Pollution Permits: Efficiency by Design

Marek Pycia and Kyle Woodward*

July 2022

Abstract

The annual adverse effects of pollution are on the order of 10% of world GDP. Many approaches are used or have been proposed to control the growing pollution problem, but none of them allows for efficient pollution control in settings in which the marginal cost of pollution is increasing and polluters are better informed than regulators about the costs of abatement. We propose a simple primary market mechanism, *True-Cost Pay as Bid*, that implements efficient pollution control and does not depend on how much information the regulators have about the abatement costs.

1 Introduction

Multi-unit auctions are frequently used to allocate emissions permits and other goods. In many such markets, the auctioneer is interested in the social efficiency of the outcome. For example, the auctioneer may want to induce market participants to limit their emissions to the point at which the cost of further emissions reductions (abatement) crosses the social marginal cost of emissions (the marginal damage curve). Less abatement is environmentally costly because pollution is socially undesirable, and more abatement is economically costly because it reduces production.

We propose a simple mechanism that we call True-Cost Pay as Bid (TCPAB) because it is a pay-as-bid auction with elastic supply set equal to the true expected marginal social cost curve and we show that this mechanism guarantees efficiency in this problem, while uniform-price auctions and other standard permit market designs do not despite all of them being frequently employed in environmental contexts. In particular, in environments where

*We would like to thank Peter Cramton, Pär Holmberg, Henrik Horn, David Levine, Axel Ockenfels, Ludvig Sinander, Melinda Suveg, Thomas Tanger, Alex Teytelboym, Jean Tirole, Alexander Wagner, Andrew Yates, and audiences at U Zurich Spring 2020 faculty lunch, Zurich Inaugural Market Design Public Lecture, Zurich Environmental Economics Research Slam, IFN Stockholm, Oxford University, and ESSET 2022.

bidders have common information an optimally-designed pay-as-bid auction will generically outperform any uniform-price auction. Furthermore, pay-as-bid auctions approximately outperform uniform-price auctions even when bidders are asymmetric. The good properties of our proposed mechanism obtain whether there is an active resale market (cap and trade) or not.

We allow polluters to be asymmetric in their cost of abatement; our substantive assumption is that the polluters are close to symmetrically informed. We focus on a seller who cares only about the quantities allocated to various bidders, conditional on the information revealed by bids. Thus, in the benchmark model, the auctioneer does not care about revenue, but can care about allocative efficiency. We relax this assumption in an extension, and show that sellers who are partially interested in revenue still prefer the pay-as-bid auction.

The case for price discovery through auctions, instead of relying on regulators' expectations of what the right price should be, was aptly summarized by *The Economist* (February 27, 2021):

> "Working out costs is tricky. The International Energy Agency (IEA), for instance, has routinely underestimated the pace of deployment of renewables. And because economies of scale drive down prices, that means that it has overestimated the costs of switching, too. In 2010 the lowest the IEA expected solar prices to drop to over the next decade was about $195 per megawatt hour. Today the price in America and Europe is $30-60."

The problem with changing costs while the allocation mechanism remains in place is commonly recognized, cf. Schmalensee et al. (1998) (history of sulphur dioxide in the US) and Blanchard et al. (2021) (central bank for permits).

The mechanism we propose is sufficiently flexible to accommodate innovation policy. Gillingham and Stock (2018) and Gates 2021 observe that a long-term view of abatement calls for discrimination among abatement approaches and de facto subsidizing technologies that will reduce emissions more effectively in the long-term. Such discrimination can be incorporated into the auctions we propose through preferential treatment of policymaker-favored technologies, e.g. by subsidizing them. Such discrimination, however, would need to be accompanied by limitations on secondary market permit trades.

We propose the TCPAB mechanism focusing on pollution types that can be controlled by a single social planner, such as sulphur dioxide or particulate matter. An important property of our mechanism is that it is well-suited to multilateral negotiations and hence to controlling the global carbon pollution. Cramton, MacKay, et al. (2017) show experimentally that multilateral negotiations over a tax are easier than negotiations over other control

mechanisms such as quantity caps because negotiating over taxes allows the negotiating parties to agree on the minimal of all proposed taxes. One could run multilateral negotiations in the context of TCPAB: negotiators would propose elastic supply curves and the outcome of the negotiation would be the pointwise minimum of the elastic supply curves proposed by individual negotiators. We think that this is an advantage of our mechanism over the discretionary central bank type solutions (of course, pollution banks could sell pollution permits using our mechanism, thus inheriting its good multilateral negotiation properties).

