}$, where $e_B \in [0, 1.25)$ is the buyer's investment, and $G(c) = c^{1.25 - e_S}$, where $e_S \in [0, 1.25)$ is the supplier's investment. Investment e has cost $e^2/2$. When $\underline{v} = 1$, we obtain $e^{FB} = e^{SB} = 0.25$, implying that first-best (and second-best) investment levels result in uniformly distributed types. For $\underline{v} = 1$, $\rho^{NE} = \max\{w_S, w_B\}$ for all bargaining weights. For $\underline{v} \in \{1/4, 1/8, 0\}$, $\rho^{NE} > \max\{w_S, w_B\}$ for all bargaining weights. For $\underline{v} \in \{1/2, 3/4\}$, $\rho^{NE} = \max\{w_S, w_B\}$ for sufficiently asymmetric bargaining weights and $\rho^{NE} > \max\{w_S, w_B\}$ otherwise.

F. Additional results

1. Horizontal mergers

A. Supplier mergers in the face of powerful buyers. — While antitrust authorities seem inclined to look more favorably upon a supplier merger when there are powerful buyers,[10] one can build on Proposition 5 to show that in the absence of bargaining power effects, a merger of suppliers that face powerful buyers can actually be worse for social surplus than a merger of powerful suppliers:

PROPOSITION F.1: *Consider a horizontal merger of suppliers i and j with $k_i^S = k_j^S = K^B$ that does not alter bargaining weights. If the pre-merger market is:*
(i) efficient, then the merger decreases expected social surplus unless $w_i = w_j = \max \mathbf{w}$;
(ii) inefficient and $G_i = G_j$, then the merger decreases expected social surplus by more if the buyers have all the bargaining power than if the merging suppliers have all the bargaining power.

Proof of Proposition F.1. A supplier merger of the type considered results in a merged entity with a weighted virtual cost function

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

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

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

Part (ii): Suppose that the pre-merger market is not efficient. Proposition 5 implies that if the merging suppliers have all the bargaining weight, then the

[10] "The Agencies consider the possibility that powerful buyers may constrain the ability of the merging parties to raise prices. ... However, the Agencies do not presume that the presence of powerful buyers alone forestalls adverse competitive effects flowing from the merger" (U.S. Department of Justice and the Federal Trade Commission, 2010, p. 27).

merger does not affect social surplus, positively or negatively. In contrast, if the buyers or a subset of buyers have all the bargaining weight, then the merger has no effect on the Lagrange multiplier on the no-deficit constraint because the buyer-optimal mechanism satisfies the no-deficit constraint when the multiplier is equal to its minimum value of $\max \mathbf{w}$. Thus, for a merger that does not alter bargaining weights, the only effect on the allocation rule comes through the effect on the merged entity's virtual cost function. Thus, a merger reduces social surplus if the merged entity's weighted virtual cost function $\Gamma_{1,2}^a$ satisfies $\Gamma_{1,2}^a(\min\{c_1, c_2\}) > \min\{\Gamma_1^a(c_1), \Gamma_2^a(c_2)\}$ for all $c_1, c_2 \in (\underline{c}, \overline{c})$ and $a \in [0, 1]$, which holds for symmetric suppliers, because then trade occurs for a strictly smaller set of realized types in the post-merger versus per-merger market. ■

Under the conditions of Proposition F.1, concerns regarding the welfare consequences of a merger are heightened when merging suppliers face a powerful buyer.

B. General conditions for weighted welfare reducing mergers. — In general, a merger removes an independent firm and creates a merged entity that draws its type from a distribution that can differ from the distributions of the pre-merger firms. We now consider conditions on the merged entity's type distribution that are sufficient for the result, as in Proposition 5, that a merger reduces expected weighted welfare. To develop intuition, consider the case of a supplier merger in which, as in Proposition 5, the merged entity draws its constant marginal cost from the distribution of the minimum of the two merging suppliers' marginal costs. Then one can essentially transfer to the pre-merger market the allocation rule of any incentive compatible post-merger mechanism by replacing the type of a merged entity that combines suppliers 1 and 2 with $\min\{c_1, c_2\}$ and allocating the merged entity's quantity to the merging supplier 1 or 2 with the lower cost. Using threshold payments, the budget surplus, not accounting for fixed payments, is then greater in the pre-merger market because the competition between suppliers 1 and 2 reduces the threshold payments to those suppliers. This means that the post-merger incomplete information bargaining mechanism is feasible in the pre-merger market—indeed has strictly greater budget surplus not accounting for fixed payments. If it also generates (weakly) greater expected weighted welfare, then it follows by a form of revealed preference argument that expected weighted welfare under the (optimized) pre-merger mechanism is (weakly) greater than under the post-merger mechanism.[11]

