A Appendix: Mechanism design foundations

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

A.1 Concepts and derivations

For ease of exposition, in this appendix we assume that $n^B = n^S = 1$. The extension to multiple buyers and/or suppliers is straightforward.

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

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

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

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

Key constraints

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

$$\hat{u}^B(v) \equiv \hat{q}^B(v)v - \hat{m}^B(v) \geq \hat{q}^B(z)v - \hat{m}^B(z), \tag{13}$$

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

$$\hat{u}^S(c) \equiv \hat{m}^S(c) - \hat{q}^S(c)c \geq \hat{m}^S(z) - \hat{q}^S(z)c. \tag{14}$$

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

$$\mathbb{E}_{v,c}\left[M^B(v,c) - M^S(v,c)\right] \geq 0.$$

Interim expected payoffs

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

Focusing on the buyers, incentive compatibility implies that

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

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

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

and for $\varepsilon < 0$ implies

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

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

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

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

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

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

Mechanism budget surplus

Using the definitions of $\hat{u}^B$ and $\hat{u}^S$ in (13) and (14), we can rewrite these as

$$\hat{m}^B(v) = \hat{q}^B(v)v - \int_{\underline{v}}^{v} \hat{q}^B(t)dt - \hat{u}^B(\underline{v}) \tag{15}$$

and

$$\hat{m}^S(c) = \hat{q}^S(c)c + \int_{c}^{\overline{c}} \hat{q}^S(t)dt + \hat{u}^S(\overline{c}). \tag{16}$$

The expected payment by the buyer is then

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

where the first equality uses the definition of the expectation, the second uses (15), the third switches the order of integration, the fourth integrates, the fifth collects terms, the sixth uses the definition of the virtual value Φ, and the last equality uses the definition of the expectation. Similarly, using (16), the expected payment to the supplier is

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

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

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

As mentioned, it is straightforward to extend these results to the case of multiple buyers and suppliers.

A.2 Myerson-Satterthwaite impossibility result

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

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

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

B Appendix: Proofs

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

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

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

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

We first show that the expected payoffs from $\langle\mathbf{Q}^{\mathbf{w}},\mathbf{M}^{\mathbf{w},\boldsymbol{\eta}}\rangle$ are Pareto undominated among the expected payoffs for any mechanism in $\mathcal{M}$. Proceeding by contradiction, suppose that $\mathbf{u}(\mathbf{w},\boldsymbol{\eta})$ is Pareto dominated by expected payoff vector $\tilde{\mathbf{u}}$ associated with a mechanism $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle\in\mathcal{M}$, i.e., $\tilde{u}_i^B\geq u_i^B(\mathbf{w},\boldsymbol{\eta})$ for all $i\in\mathcal{N}^B$ and $\tilde{u}_i^S\geq u_i^S(\mathbf{w},\boldsymbol{\eta})$ for all $i\in\mathcal{N}^S$, with a strict inequality for at least one agent.

If there exists $i\in\mathcal{N}^B\cup\mathcal{N}^S$ and $j\in\{B,S\}$ such that $\tilde{u}_i^j>u_i^j(\mathbf{w},\boldsymbol{\eta})$ and $w_i^j>0$, then

$$\sum_{i\in\mathcal{N}^B} w_i^B\tilde{u}_i^B+\sum_{i\in\mathcal{N}^S} w_i^S\tilde{u}_i^S>\sum_{i\in\mathcal{N}^B} w_i^B u_i^B(\mathbf{w},\boldsymbol{\eta})+\sum_{i\in\mathcal{N}^S} w_i^S u_i^S(\mathbf{w},\boldsymbol{\eta}), \tag{17}$$

which contradicts $\langle\mathbf{Q}^{\mathbf{w}},\mathbf{M}^{\mathbf{w},\boldsymbol{\eta}}\rangle$ maximizing expected weighted welfare over mechanisms in $\mathcal{M}$. So, for all $i\in\mathcal{N}^B\cup\mathcal{N}^S$ and $j\in\{B,S\}$ such that $\tilde{u}_i^j>u_i^j(\mathbf{w},\boldsymbol{\eta})$, we must have $w_i^j=0$. It follows that (17) holds with equality, which says that $\langle\tilde{\mathbf{Q}},\tilde{\mathbf{M}}\rangle$ maximizes expected weighted welfare. By the uniqueness of the allocation rule identified in Lemma 1, we have $\tilde{\mathbf{Q}} = \mathbf{Q}^{\mathbf{w}}$. Thus, using the payoff equivalence theorem, for any $i \in \mathcal{N}^B \cup \mathcal{N}^S$ and $j \in \{B, S\}$ such that $\tilde{u}_i^j > u_i^j(\mathbf{w}, \boldsymbol{\eta})$, the difference $\tilde{u}_i^j - u_i^j(\mathbf{w}, \boldsymbol{\eta})$ reflects an increase in the interim expected payment to that agent's worst-off type. It follows that there exists a mechanism in $\mathcal{M}$ that has allocation rule $\mathbf{Q}^{\mathbf{w}}$ and payment rule based on $\tilde{\mathbf{M}}$, but that redirects that agent's fixed payment to an agent with positive bargaining weight, that has greater expected weighted welfare than $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$, which is a contradiction. Thus, we conclude that there is no mechanism in $\mathcal{M}$ that Pareto dominates $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$, completing the first part of the proof.