The advantages of elastic supply have been recognized in many markets (Cramton and Stoft (2005) electricity; P. Klemperer (2010), toxic assets).[1] In the context of pollution, the potential advantages of elastic supply were recognized before us by Karp and Traeger (2021), who recognized its beneficial properties in the context of perfect Walrasian competition. We go beyond their work by proposing a well-behaved auction mechanism and studying its properties.

We contribute as well to the large literature on multi-unit auctions. Equilibrium existence for our setting was studied in Woodward (2016) and Pycia and Woodward (2022a). The latter paper also studies revenue maximization in multi-unit auctions; while they focus on auctions with inelastic supply, they also establish the essential uniqueness of equilibria of pay as bid with deterministic elastic supply. For earlier work on multi-unit auctions, see e.g. Ausubel et al. (2014).

Finally, we contribute to the literature on trade-offs between price and quantity initiated by Weitzman (1974): we show that by using TCPAB one can achieve the unconstrained optimum, thus eliminating the need for the trade off.

2 Permit Mechanisms around the World

In this paper we show that the True-Cost Pay as Bid auction is efficient and approximately robust to asymmetric information. The True-Cost Pay as Bid auction is a specific implementation of a pay as bid auction, with elastic supply set equal to the true marginal damage curve associated with emissions. We argue in favor of True-Cost Pay as Bid auction in part by showing that other real-world permit mechanisms do not share its desirable qualities. To illustrate these points, we consider standard pollution mechanisms in a simple example of an emissions permit market with two polluters. Our polluters procure emissions permits to offset their emissions from each year's production. When the bidders' common signal is s, a

[1] LiCalzi and Pavan (2005) discuss the role elastic supply might play in eliminating tacit collusion in uniform price; cf. also McAdams (2007), who proposed ex post adjustable supply as a mean to fight collusion. Our pay as bid proposal does not run into tacit collusion problems that might occur in uniform-price auctions.

polluter's value for emissions quantity q is

$$v\left(q;s\right)=\left(s-q\right)_{+}$$

where $(x)_+$ denotes $\max\{x,0\}$. The social cost of emissions is $C(Q;\omega)=Q$, and does not depend on the state of the world. Conditional on polluters' private information, the efficient level of emissions is $Q=2s/3$. This introductory section is mostly self-contained; for formal definitions of e.g. uniform-price and pay-as-bid auctions, see Section 3 that follows.

2.1 Command and Control

Command and control is used to control pollution in many markets, including car emissions in the U.S., Europe, Mexico, and many other countries. In a command and control framework, the regulator sets explicit conditions on emissions: maximum levels, reduction technologies, etc. We consider a regulator who sets an annual emissions quantity Q. Because the marginal value of emissions is weakly positive, we may assume that polluters emit the entire allocated quantity, and therefore this is a straightforward decision problem for the regulator,

$$\max_{Q}\mathbb{E}\left[2\int_{0}^{\frac{1}{2}Q}v\left(q;s\right)dq-\int_{0}^{Q}C\left(Q';\omega\right)dQ'\right].$$

This leaves

$$\max_{Q}\left(\mathbb{E}\left[s\right]+1\right)Q-\frac{3}{4}Q^{2}.$$

Thus optimal supply is

$$Q^{\star}=\frac{2}{3}\left(1+\mathbb{E}\left[s\right]\right). \tag{1}$$

In implementation, command and control has an observed weakness of inflexibility, cf. Weitzman (1974). This can be readily seen in 1. The mechanism's optimal level of emissions depends on bidders' expected information, and there will be an efficiency gap associated with the regulator's incomplete information.

2.2 Taxes

Taxes are used—or scheduled to be used—to control carbon dioxide pollution in many markets including the U.K., France, Canada, Argentina, and many other countries (cf., e.g., Cramton, MacKay, et al. 2017, Blanchard et al. 2021, and Santikarn et al. 2021). Under taxation, the regulator sets a fixed price at which polluters may purchase an arbitrary quantity of emissions permits. Profit maximization implies that if the tax on pollution is $\tau^{\star}$, each

polluter will purchase $q = (s - \tau^\star)_+$ emissions permits. As under command and control, this is a straightforward decision problem for the regulator,

$$\max_{\tau} \mathbb{E}\left[2\int_0^{(s-\tau)_+} v\left(q;s\right) dq - \int_0^{2(s-\tau)_+} C\left(Q';\omega\right) dQ'\right].$$

From this, it is straightforward to see that

$$\mathbb{E}\left[(s-\tau)_+ - \frac{1}{2}s\right] = 0. \tag{2}$$

Implementation of optimal taxation faces the difficulty of determining optimal taxes. In environments which are dynamically variant or statically unstable, a single optimal tax will not be able to enforce efficient market outcomes, cf. Weitzman (1974). This can be readily seen in 2, where the single-dimensional tax must aggregate all expected information about polluters' signals.