[11]If the merging suppliers do not have the maximum bargaining weight, then pre-merger expected weighted welfare can be increased by distributing the savings from reduced payments to the merging suppliers to firms with higher bargaining weights, yielding the result that expected weighted welfare is greater pre-merger. If the merging suppliers have the maximum bargaining weight and all other firms have lower bargaining weights, then one can achieve the same expected weighted welfare in the pre-merger market, and potentially greater expected weighted welfare once the mechanism is optimized for the pre-merger market. If the merging suppliers have the maximum bargaining weight and all other

We provide general conditions for this argument to apply in the following lemma. As stated in the lemma, we require that the post-merger distribution, $G_{1,2}$ in the case of a merger of suppliers 1 and 2, is equal to the distribution of some nondecreasing function of h the merging firms' types. The remaining conditions ensure that one can rank the threshold payments of the merging firms relative to the threshold payment of the merged entity.

LEMMA F.1: *A merger of suppliers* 1 *and* 2 *that does not alter bargaining weights or shares weakly reduces expected weighted welfare if the merged entity's cost distribution* $G_{1,2}$ *and capacity* $k^S_{1,2}$ *are such that there exists a continuous, nondecreasing function* $h : [\underline{c}, \overline{c}]^2 \to [\underline{c}, \overline{c}]$ *satisfying:*

(*i*) *for all* $z \in [\underline{c}, \overline{c}]$,
$$\Pr_{c_1,c_2}(h(c_1, c_2) \le z) = G_{1,2}(z),$$

(*ii*) *for all* $c_1, c_2 \in [\underline{c}, \overline{c}]$, $\min\{c_1, c_2\} \le h(c_1, c_2)$,

(*iii*) *if* $k^S_i < k^S_{1,2}$ *for* $i \in \{1, 2\}$, *then for all* $c_1, c_2 \in [\underline{c}, \overline{c}]$, $c_{3-i} \le h(c_1, c_2)$*;*

and analogously for a merger of buyers.

Proof of Lemma F.1. The proof of the lemma proceeds along similar lines as the proof of Proposition 5. That is, it starts with the mechanism that is optimal post merger and constructs a mechanism that replicates the allocations and payments for the nonmerging firms pre merger and shows that it generates weakly more revenue from the merging firms. Because that mechanism is typically not optimal, it then follows that the optimal mechanism pre merger must generate at least as much weighted welfare as the optimal post-merger mechanism.

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

$$\tilde{Q}^S_j(\mathbf{v}, \mathbf{c}) \equiv \hat{Q}^S_j(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}}) \quad \text{and} \quad \tilde{Q}^B_i(\mathbf{v}, \mathbf{c}) \equiv \hat{Q}^B_i(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}}).$$

By condition (i) of the lemma, the interim expected allocations of the nonmerging agents are the same under $\hat{\mathbf{Q}}$ and $\tilde{\mathbf{Q}}$, and so their expected thresholds payments are the same as well.

For supplier 1, define the allocation rule by

$$\text{(F.1)} \quad \tilde{Q}^S_1(\mathbf{v}, \mathbf{c}) \equiv \min\left\{k^S_1, \hat{Q}^S_{1,2}(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}})\right\} \cdot \mathbf{1}_{c_1 \le c_2} + \max\left\{0, \hat{Q}^S_{1,2}(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}}) - k^S_2\right\} \cdot \mathbf{1}_{c_1 > c_2},$$

firms have a bargaining weight of zero, then no further optimization is possible, and so the merger has no effect on expected weighted welfare.

and for supplier 2, define

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

which implies that

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

This allocation rule is monotone—the assumptions on h and the monotonicity of $\tilde{\mathbf{Q}}$ imply that $\tilde{Q}_1^S$ is nonincreasing in c_1 and $\tilde{Q}_2^S$ is nonincreasing in c_2.