We now turn to the proof that any Pareto undominated payoff vector can be achieved using $\langle \mathbf{Q}^{\mathbf{w}}, \mathbf{M}^{\mathbf{w},\boldsymbol{\eta}} \rangle$ with appropriately chosen $\mathbf{w}$ and $\boldsymbol{\eta}$. Let $\tilde{\mathbf{u}}$ be a payoff vector that is Pareto undominated among expected payoff vectors for mechanisms in $\mathcal{M}$, and let $\langle \tilde{\mathbf{Q}}, \tilde{\mathbf{M}} \rangle \in \mathcal{M}$ be a mechanism that generates payoff vector $\tilde{\mathbf{u}}$.

By the assumption of Pareto undominatedness, $\tilde{u}_i^B$ solves

$$\max_{\langle \mathbf{Q}, \mathbf{M} \rangle \in \mathcal{M}} u_i^B \text{ s.t. for all } j \in \mathcal{N}^B \backslash \{i\},\ u_j^B \geq \tilde{u}_j^B, \text{ and for all } j \in \mathcal{N}^S,\ u_j^S \geq \tilde{u}_j^S \tag{18}$$

and $\tilde{u}_i^S$ solves

$$\max_{\langle \mathbf{Q}, \mathbf{M} \rangle \in \mathcal{M}} u_i^S \text{ s.t. for all } j \in \mathcal{N}^B,\ u_j^B \geq \tilde{u}_j^B, \text{ and for all } j \in \mathcal{N}^S \backslash \{i\},\ u_j^S \geq \tilde{u}_j^S. \tag{19}$$

To ensure the applicability of the Karush-Kuhn-Tucker Theorem, in what follows, if there exists a buyer ℓ whose payoff $\tilde{u}_i^B$ is equal to that buyer ℓ's maximum payoff for all mechanisms in $\mathcal{M}$ then we focus on the problem (18) for buyer ℓ, and, if not, then if there exists a seller ℓ whose payoff $\tilde{u}_i^S$ is equal to that seller ℓ's maximum payoff for all mechanisms in $\mathcal{M}$ then we focus on the problem (19) for seller ℓ, and otherwise we arbitrarily let $\ell = 1$ and focus on the problem (18) for buyer ℓ. This guarantees that it is feasible for the inequality constraints to hold strictly and so Karush-Kuhn-Tucker conditions are necessary and sufficient for an optimum. (This choice of ℓ ensures that we are not focusing on the problem for an agent ℓ where $\tilde{\mathbf{u}}$ can only be obtained through incomplete information bargaining by giving agent ℓ a bargaining weight of zero. Because we assume that at least one agent has a positive bargaining weight, a choice of ℓ that avoids this issue is always possible.)

In what follows, we assume that agent ℓ is a buyer. Analogous arguments apply if agent ℓ is a seller.

Using incentive compatibility and individual rationality, the problem (18) for buyer ℓ can be recast as choosing $\mathbf{Q}$, $\hat{u}_j^B(\underline{v}) \geq 0$ for all $j \in \mathcal{N}^B$, and $\hat{u}_j^S(\overline{c}) \geq 0$ for all $j \in \mathcal{N}^S$. Noting that buyer j's expected payoff is the expectation of $u_j^B(\mathbf{v}, \mathbf{c}) \equiv (v_j - \Phi_j(v_j)) Q_j^B(\mathbf{v}, \mathbf{c}) + \hat{u}_j^B(\underline{v})$

and seller j's expected payoff is the expectation of $u_j^S(\mathbf{v},\mathbf{c}) \equiv (\Gamma_j(c_j)-c_j)Q_j^S(v,\mathbf{c}) + \hat{u}_j^S(\overline{c})$, and letting $\rho \geq 0$ be the multiplier on the no-deficit constraint and $\mu_j^x \geq 0$ be the multiplier on the constraint that $u_j^x \geq \tilde{u}_j^x$, the associated Lagrangian is

$$\begin{aligned}\mathcal{L} \;\equiv\; & \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[u_\ell^B(\mathbf{v},\mathbf{c}) + \sum_{j\in\mathcal{N}^B\backslash\{\ell\}} \mu_j^B(u_j^B(\mathbf{v},\mathbf{c}) - \tilde{u}_j^B) + \sum_{j\in\mathcal{N}^S} \mu_j^S\left(u_j^S(\mathbf{v},\mathbf{c}) - \tilde{u}_j^S\right)\right] \\ & +\rho\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{j\in\mathcal{N}^B}\left(\Phi_j(v_j)Q_j^B(\mathbf{v},\mathbf{c}) - \hat{u}_j^B(\underline{v})\right) - \sum_{j\in\mathcal{N}^S}\left(\Gamma_j(c_j)Q_j^S(\mathbf{v},\mathbf{c}) - \hat{u}_j^S(\overline{c})\right)\right],\end{aligned}$$

which we can rewrite, up to terms that do not involve the allocation rule, as

$$\begin{aligned}& \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\left(v_\ell - \Phi_\ell(v_\ell) + \rho\Phi_\ell(v_\ell)\right)Q_\ell^B(\mathbf{v},\mathbf{c}) + \sum_{j\in\mathcal{N}^B\backslash\{\ell\}}\left(\mu_j^B(v_j - \Phi_j(v_j)) + \rho\Phi_j(v_j)\right)Q_j^B(\mathbf{v},\mathbf{c})\right] \\ & +\mathbb{E}_{\mathbf{v},\mathbf{c}}\left[\sum_{j\in\mathcal{N}^S}\left(\mu_j^S(\Gamma_j(c_j) - c_j) - \rho\Gamma_j(c_j)\right)Q_j^S(\mathbf{v},\mathbf{c})\right].\end{aligned}$$