2.3 Uniform-Price

Uniform-price auctions are used—or scheduled to be used—to allocate pollution permits (e.g., to control carbon dioxide pollution) in many markets including the U.S., China, Germany, Canada, and many others (Santikarn et al. 2021). While the implemented auctions are mostly inelastic let us consider the more general class of uniform-price auctions allowing elastic supply.[2] Unlike command and control or taxation, elastic supply permits the regulator to implement a mechanism with more than a single parameter: the elastic supply curve corresponds to a continuum of parameters chosen against individual prices. While this can allow the permit efficient allocations to arise, it also introduces ambiguity into the allocation procedure because of the multiplicity of equilibria (c.f., Wilson (1979), P. D. Klemperer and Meyer (1989), and Wang and Zender (2002)). We say that an auction *implements allocation* $\tilde{q}$ if, for all bidding equilibria $b^\star$, $q(b^\star(s)) = \tilde{q}(s)$; we then also write $q \circ b^\star = q^\star$.

Lemma 1. [Uniform-Price Equilibrium Selection] *Fix an optimal allocation correspondence $q^\star$. If aggregate optimal supply $Q^\star$ is nonconstant, then for any supply curve S the uniform-price auction admits bidding equilibria in which $q \circ b^\star \neq q^\star$.*

A few notes are in order. First, we assume that the seller's marginal cost curve is nonconstant. If it were, a fixed reserve price (letting $S(Q) = c$ for all Q) would yield efficient

[2] In practice, supply is typically either fixed or subject to a *price collar*, a fixed supply with the addition of a price floor and a price ceiling.

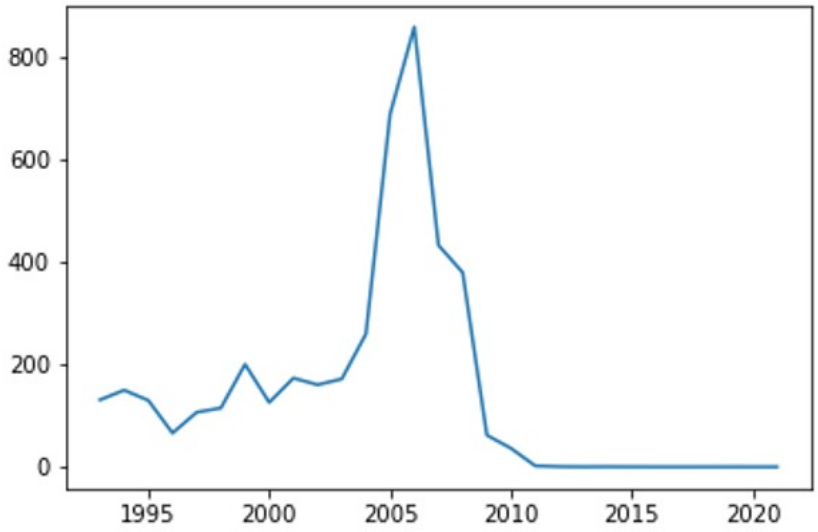


Figure 1: US SO2 spot permit prices in EPA annual auctions (vertical axis: market claring prices in USD; horizontal axis: auction years). Data: US Environmental Protection Agency.

implementation in the uniform-price auction; this would be equivalent to a tax. Thus in settings where the seller's marginal cost curve is roughly constant, optimal taxation will achieve near efficiency. Second, we assume that the support of optimal aggregate supply is convex and nondegenerate. If it were degenerate, fixed supply would yield efficient outcomes in the uniform-price auction. To see the role of convexity, consider an optimal aggregate supply which may take only two values. In this case, provided the two quantities are sufficiently far apart relative to the seller's change in marginal cost (and also taking into consideration bidders' values), the seller can implement efficient outcomes by setting supply equal to a step function with jumps at each of the two relevant quantities. Third, the optimal uniform-price auction depends on the selected bidding equilibrium; cf. the discussion of equilibrium multiplicity above.