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

$$q_1^* \equiv \tilde{Q}_1^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) > 0$$

denote the quantity traded by supplier 1 under $\tilde{\mathbf{Q}}$. Thus, supplier 1 has q_1^* threshold types under $\tilde{\mathbf{Q}}$. For $q \in \{1,\dots,q_1^*\}$, let x_q^* be supplier 1's q-th lowest threshold type unit under $\tilde{\mathbf{Q}}$:

$$x_q^*(\mathbf{v},\mathbf{c}_{-1}) \equiv \sup\{x \in [\underline{c},\overline{c}] \mid \tilde{Q}_1^S(\mathbf{v}, x, c_2, \mathbf{c}_{-\{1,2\}}) \geq q_1^* + 1 - q\}, \tag{F.3}$$

which implies that

$$x_1^*(\mathbf{v},\mathbf{c}_{-1}) \leq x_2^*(\mathbf{v},\mathbf{c}_{-1}) \leq \cdots \leq x_{q_1^*}^*(\mathbf{v},\mathbf{c}_{-1}).$$

By (F.2), the merged entity's quantity $q_{1,2}^* \equiv \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1,c_2), \mathbf{c}_{\{1,2\}})$ is such that $q_{1,2}^* \geq q_1^*$. We can define $q_{1,2}^*$ threshold types for the merged entity. That is, for $q \in \{1,\dots,q_{1,2}^*\}$, we can define z_q^* to be the merged entity's q-th lowest threshold type under $\hat{\mathbf{Q}}$:

$$z_q^*(\mathbf{v},\mathbf{c}_{-\{1,2\}}) \equiv \sup\{z \in [\underline{c},\overline{c}] \mid \hat{Q}_{1,2}^S(\mathbf{v}, z, \mathbf{c}_{-\{1,2\}}) \geq q_{1,2}^* + 1 - q\},$$

where

$$z_1^*(\mathbf{v},\mathbf{c}_{-\{1,2\}}) \leq \cdots \leq z_{q_1^*}^*(\mathbf{v},\mathbf{c}_{-\{1,2\}}) \leq \cdots \leq z_{q_{1,2}^*}^*(\mathbf{v},\mathbf{c}_{-\{1,2\}}).$$

We now show that for all $q \in \{1,\dots,q_1^*\}$, $x_q^*(\mathbf{v},\mathbf{c}_{-1}) \leq z_q^*(\mathbf{v},\mathbf{c}_{-\{1,2\}})$. This implies that the sum of the threshold types of supplier 1 under $\tilde{\mathbf{Q}}$ are less than or equal to the sum of the merged entity's smallest q_1^* threshold types under $\hat{\mathbf{Q}}$:

$$\sum_{q=1}^{q_1^*} x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq \sum_{q=1}^{q_1^*} z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}}).$$

Because the analogous result holds for supplier 2, it follows that the sum of both suppliers' threshold types under $\tilde{\mathbf{Q}}$ is less than or equal to the sum of all of the merged entity's threshold types under $\hat{\mathbf{Q}}$. That is, letting q_2^* be supplier 2's quantity under $\tilde{\mathbf{Q}}$ and $y_q^*(\mathbf{v}, \mathbf{c}_{-2})$ denote its q-th threshold type, we have

$$\sum_{q=1}^{q_1^*} x_q^*(\mathbf{v}, \mathbf{c}_{-1}) + \sum_{q=1}^{q_2^*} y_q^*(\mathbf{v}, \mathbf{c}_{-2}) \leq \sum_{q=1}^{q_{1,2}^*} z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}}).$$

To show that for all $q \in \{1, \ldots, q_1^*\}$, $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}})$, we proceed by contradiction. That is, given $q \in \{1, \ldots, q_1^*\}$, suppose to the contrary that $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) > z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}})$. Then by the definition of $z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}})$, we have

$$\hat{Q}_{1,2}^S(\mathbf{v}, x_q^*(\mathbf{v}, \mathbf{c}_{-1}), \mathbf{c}_{-\{1,2\}}) < q_{1,2}^* + 1 - q. \tag{F.4}$$

We next show that for all $x < x_q^*(\mathbf{v}, \mathbf{c}_{-1})$,

$$\hat{Q}_{1,2}^S(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}) \geq q_{1,2}^* + 1 - q \tag{F.5}$$

and that

$$x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq h(x_q^*(\mathbf{v}, \mathbf{c}_{-1}), c_2) \tag{F.6}$$

hold. To do so, we distinguish between the following three possible cases:

Case 1: $x_1^*(\mathbf{v}, \mathbf{c}_{-1}) \leq x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq c_2$ with at least one inequality strict.