Let $\tilde{\mathbf{Q}}$ denote the solution value for the allocation rule. Denote the expected budget surplus under $\tilde{\mathbf{Q}}$ not including payments to worst-off types by

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

which is nonnegative because any solution satisfies the no-deficit constraint. Then the solution values for the payments to the worst-off types are given by, for all $j \in \mathcal{N}^B\backslash\{\ell\}$,

$$\hat{u}_j^B(\underline{v}) = \tilde{u}_j^B - \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[(v_j - \Phi_j(v_j))\tilde{Q}_j^B(\mathbf{v},\mathbf{c})\right],$$

and for all $j \in \mathcal{N}^S$,

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

and using the Pareto optimality of $\tilde{\mathbf{Q}}$, which implies that the mechanism must pay out all of $\tilde{R}$ to the agents,

$$\hat{u}_\ell^B(\underline{v}) = \tilde{R} - \sum_{j\in\mathcal{N}^B/\{\ell\}}\hat{u}_j^B(\underline{v}) - \sum_{j\in\mathcal{N}^S}\hat{u}_j^S(\overline{c}).$$

Given the solution values for the multipliers $\boldsymbol{\mu}$, let $\overline{w} \equiv \max\{\mu_1^B, ..., \mu_{\ell-1}^B, 1, \mu_{\ell+1}^B, ..., \mu_m^B, \mu_1^S, \ldots, \mu_n^S\}$ and define $\tilde{\mathbf{w}} \equiv (\frac{\mu_1^B}{\overline{w}}, ..., \frac{\mu_{\ell-1}^B}{\overline{w}}, \frac{1}{\overline{w}}, \frac{\mu_{\ell+1}^B}{\overline{w}}, ..., \frac{\mu_m^B}{\overline{w}}, \frac{\mu_1^S}{\overline{w}}, \ldots, \frac{\mu_n^S}{\overline{w}})$. Recall that $\mathbf{Q}^{\mathbf{w}}$ has associated Lagrangian, up to terms that do not involve the allocation rule, given by

$$\mathcal{L}^{\mathbf{w}} \equiv \mathbb{E}_{\mathbf{v},\mathbf{c}} \left[\sum_{j\in\mathcal{N}^B} \left(w_j^B \left(v_j - \Phi_j(v_j)\right) + \rho\Phi_j(v_j)\right) Q_j^B(\mathbf{v},\mathbf{c}) + \sum_{j\in\mathcal{N}^S} \left(w_j^S(\Gamma_j(c_j) - c_j) - \rho\Gamma_j(c_j)\right) Q_j^S(\mathbf{v},\mathbf{c}) \right].$$

It follows that $\mathcal{L}$ is equal to $\overline{w}\mathcal{L}^{\tilde{\mathbf{w}}}$ plus terms that do not involve $\mathbf{Q}$, and so $\tilde{\mathbf{Q}}$ is the same allocation rule as $\mathbf{Q}^{\tilde{\mathbf{w}}}$. Further, the payoffs $\tilde{\mathbf{u}}$ are replicated by $\langle \mathbf{Q}^{\tilde{\mathbf{w}}}, \mathbf{M}^{\tilde{\mathbf{w}},\tilde{\boldsymbol{\eta}}}\rangle$ with any specification of $\tilde{\boldsymbol{\eta}}$ if $\tilde{R}=0$ and otherwise with for $j\in\mathcal{N}^S$,

$$\tilde{\eta}_j^S = \begin{cases} \frac{1}{\tilde{R}}\left(\tilde{u}_j^S - \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[(\Gamma_j(c_j) - c_j)\, \tilde{Q}_j^S(\mathbf{v},\mathbf{c})\right]\right) & \text{if } \mu_j^S = \overline{w}, \\ 0 & \text{otherwise,} \end{cases} \tag{20}$$

for $j \in \mathcal{N}^B \backslash \{\ell\}$,

$$\tilde{\eta}_j^B = \begin{cases} \frac{1}{\tilde{R}}\left(\tilde{u}_j^B - \mathbb{E}_{\mathbf{v},\mathbf{c}}\left[(v_j - \Phi_j(v_j))\, \tilde{Q}_j^B(\mathbf{v},\mathbf{c})\right]\right) & \text{if } \mu_j^B = \overline{w}, \\ 0 & \text{otherwise,} \end{cases} \tag{21}$$

and $\eta_\ell^B = 1 - \sum_{j\in\mathcal{N}^B\backslash\{\ell\}} \eta_j^B - \sum_{j\in\mathcal{N}^S}\eta_j^S$, which completes the proof. ■

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

Case 1. Suppose that $K^B \le K^S$ and (7) fails to hold. (Analogous arguments apply if $K^B > K^S$ and (7) 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 Proposition 4. To begin, note that under the assumption that all buyers have bargaining weight w^B and all suppliers have bargaining weight w^S, the allocation rule $\mathbf{Q}^{\mathbf{w}}$ depends only on the bargaining differential $\Delta \equiv \frac{w^B - w^S}{\max\{w^B, w^S\}}$. Given Δ, we denote the associated point on the frontier as $(u_S(\Delta), u_B(\Delta))$, where u_S is the sum of all suppliers' expected payoffs and u_B is the sum of all buyers' expected payoffs.