As an intermediate step, we show that $x_1^*(\mathbf{v}, \mathbf{c}_{-1}) < c_2$, which holds by the hypothesis of this case, implies that $q_{1,2}^* = q_1^*$. To see this, note that using $q_1^* = \tilde{Q}_1^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}})$ and the definition of $x_1^*(\mathbf{v}, \mathbf{c}_{-1})$, we have $c_1 \leq x_1^*(\mathbf{v}, \mathbf{c}_{-1})$. Let $\varepsilon \in (0, c_2 - x_1^*(\mathbf{v}, \mathbf{c}_{-1}))$. It follows from the definition of $x_1^*(\mathbf{v}, \mathbf{c}_{-1})$ that $\tilde{Q}_1^S(\mathbf{v}, x_1^*(\mathbf{v}, \mathbf{c}_{-1}) + \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}) < q_1^* \leq k_1^S$, which implies that for supplier 2, we have $\tilde{Q}_2^S(\mathbf{v}, x_1^*(\mathbf{v}, \mathbf{c}_{-1}) + \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}) = 0$. Using the monotonicity of $\tilde{Q}_2^S$ in c_1, it follows that $\tilde{Q}_2^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) = 0$. Thus,

$$q_{1,2}^* = \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{\{1,2\}}) = \tilde{Q}_1^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) + \tilde{Q}_2^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) = q_1^*.$$

Using the hypothesis of this case that $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq c_2$ and the definition of $\tilde{Q}_1^S$ in (F.1), we have

$$\text{(F.7)}\quad x_q^*(\mathbf{v}, \mathbf{c}_{-1}) = \sup\left\{x \in [\underline{c}, \overline{c}] \mid \min\left\{k_1^S, \hat{Q}_{1,2}^S(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}})\right\} \geq q_1^* + 1 - q\right\},$$

which implies that for all $x < x_q^*(\mathbf{v}, \mathbf{c}_{-1})$,

$$\hat{Q}_{1,2}^S(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}) \geq q_1^* + 1 - q = q_{1,2}^* + 1 - q,$$

where the final equality uses $q_{1,2}^* = q_1^*$. Hence, (F.5) holds. In addition, because by the hypothesis of this case, we have $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq c_2$, it follows by condition (ii) of the lemma that (F.6) holds.

Case 2: $x_1^*(\mathbf{v}, \mathbf{c}_{-1}) = x_q^*(\mathbf{v}, \mathbf{c}_{-1}) = c_2$.

As mentioned in Case 1, using $q_1^* = \tilde{Q}_1^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}})$ and the definition of $x_1^*(\mathbf{v}, \mathbf{c}_{-1})$, we have $c_1 \leq x_1^*(\mathbf{v}, \mathbf{c}_{-1})$. Thus, ignoring the possibility of a tie between c_1 and c_2 (which is a zero probability event), the hypothesis of this case implies that $c_1 < c_2$. In addition, the hypothesis of this case implies that for all $q' \in \{1, \ldots, q\}$,

$$x_{q'}^*(\mathbf{v}, \mathbf{c}_{-1}) = c_2.$$

This implies that for $\varepsilon \in (0, \min\{c_2 - c_1, \hat{\varepsilon}\})$ with $\hat{\varepsilon} > 0$ sufficiently small,

$$q_1^* = \tilde{Q}_1^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) = \tilde{Q}_1^S(\mathbf{v}, c_2 - \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}) > \tilde{Q}_1^S(\mathbf{v}, c_2 + \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}).$$

That is, as supplier 1's cost moves from just below c_2 to just above c_2, supplier 1's quantity under $\tilde{Q}_1^S$ is reduced from q_1^*. Changes in supplier 1's type, while still remaining below c_2, that do not affect supplier 1's quantity, also do not affect supplier 2's quantity, so we have

$$q_2^* = \tilde{Q}_2^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) = \tilde{Q}_2^S(\mathbf{v}, c_2 - \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}).$$

Thus,

$$\begin{aligned}
q_{1,2}^* &= \hat{Q}_{1,2}^S(\mathbf{v}, h(c_1, c_2), \mathbf{c}_{-\{1,2\}}) \\
&= \tilde{Q}_1^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) + \tilde{Q}_2^S(\mathbf{v}, c_1, c_2, \mathbf{c}_{-\{1,2\}}) \\
&= \tilde{Q}_1^S(\mathbf{v}, c_2 - \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}) + \tilde{Q}_2^S(\mathbf{v}, c_2 - \varepsilon, c_2, \mathbf{c}_{-\{1,2\}}) \\
&= \hat{Q}_{1,2}^S(\mathbf{v}, h(c_2 - \varepsilon, c_2), \mathbf{c}_{-\{1,2\}}) \\
&= \hat{Q}_{1,2}^S(\mathbf{v}, h(x_q^*(\mathbf{v}, \mathbf{c}_{-1}) - \varepsilon, c_2), \mathbf{c}_{-\{1,2\}}),
\end{aligned}$$

which implies that for all $x < x_q^*(\mathbf{v}, \mathbf{c}_{-1})$,

$$\hat{Q}^S_{1,2}(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}) \geq q^*_{1,2}, \tag{F.8}$$

and hence (F.5) holds. In addition, because by the hypothesis of this case we have $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) = c_2$, it follows by condition (ii) of the lemma that (F.6) holds.

Case 3: $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) > c_2$. In this case, using the definition of $\tilde{Q}^S_1$ in (F.1), we have

$$x_q^*(\mathbf{v}, \mathbf{c}_{-1}) = \sup\left\{x \in [\underline{c}, \overline{c}] \mid \max\left\{0, \hat{Q}^S_{1,2}(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}) - k^S_2\right\} \geq q^*_1 + 1 - q\right\}. \tag{F.9}$$

Further, given that supplier 1 trades $q^*_1 + 1 - q > 0$ units when its cost is (arbitrarily close to) $x_q^*(\mathbf{v}, \mathbf{c}_{-1})$, even though that cost is greater than c_2, it follows that $\hat{Q}^S_{1,2}(\mathbf{v}, h(x_q^*(\mathbf{v}, \mathbf{c}_{-1}), c_2), \mathbf{c}_{-\{1,2\}}) - k^S_2 \geq q^*_1 + 1 - q > 0$. Thus, we can write (F.9) as

$$x_q^*(\mathbf{v}, \mathbf{c}_{-1}) = \sup\left\{x \in [\underline{c}, \overline{c}] \mid \hat{Q}^S_{1,2}(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}) - k^S_2 \geq q^*_1 + 1 - q\right\},$$

which implies that for all $x < x_q^*(\mathbf{v}, \mathbf{c}_{-1})$, $\hat{Q}^S_{1,2}(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}) \geq q^*_1 + 1 - q + k^S_2 \geq q^*_{1,2} + 1 - q$, where the final inequality uses $q^*_1 + k^S_2 \geq q^*_1 + q^*_2 = q^*_{1,2}$. Hence, (F.5) holds. In addition, because under $\tilde{\mathbf{Q}}$ supplier 1 trades the positive quantity q when its cost is $x_q^*(\mathbf{v}, \mathbf{c}_{-1})$ and supplier 2's cost is $c_2 < x_q^*(\mathbf{v}, \mathbf{c}_{-1})$, it must be that $k^S_2 < k^S_{1,2}$. It then follows by condition (iii) of the lemma that (F.6) holds.

We have now shown that, under the supposition that $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) > z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}})$, (F.5) and (F.6) hold. We can combine (F.4) and (F.5), to get for all $x < x_q^*(\mathbf{v}, \mathbf{c}_{-1})$,

$$\hat{Q}^S_{1,2}(\mathbf{v}, x_q^*(\mathbf{v}, \mathbf{c}_{-1}), \mathbf{c}_{-\{1,2\}}) < \hat{Q}^S_{1,2}(\mathbf{v}, h(x, c_2), \mathbf{c}_{-\{1,2\}}),$$

which, since $\hat{Q}^S_{1,2}$ is nonincreasing in the merged entity's type, implies that

$$x_q^*(\mathbf{v}, \mathbf{c}_{-1}) > h(x_q^*(\mathbf{v}, \mathbf{c}_{-1}), c_2),$$

which contradicts (F.6), thereby allowing us to conclude that $x_q^*(\mathbf{v}, \mathbf{c}_{-1}) \leq z_q^*(\mathbf{v}, \mathbf{c}_{-\{1,2\}})$.