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

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

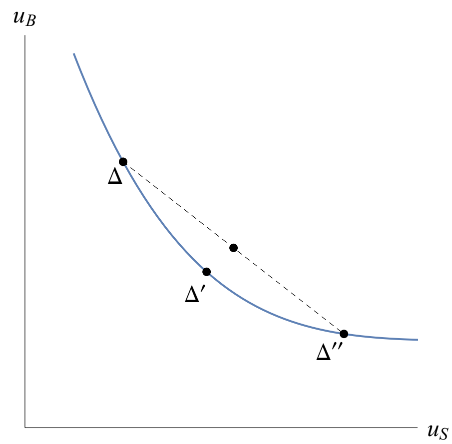

Then there exist points on the frontier, which we denote by their associated bargaining differentials Δ, Δ', and Δ'', with $\Delta > \Delta' > \Delta''$, such that $(u_S(\Delta) + u_S(\Delta''))/2 > u_S(\Delta')$ and $(u_B(\Delta)+u_B(\Delta''))/2 > u_B(\Delta')$. Let $\mu(\Delta)$ denote the incomplete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

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

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

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

We begin with a brief preamble to establish some notation and useful relations. Let $\hat{u}_{i,\mathbf{Q}}^B(v_i; \mathbf{e}_{-i}^B, \mathbf{e}^S)$ 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 $(\mathbf{e}_{-i}^B, \mathbf{e}^S)$. Define $\hat{u}_{i,\mathbf{Q}}^S(c_i; \mathbf{e}^B, \mathbf{e}_{-i}^S)$ analogously. Let $u_{i,\mathbf{Q}}^B(\mathbf{e})$ and $u_{i,\mathbf{Q}}^S(\mathbf{e})$ 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 $q_i^B(v_i; \mathbf{e}_{-i}^B, \mathbf{e}^S) \equiv \mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}|\mathbf{e}_{-i}^B,\mathbf{e}^S}\left[Q_i^B(\mathbf{v}, \mathbf{c})\right]$ and $q_i(c_i; \mathbf{e}^B, \mathbf{e}_{-i}^S) \equiv$

$\mathbb{E}_{\mathbf{v},\mathbf{c}_{-i}|\mathbf{e}^B,\mathbf{e}^S_{-i}}\left[Q_i^S(\mathbf{v},\mathbf{c})\right]$. As discussed in Appendix A.1, by the payoff equivalence theorem, we have, up to a constant,

$$\hat{u}^B_{i,\mathbf{Q}}(v_i;\mathbf{e}^B_{-i},\mathbf{e}^S) = \int_{\underline{v}}^{v_i} q_i^B(x;\mathbf{e}^B_{-i},\mathbf{e}^S)dx, \tag{22}$$

and, taking expectations with respect to v_i, one obtains

$$u^B_{i,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{v}}^{\overline{v}} q_i^B(x;\mathbf{e}^B_{-i},\mathbf{e}^S)(1-F(x;e_i^B))dx \tag{23}$$

up to a constant, and, analogously,

$$u^S_{i,\mathbf{Q}}(\mathbf{e}) = \int_{\underline{c}}^{\overline{c}} q_i^S(x;\mathbf{e}^B,\mathbf{e}^S_{-i})G_i(x;e_i^S)dx \tag{24}$$

up to a constant.

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

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

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

$$\overline{e}_i^B \in \arg\max_{e_i^B} u^B_{i,\mathbf{Q}^{FB}}(e_i^B,\overline{\mathbf{e}}^B_{-i},\overline{\mathbf{e}}^S) - \Psi_i^B(e_i^B) \tag{25}$$

and

$$\overline{e}_j^S \in \arg\max_{e_j^S} u^S_{j,\mathbf{Q}^{FB}}(\overline{\mathbf{e}}^B,e_j^S,\overline{\mathbf{e}}^S_{-j}) - \Psi_j^S(e_j^S). \tag{26}$$

Assume that (9)–(10) hold. Let $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ denote the incomplete information bargaining allocation rule given in Lemma 1, but with the virtual types defined in terms of the type distributions associated with investment $\overline{\mathbf{e}}$, and let $\rho^{\mathbf{w}}_{\overline{\mathbf{e}}}$ 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 $j\in\mathcal{N}^S$,

$$\overline{e}_i^B \in \arg\max_{e_i^B} u^B_{i,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}(e_i^B,\overline{\mathbf{e}}^B_{-i},\overline{\mathbf{e}}^S) - \Psi_i^B(e_i^B) \tag{27}$$

and

$$\overline{e}_j^S \in \arg\max_{e_j^S} u^S_{j,\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}}(\overline{\mathbf{e}}^B,e_j^S,\overline{\mathbf{e}}^S_{-j}) - \Psi_j^S(e_j^S). \tag{28}$$