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

Under the conditions of Lemma F.1, one can construct a pre-merger mechanism that mimics the post-merger allocation and payments, but with weakly lower payments to merging suppliers (weakly higher payments to merging buyers), resulting in an incentive compatible, individually rational, no-deficit mechanism for the pre-merger market that has the same or greater expected weighted welfare. Optimizing that mechanism for the pre-merger market only reinforces the result that expected weighted welfare is greater pre merger than post merger.

2. Vertical integration

A. MAXIMUM GAIN FROM VERTICAL INTEGRATION. — As noted, for $\underline{v} \geq \overline{c}$ and equal bargaining weights pre-integration, vertical integration cannot possibly increase social surplus. This suggests that, as the overlap of supports becomes smaller, the social surplus gains from vertical integration may decrease as well.

To formalize and substantiate this notion, fix $[\underline{c}, \overline{c}] = [0,1]$ and $\overline{v} = 1$, and define the *maximum gain from vertical integration* associated with $\underline{v} \in [0,1)$ by $\mathcal{G}(\underline{v}) \equiv (W^{FB}(\underline{v}) - W^{SB}(\underline{v}))/W^{FB}(\underline{v})$, where $W^{FB}(\underline{v})$ and $W^{SB}(\underline{v})$ denote first-best and second-best social surplus, respectively, as a function of $\underline{v}$. The amount $\mathcal{G}(\underline{v})$ provides only an upper bound for the gain from vertical integration because vertical integration does not necessarily make the first-best possible when it is not possible absent vertical integration. Then we have:[12]

PROPOSITION F.2: *Assuming $n^B = 1$ and symmetric suppliers, $\mathcal{G}(\underline{v})$ decreases in $\underline{v}$ whenever $\mathcal{G}(\underline{v}) > 0$.*

Proof of Proposition F.2. If $\underline{v} = 0$, then the first-best is not possible without vertical integration; for $n^B = n^S = 1$, see Myerson and Satterthwaite (1983), and for $n^B = 1 < n^S$, see Williams (1999), whose results imply that a necessary condition for first-best to be possible is that $\underline{c}$ is strictly smaller than the lower bound of the support of the buyer's distribution. Because the suppliers are assumed symmetric, we drop the supplier subscripts on the distribution G and on the unweighted and weighted virtual cost functions Γ and Γ^a. Let $L_{n^S}(c) \equiv 1-(1-G(c))^{n^S}$ with density $l_{n^S}(c)$ denote the distribution of the lowest cost draw of a seller. Because we assume 1 buyer, we drop the buyer subscript on the distribution F and on the unweighted and weighted virtual value functions Φ and Φ^a.

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

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

[12]Not surprisingly, a result analogous to Proposition F.2 obtains for the case of one single-unit supplier if one fixes the buyers' support at $[0,1]$ and $\underline{c} = 0$ and varies $\overline{c} \in (0,1]$. In that case, the maximum gain from vertical integration is decreasing in $\overline{c}$.

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

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

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

Analogous to how we define $\pi^{\mathbf{w}}$ in (7), given $\underline{v}$ and ρ, define $\pi(\underline{v}, \rho)$ to be the budget surplus of the mechanism with the allocation rule that solves (6) for $\mathbf{w} = \mathbf{1}$, not including fixed payments:

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

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

$$\frac{\partial \rho_{\underline{v}}}{\partial \underline{v}} \leq 0, \tag{F.10}$$

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

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

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

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

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

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

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

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

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

which is the desired result.

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

$$\frac{\partial\Gamma^{1/\rho^{-1}}(\Phi^{1/\rho}(v))}{\partial\rho} < 0. \tag{F.11}$$

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

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

and

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

Thus, for $\underline{v}\in(\Phi^{1/\rho_0}{}^{-1}(0),1)$, the sign of $\mathcal{G}'(\underline{v})$ is equal to the sign of $h(\underline{v})k'(\underline{v}) - h'(\underline{v})k(\underline{v})$. Using $0 < h(\underline{v}) < k(\underline{v})$, a sufficient condition for $\mathcal{G}'(\underline{v}) < 0$ is that

$k'(\underline{v}) < h'(\underline{v})$. To see that this holds, first note that if $\rho_{\underline{v}} > 1$, then for all $v \in [\underline{v}, 1]$,

$$\Gamma^{1/\rho_{\underline{v}}^{-1}}(\Phi^{1/\rho_{\underline{v}}}(v)) < v, \tag{F.12}$$