Assumptions (9)–(10) ensure that the first-best investments are characterized by their first-order conditions. Thus, using (23) and (25), we have for all $i \in \mathcal{N}^B$,

$$-\int_{\underline{v}}^{\overline{v}} q_i^{FB,B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx - \Psi_i^{B\prime}(\overline{e}_i^B) = 0. \tag{29}$$

Similarly, using (23) and (27), we have

$$-\int_{\underline{v}}^{\overline{v}} q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx - \Psi_i^{B\prime}(\overline{e}_i^B) = 0. \tag{30}$$

Combining (29) and (30), we have

$$\int_{\underline{v}}^{\overline{v}} (q_i^{FB,B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S) - q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x; \overline{\mathbf{e}}_{-i}^B, \overline{\mathbf{e}}^S)) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx = 0. \tag{31}$$

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

$$\mathbb{E}_{\mathbf{v}_{-i},\mathbf{c}|\overline{\mathbf{e}}_{-i}^B,\overline{\mathbf{e}}^S} \left[\int_{\underline{v}}^{\overline{v}} (Q_i^{FB,B}(x, \mathbf{v}_{-i}, \mathbf{c}) - Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(x, \mathbf{v}_{-i}, \mathbf{c})) \frac{\partial F_i(x; \overline{e}_i^B)}{\partial e} dx \right] = 0. \tag{32}$$

Steps analogous to those leading to (31) imply that for all $i \in \mathcal{N}^S$,

$$\int_{\underline{c}}^{\overline{c}} (q_i^{FB,S}(x; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-i}^S) - q_i^{\mathbf{w},\overline{\mathbf{e}},S}(x; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-i}^S)) \frac{\partial G_i(x; \overline{e}_i^S)}{\partial e} dx = 0. \tag{33}$$

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 $v \in (\underline{v}, \overline{v})$, (32) then implies that

$$\sum_{i \in \mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v}, \mathbf{c}) = \sum_{i \in \mathcal{N}^B} Q_i^{FB,B}(\mathbf{v}, \mathbf{c}) \equiv \xi(\mathbf{v}, \mathbf{c}) \tag{34}$$

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. It 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 (34). 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 supplier 1 trades in the first-best but not under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$, while supplier 2 trades under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ but not under the first-best, i.e., $Q_1^{FB,S}(\hat{\mathbf{v}}, \hat{\mathbf{c}}) > 0$, $Q_1^{\mathbf{w},\overline{\mathbf{e}},S}(\hat{\mathbf{v}}, \hat{\mathbf{c}}) = 0$, $Q_2^{\mathbf{w},\overline{\mathbf{e}},S}(\hat{\mathbf{v}}, \hat{\mathbf{c}}) > 0$, and $Q_2^{FB,S}(\hat{\mathbf{v}}, \hat{\mathbf{c}}) = 0$. This implies that $\hat{c}_1$ is among the $\xi(\hat{\mathbf{v}}, \hat{\mathbf{c}})$ lowest elements of $\hat{\mathbf{c}}$, but $\hat{c}_2$ is not, and that $\Gamma_2^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_2; \overline{e}_2^S)$ is among the $\xi(\hat{\mathbf{v}}, \hat{\mathbf{c}})$ lowest elements of $\mathbf{\Gamma}^{\mathbf{w}^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{\mathbf{c}}, \overline{\mathbf{e}}^S)$, but $\Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_1; \overline{e}_1^S)$ is not. It follows that

$$\hat{c}_1 < \hat{c}_2 \leq \Gamma_2^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_2; \overline{e}_2^S) \leq \Gamma_1^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_1; \overline{e}_1^S) \text{ and } \Gamma_2^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{c}_2; \overline{e}_2^S) \leq \mathbf{\Gamma}^{\mathbf{w}^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\hat{\mathbf{c}}, \overline{\mathbf{e}}^S)_{[\xi(\hat{\mathbf{v}}, \hat{\mathbf{c}})]}.$$

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

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

which implies that no trades occur under $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ and exactly 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 $(\tilde{\mathbf{v}}, \tilde{\mathbf{c}})$, $\sum_{i \in \mathcal{N}^B} Q_i^{\mathbf{w},\overline{\mathbf{e}},B}(\mathbf{v}, \mathbf{c}) \neq \sum_{i \in \mathcal{N}^B} Q_i^{FB,B}(\mathbf{v}, \mathbf{c})$, which contradicts (34). 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 (34) 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 (11) holds. (Analogous arguments apply to the case with $K^B > K^S$ and (12).) Then all buyers consume their full demands under the first-best. If $w_1^S = \cdots = w_{n^S}^S$, then (11) 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 (34), $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ induces the same suppliers to produce as does $\mathbf{Q}^{FB}$, and again we are done.