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

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

where the first inequality uses $\rho_{\underline{v}} > 1$ and (F.12), and the second inequality uses $\rho_{\underline{v}} > 1$, (F.10), (F.11), and (F.12), completing the proof that for all $\underline{v} \in (\Phi^{1/\rho_0^{-1}}(0), 1)$ such that $\rho_{\underline{v}} > 1$, $\mathcal{G}'(\underline{v}) < 0$. ■

Proposition F.2 provides a monotonicity result relating differences in supports to the maximum gain from vertical integration. Reduced overlap of the supports reduces the maximum gain from vertical integration. Intuitively, the social benefit from vertical integration is reduced when gains from trade are more certain because then market-based transactions between nonintegrated firms work better.

B. COMPARATIVE STATICS FOR VERTICAL INTEGRATION. — In this section, we begin by considering the possibility, captured by Proposition 6 that with overlapping supports, the social surplus effects of vertical integration depend, in general, on the number of firms. Specifically, consider the case of one buyer and multiple suppliers, each of which draws it cost from the same distribution. We know from Williams (1999) that the first-best is possible if $\overline{v} > \overline{c}$ and n^S is large enough.[13] Because vertical integration induces the buyer's willingness to pay to be $y = \min\{v, c\}$, the support of y is $[\min\{\underline{v}, \underline{c}\}, \overline{c}]$. The results of Williams (1999) for this case imply that the first-best is not possible. Hence, vertical integration is socially harmful whenever n^S and the supports are such that the first-best is possible without vertical integration. When $\underline{c} = \underline{v}$ and $\overline{c} = \overline{v}$, the first-best is not possible absent vertical integration, nor with vertical integration if $n^S > 1$ (see, e.g., Williams, 1999).

So the question arises whether vertical integration is more or less likely to be socially harmful if n^S is larger. The intuition and insights from oligopoly mod-

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

els, discussed in the body of the paper, together with the point that the double markup under linear pricing in those models corresponds to having to pay information rents to both the buyer and the suppliers in incomplete information bargaining, suggests that vertical integration is more likely to be harmful the larger is n^S. However, establishing this in general is challenging and beyond the scope of the present paper. That said, Figure F.1(a) illustrates a case in which the oligopoly intuition carries over to our setup. In particular, in a one-to-many market with overlapping supports, vertical integration eliminates a double-markup (of information rents), but it also makes the outside market less competitive, and possibly less efficient, and affects the virtual type function of the integrated firm. As the number of nonintegrated suppliers grows large, the probability that the vertically integrated firm sources internally goes to zero, and the outside market is close to efficient (see Appendix F.2.c for details). Because all effects become small, it is hard to prove general results analytically. As shown in Figure F.1(a), for the case of uniformly distributed types on $[0, 1]$, the change in social surplus due to vertical integration is nonmonotone in the number of outside suppliers and, in the limit, approaches zero from below.

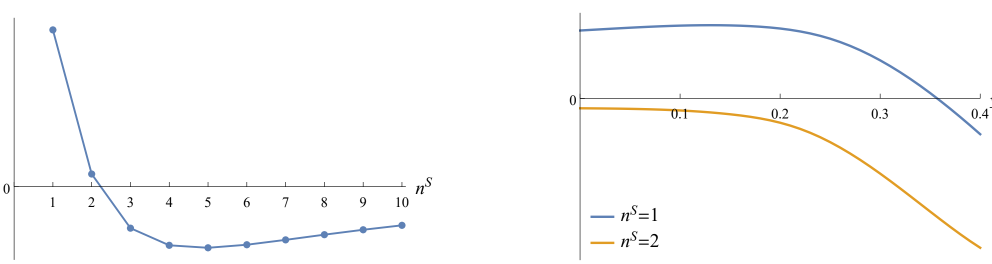

(a) Social surplus effect of vertical integration given $[\underline{c}, \overline{c}] = [\underline{v}, \overline{v}]$ as a function of n^S (illustration of Proposition 6)

(b) Social surplus effect of vertical integration given $[\underline{c}, \overline{c}] = [0, 1]$ and $\overline{v} = 1.2$ as a function of $\underline{v}$ (illustration of Proposition 7)

Figure F.1. Change in expected social surplus as a result of vertical integration.