So, suppose that $K^B < K^S$, (11) holds, and the suppliers do not all have the same bargaining weight. Let the suppliers be numbered such that $w_1^S = \min_{i\in\mathcal{N}^S} w_i^S < w_2^S$. Using (11), we drop the agent subscript on the virtual types and the investment argument in the virtual types. We have, for all $c \in (\underline{c}, \overline{c})$,

$$\Gamma^{w_2^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c) < \max_{i\in\mathcal{N}^S} \Gamma^{w_i^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c) = \Gamma^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c), \tag{35}$$

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

$$\begin{aligned}
&q_1^{\mathbf{w},\overline{\mathbf{e}},S}(c_1; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-1}^S)\\
&= \mathbb{E}_{\mathbf{v},\mathbf{c}_{-1}|\overline{\mathbf{e}}^B,\overline{\mathbf{e}}_{-1}^S}[\text{repetitions of } \Gamma^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c_1) \text{ in the } \xi(\mathbf{v},\mathbf{c})\text{-smallest elements of } \mathbf{\Gamma}^{\mathbf{w}^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\mathbf{c}, \overline{\mathbf{e}}^S)]\\
&= \mathbb{E}_{\mathbf{v},\mathbf{c}_{-1}|\overline{\mathbf{e}}^B,\overline{\mathbf{e}}_{-1}^S}[\text{repetitions of } c_1 \text{ in the } \xi(\mathbf{v},\mathbf{c})\text{-smallest elements of } \Gamma^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}-1}}(\mathbf{\Gamma}^{\mathbf{w}^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(\mathbf{c}, \overline{\mathbf{e}}^S))]\\
&< \mathbb{E}_{\mathbf{v},\mathbf{c}_{-1}|\overline{\mathbf{e}}^B,\overline{\mathbf{e}}_{-1}^S}[\text{repetitions of } c_1 \text{ in the } \xi(\mathbf{v},\mathbf{c})\text{-smallest elements of } \mathbf{c}]\\
&= q_1^{FB,S}(c_1; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-1}^S),
\end{aligned}$$

where the first equality uses the definition of $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}}$ and (34), the second equality uses the assumption that $\Gamma^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}-1}}$ is increasing, the inequality follows from (35), which implies that for all $i \in \mathcal{N}^S$,

$$\Gamma^{w_1^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}-1}}(\Gamma^{w_i^S/\rho_{\overline{\mathbf{e}}}^{\mathbf{w}}}(c_i)) \leq c_i,$$

and the final equality uses the definition of $\mathbf{Q}^{FB}$ and (34). This result that $q_1^{\mathbf{w},\overline{\mathbf{e}},S}(c_1; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-1}^S) < q_1^{FB,S}(c_1; \overline{\mathbf{e}}^B, \overline{\mathbf{e}}_{-1}^S)$ for all $c \in (\underline{c}, \overline{c})$ contradicts (33), and so we conclude that all suppliers must have the same bargaining weight, and so $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}} = \mathbf{Q}^{FB}$. Analogously, if $K^B > K^S$ and (12) holds, then we can use (31) to show that all buyers must have the same bargaining weight and so again $\mathbf{Q}^{\mathbf{w},\overline{\mathbf{e}}} = \mathbf{Q}^{FB}$, completing the proof. ■

Proof of Proposition 7. With nonoverlapping supports and symmetric bargaining weights, the pre-integration market achieves the first-best. Suppose that $n^B = 1$ and $n^S \geq 2$. After integration between the buyer and supplier 1, the buyer's willingness to pay is the cost realization of the integrated supplier, that is, c_1, 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. An analogous argument applies for the case of $n^S = 1$ and $n^B \geq 2$. ■

C Appendix: Extensions

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

C.1 Heterogeneous outside options

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

Fixed costs of participating in the mechanism

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

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 $x_i \geq 0$ for supplier i. These outside options are best thought of as fixed costs of participating in the mechanism because they have to be borne regardless of whether an agent trades. In this case, the price-formation mechanism with weights $\mathbf{w}$ is the solution to

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

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

$$\mathbb{E}_{v,\mathbf{c}}\left[\sum_{i\in\mathcal{N}} \left(\Phi(v) - \Gamma_i(c_i)\right) \cdot \mathbf{1}_{\Phi^{w^B/\rho}(v) \geq \Gamma_i^{w_i/\rho}(c_i) = \min_{j\in\mathcal{N}} \Gamma_j^{w_j/\rho}(c_j)}\right] \geq x_B + \sum_{i\in\mathcal{N}} x_i, \tag{36}$$

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

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

Production-relevant outside options

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

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

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

$$\Gamma_{i,a}^o(c) \equiv c + (1 - a)\frac{G_i(c - y_i)}{g_i(c - y_i)} = \Gamma_{i,a}(c - y) + y < \Gamma_{i,a}(c), \tag{37}$$