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

We provide comparative statics related to Proposition 7 in Figure F.1(b), where we illustrate that given $\overline{c} < \overline{v}$, as $\underline{v}$ increases, eventually vertical integration decreases expected social surplus (again, see Appendix 2.c for details). While Proposition 7 is straightforward to prove by taking the case of $\underline{v} = \overline{c}$, Figure F.1 provides examples in which having as few as two independent suppliers in the post-integration market is sufficient for vertical integration to reduce social surplus even when $\underline{v} = \underline{c}$. This emphasizes the salient possibility of anticompetitive vertical integration in a variety of settings.

Interesting and challenging (and still open) issues arise with vertical integration in many-to-many settings. As discussed in and around footnote 22, in that case, the integrated firm may be a buyer or a supplier vis-à-vis the outside firms in the post-integration market, or not trade at all.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Social surplus under vertical integration given n independent suppliers, denoted

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

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

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

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

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

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

3. *Investment effects of vertical integration*

Using Proposition 8, we can connect investment with vertical integration. We assume that vertical integration does not affect the cost of investment for the integrated firm, so if buyer i and supplier j integrate and invest e_i^B in the integrated buyer's distribution and e_j^S in the integrated supplier's distribution, the cost of investment is $\Psi_i^B(e_i^B) + \Psi_j^S(e_j^S)$. With one buyer and one supplier in the pre-integration market and overlapping supports, incomplete information bargaining is inefficient, which under conditions (10) and (11), implies 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, say, one buyer and two or more symmetric suppliers and nonoverlapping supports, incomplete information bargaining is efficient for some bargaining weights, including symmetric ones, without vertical integration, which implies that investments are efficient without vertical integration. But following vertical integration, incomplete information bargaining is inefficient, and so, under (10) and (11), 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 F.1: *Assuming that* (10) *and* (11) *hold, for a one-to-one pre-integration market with overlapping supports, vertical integration promotes efficient investment; but for a one-to-many pre-integration market with nonoverlapping supports, if at least one of* (ii)–(iv) *as stated following* (11) *in the body of the paper holds, then vertical integration disrupts efficient investment if bargaining is efficient prior to vertical integration (which occurs, for example, with symmetric bargaining weights).*

References for the online appendix

Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere. 2002. "Bargaining with Incomplete Information." In *Handbook of Game Theory.* Vol. 3, , ed. R.J. Aumann and S. Hart, 1897–1945. Elsevier Science B.V.

Backus, Matthew, Thomas Blake, Bradley Larsen, and Steven Tadelis. 2020. "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions." *Quarterly Journal of Economics*, 135: 1319–1361.

Boyd, Stephen, and Lieven Vandenberghe. 2004. *Convex Optimization.* Cambridge University Press.

Chatterjee, Kalyan, and William Samuelson. 1983. "Bargaining under Incomplete Information." *Operations Research*, 31: 835–851.

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

Crémer, Jacques, and Michael H. Riordan. 1985. "A Sequential Solution to the Public Goods Problem." *Econometrica*, 53(1): 77–84.

Delacrétaz, David, Simon Loertscher, Leslie M. Marx, and Tom Wilkening. 2019. "Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade." *Journal of Economic Theory*, 179: 416–454.

Krishna, Vijay. 2010. *Auction Theory.* . Second ed., Elsevier.

Larsen, Bradley J. 2021. "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions." *Review of Economic Studies*, 88(2): 851–882.

Loertscher, Simon, and Andras Niedermayer. 2019. "An Optimal Pricing Theory of Transaction Fees in Thin Markets." Working Paper, University of Melbourne.

Makowski, L, and C Mezzetti. 1993. "The Possibility of Efficient Mechanisms for Trading and Indivisible Object." *Journal of Economic Theory*, 59(2): 451–465.

McAfee, R. Preston, and John McMillan. 1987. "Auctions and Bidding." *Journal of Economic Literature*, 25(2): 699–738.

Myerson, Roger B., and Mark A. Satterthwaite. 1983. "Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory*, 29(2): 265–281.

Nash, John F. 1950. "The Bargaining Problem." *Econometrica*, 18(2): 155–162.

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

U.S. Department of Justice and the Federal Trade Commission. 2010. "Horizontal Merger Guidelines." https://www.justice.gov/atr/horizontal-merger-guidelines-08192010.

Williams, Steven R. 1987. "Efficient Performance in Two Agent Bargaining." *Journal of Economic Theory*, 41(1): 154–172.

Williams, Steven R. 1999. "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms." *Economic Theory*, 14: 155–180.