where the inequality holds for all $a < 1$ because the monotone hazard rate assumption implies that $\Gamma_{i,a}'(c) > 1$ for all $a < 1$. This in turn has two, somewhat subtle implications. Let z be the threshold for supplier i to trade when its outside option is zero, i.e., keeping z fixed, i trades if and only if $\Gamma_{i,a}(c) \leq z$. (Note that z will be the minimum of the buyer's weighted virtual value and the smallest weighted virtual cost of i's competitors, but this does not matter for the argument that follows.) Assuming that $a < 1$ and $y_i < \bar{c} - \underline{c}$, which implies that $\underline{c}_i < \bar{c}$, it follows that there are costs $c \in [\underline{c}_i, \bar{c}]$ and thresholds z such that supplier i trades when it has the outside option and not without it, that is, $\Gamma^o_{i,a}(c) < z < \Gamma_{i,a}(c)$. This reflects the reasonably well-known result that optimal mechanisms tend to discriminate in favor of weaker agents (McAfee and McMillan, 1987), which in this case is the agent with the positive outside option. It also resonates with intuition from complete information models: keeping costs fixed, the agent with the better outside option is treated more favorably, indeed, it is evaluated according to a smaller weighted virtual cost. However, from an ex ante perspective, the larger is the value of the outside option, the less likely is the agent to trade. To see this, consider a fixed realization of z. (The distribution of these thresholds is not be affected by i's outside option and hence our argument extends directly once one integrates over z and its density.) Given y_i, supplier i trades if and only if its cost c is below $\tau(y)$ satisfying $\Gamma^o_{i,a}(\tau(y)) = z$. Using (37), this is equivalent to $\Gamma_{i,a}(\tau(y) - y) + y = z$, which in turn is equivalent to $\tau(y) = \Gamma^{-1}_{i,a}(z - y) + y$, whose derivative for $a < 1$ satisfies

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

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

C.2 Preferences over suppliers and bargaining externalities

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

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

We define the virtual surplus $\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}$ associated with trade between the buyer and supplier i, accounting for the agents' bargaining weights $\mathbf{w}$ and the buyer's preferences $\boldsymbol{\theta}$, with $\rho^{\mathbf{w},\boldsymbol{\theta}}$ defined analogously to before as $\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}(v, c_i) \equiv \theta_i \Phi^{w^B/\rho^{\mathbf{w},\boldsymbol{\theta}}}(v) - \Gamma_i^{w_i^S/\rho^{\mathbf{w},\boldsymbol{\theta}}}(c_i)$. Let $\boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv (\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}(v, c_i))_{i \in \mathcal{N}^S}$ 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 2. *Assuming that $n^B = 1$ and $n^S \geq 2$, in the generalized setup with buyer preferences $\boldsymbol{\theta}$, incomplete information bargaining with weights $\mathbf{w}$ has the allocation rule for $i \in \mathcal{N}^S$, $Q_i^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv 1$ if $\Lambda_i^{\mathbf{w},\boldsymbol{\theta}}(v, c_i) \geq \max\{0, \boldsymbol{\Lambda}^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c})_{(K^B)}\}$, and otherwise $Q_i^{\mathbf{w},\boldsymbol{\theta}}(v, \mathbf{c}) \equiv 0$.*

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

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

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

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 i trades, $\Pr(\theta_i \Phi^{w^B/\rho^{\mathbf{w},\boldsymbol{\theta}}}(v) \geq \Gamma_i^{w_i^S/\rho^{\mathbf{w},\boldsymbol{\theta}}}(c_i))$, 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 9. *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 9 does not necessarily extend to the case with $K^B < n$, as shown in the following example.

Example with bargaining externalities

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

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

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

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

C.3 Generalization

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

Let $\mathcal{P}$ be the set of subsets of $\mathcal{N}^S$ with no more than K^B elements (including the empty set) and let $\boldsymbol{\theta} = \{\theta_X\}_{X\in\mathcal{P}}$ be a commonly known vector of taste parameters of the buyer satisfying the "size-dependent discounts" condition of Delacrétaz et al. (2019). Specifically, let there be supplier-specific preferences $\{\hat{\theta}_i\}_{i\in\mathcal{N}}$ and size-dependent discounts $\{\delta_i\}_{i\in\mathcal{N}}$ with $0 = \delta_0 = \delta_1 \leq \delta_2 \leq \cdots \leq \delta_n$ such that for all $X \in \mathcal{P}$, $\theta_X = \sum_{i\in X} \hat{\theta}_i - \delta_{|X|}$. Thus, the buyer's value for purchasing from suppliers in $X \in \mathcal{P}$ when its type 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 $\hat{\theta}_i = \theta$ for some common θ and for $i \in \mathcal{N}$, δ_i either all zero for constant marginal value or increasing in i for decreasing marginal value; (ii) differentiated products by letting $\hat{\theta}_i$ differ by i and setting all δ_i to zero; (iii) a one-buyer version of the Shapley-Shubik model by setting $K^B = 1$; 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 ironed weighted virtual value, scaled by $\theta_{X^*_\rho(v,\mathbf{c})}$, and the ironed 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 i is

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

D Appendix: Foundations

In Section D.1 we show that our model encompasses the k-double auction as a special case. In Section D.2, we provide an axiomatic approach to incomplete information bargaining. In Section D.3, we provide an extensive-form game that, at least under some assumptions, has the incomplete information bargaining outcome as a Nash equilibrium outcome.

D.1 Appendix: k-double auction as a special case

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

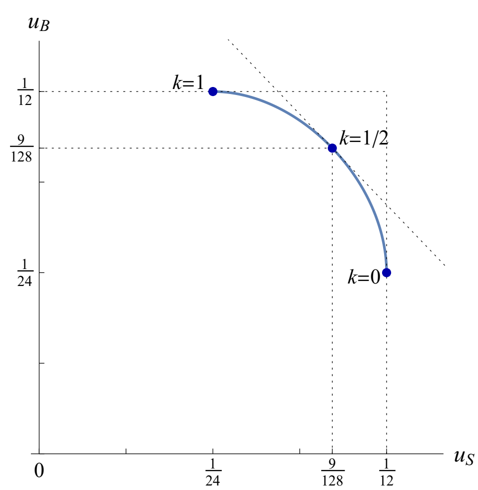

Figure D.1: Payoffs in the k-double auction for all $k \in [0,1]$. Assumes that there is one supplier and that the buyer's value and the supplier's cost are uniformly distributed on $[0,1]$.

To see that incomplete information bargaining encompasses the k-double auction as a

[2]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.

special case, note that for the case of one supplier, the allocation $Q_i^{\mathbf{w}}(v, \mathbf{c})$ is the same for all $\mathbf{w}$ with the same bargaining differential Δ defined by

$$\Delta \equiv \frac{w^B - w_1^S}{\max\{w^B, w_1^S\}} \in [-1, 1].$$

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

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

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

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

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

Making the substitution $\Delta = w^B - w^S$ and writing $\rho^{\mathbf{w}}$ as a function of Δ, we have $\rho^\Delta = 1$ for $\Delta \in \{-1, 1\}$ and otherwise ρ^Δ is given by:

$$\rho^\Delta = \frac{3(1 - |\Delta|)}{4 - 2|\Delta| - 2\sqrt{1 - |\Delta| + \Delta^2}}. \tag{38}$$

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

$$\Delta_k \equiv \frac{1 - 2k}{k^2 - \max\{1, 2k\}}. \tag{39}$$

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 $Q_1^{(1+\Delta,1)}(v, c)$ if $\Delta \in [-1, 0]$ and for $Q_1^{(1,1-\Delta)}(v, c)$ if $\Delta \in (0, 1]$ to get

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

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

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

D.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 a case in point, Ausubel et al. (2002) note that asking for efficient outcomes in bargaining is "fruitless," given the impossibility theorem of Myerson and Satterthwaite (1983).

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

Observe that the payoff equivalence theorem is *distribution free* (or detail free) insofar as it holds for any distributions $F_1, \ldots, F_{n^B}$ 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 assume regularity (i.e., that virtual value and cost functions are increasing) in order to avoid the technicalities of ironing. We do the same here, although all results continue to hold without regularity assumptions when the weighted and unweighted virtual value and cost functions are replaced by their ironed counterparts.

The first three axioms ensure that the incomplete information bargaining mechanism is *feasible*, which means that beyond satisfying resource constraints, the mechanism satisfies incentive compatibility, individual rationality, and does not run a deficit.[3]

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 buyer and suppliers symmetrically.

Axiom 4: Efficiency: The mechanism maximizes expected social surplus subject to the

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

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 introduced above with $p = (\underline{v} + \overline{c})/2$. Notice the similarity to Nash bargaining here—the posted price is the same price at which a buyer with value $\underline{v}$ and a supplier with cost $\overline{c}$ would trade under Nash's axioms and assumptions.

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

We now state our characterization result.

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

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

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

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

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

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

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

D.3 Implementation

In many cases, economists have achieved greater comfort with models of price-formation processes when the literature has shown that there exists a noncooperative game that, at least under some assumptions, has an equilibrium outcome that is the same as the outcome delivered by the model under consideration. Indeed, this comfort often extends well beyond the narrow confines of the foundational game. For example, the existence of microfoundations are regularly invoked to support empirical estimation of a model even when the data-generation process does not conform to the extensive-form game providing the microfoundation.[4] As

[4]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.

another case in point, to support the model of perfectly competitive markets, one might view price-taking buyers and suppliers as submitting demand and supply schedules to a (fictitious) Walrasian auctioneer who then sets market clearing prices. Similarly, in the Cournot model, one might view firms as submitting quantities to an auctioneer or market maker who sets the market clearing price.[5] Under assumptions on the alternation of offers and taking the limit as the time between offers goes to zero, Rubinstein bargaining delivers Nash bargaining outcomes (Rubinstein, 1982); and under additional assumptions, including conditions on firms' marginal contributions and passive beliefs, the limit of an alternating-offers game approximates Nash-in-Nash outcomes (Collard-Wexler et al., 2019).

In light of this, it is perhaps useful to note that, as mentioned above and discussed in Appendix D.1, for the case of one supplier 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 now, provide micro-foundations for incomplete information bargaining. Specifically, for general distributions and any number of suppliers, the incomplete information bargaining outcome arises in equilibrium in an extensive-form game involving a buyer, suppliers, and a fee-setting broker.

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

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

intermediary.

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

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

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

Proposition 10. *The outcome of incomplete information bargaining with bargaining weights* $\mathbf{w}$ *and shares* $\boldsymbol{\eta}$ *is a Bayes Nash equilibrium outcome of the fee-setting extensive-form game with intermediary preference weights* $\mathbf{w}$ *and profit shares* $\boldsymbol{\eta}$*.*

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

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

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

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

i wins if and only if

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

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

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

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

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

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

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

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

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

where the final equality uses the change of variables $y = \Gamma_i(c_i)$. Thus, the buyer with value

v maximizes its interim expected payoff by choosing r to solve

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

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

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

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

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

This is reminiscent of the role of intermediaries in the wholesale used car market as described by Larsen (2020). There, auction houses run auctions, facilitate further bargaining in the substantial number of cases in which the auction does not result in trade, and collect fees from traders.

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

References for the online appendix

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

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

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

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

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

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

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

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

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

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

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

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

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