2.4 Inelastic Pay as Bid

Pay-as-bid auctions are used to sell e.g. sulphur dioxide permits. They are implemented with fixed and inelastic supply. The resulting pricing problems of this approach can be seen in the prices of U.S. SO2 permits sold by U.S. EPA. The inelastic supply is unable to generate sensible prices when demand for permits substantially changes (for an analysis of the demand change, see, e.g., Schmalensee et al. (1998)).

The nice properties we derive for the True-Cost Pay-as-Bid auction depend both on the

incentives induced by the pay-as-bid mechanism, and on the flexibility afforded by elastic supply. Removing one or the other feature of the mechanism allows suboptimal outcomes to arise again.

3 Model

There are $n \geq 2$ polluters, $i \in I = \{1, ..., n\}$. Polluter i's marginal valuation for quantity q is denoted $v^i(q; s_i)$, where s_i is a signal privately known to polluter i; in particular, we allow the polluters to be ex ante asymmetric. Without loss of generality, we decompose polluter i's signal as $s_i = (s, \theta_i)$, where s is common to all polluters and θ_i is private to polluter i. We assume that each $v^i(\cdot; s_i)$ is strictly decreasing where it is strictly positive, Lipschitz continuous, and almost-everywhere differentiable in q. We allow arbitrary dimensionality of s_i, and an arbitrary integrable $v^i(q; \cdot)$. We focus on environments in which s_i are highly correlated across bidders but the regulator does not know the polluters' information. Our results take the simplest form when the correlation across polluters is perfect, $s_1 = ... = s_n = (s, 0)$, without imposing any further assumptions on the distribution of s. Under perfect correlation, signal s has no strategic importance for bidders participating in an auction, and thus when studying the equilibrium among such bidders we fix s and denote the bidders' marginal valuation by $v^i(q; s_i) = v^i(q)$.[3] We assume that bidders' values are regular in the sense of Pycia and Woodward (2022a).

The regulator's value for a social allocation $(q_i)_{i=1}^n$ is $\pi(q; s, \omega)$, and depends on the level of emissions, the bidders' information s, and an unknown state of the world $\omega \sim F_\omega$. Note that the regulator's value can depend on polluters' common information. This will be the case when, for example, the seller is interested in allocative efficiency. Given polluters' common signal s, we denote the seller's optimal allocation by $q^\star(s) = \arg\max_q \mathbb{E}_\omega[\pi(q; s, \omega)]$, and we assume that each $q^\star(s)$ is single-valued. We allow the seller's value to depend on endogenous equilibrium outcomes, such as the revenue obtained from permit allocation, but for the bulk of our analysis we assume that the seller is interested in efficiency, so that

$$\pi\left(q; s, \omega\right) = \sum_{i=1}^{n} \int_0^{q_i} v^i\left(x; s\right) dx - C\left(\sum_{i=1}^{n} q_i; \omega\right),$$

where $C(Q; \omega)$ is the economic damage associated with emissions Q when the state of the world is ω. Setting $Q^\star\left(s\right) = \sum_{i=1}^n q_i^\star\left(s\right)$, we assume throughout that the set of welfare-

[3]The seller may not know the bidders' information if, for example, the seller needs to commit to the auction mechanism before this information is revealed. Alternatively, the seller may want to fix a single design for multiple auctions.

maximizing aggregate quantities $\{Q^\star(s) : s \in \mathcal{S}\}$ is a non-trivial interval. In particular, this assumption rules out discrete bidder signal spaces. We also assume that the seller's marginal cost curve $c = dC/dQ$ is nonconstant on $\operatorname{Supp} Q^\star$.

3.1 Elastic pay as bid

In an elastic pay-as-bid auction the regulator first announces a supply curve S, then solicits decreasing bid curves $b^i : \mathbb{R}_+ \to \mathbb{R}_+$ from each polluter. After receiving bid curves, the regulator computes the market-clearing price given reported aggregate demand and announced supply,

$$p^\star = \inf \left\{ p \colon \exists q \geq 0 \text{ s.t. } b^i(q_i) \leq p \; \forall i, \text{ and } S\left(\sum_{i \in I} q_i\right) \leq p \right\}.$$

The regulator allocates to each polluter the number of permits they demand at the market-clearing price, $q_i = \varphi^i(p^\star)$, where φ^i is the demand function of polluter i (the inverse of the submitted bid function b^i.[4] In exchange, the polluter transfers to the regulator the entire area under their submitted bid curve, $t_i = \int_0^{q_i} b^i(x)dx$.

4 True-Cost Pay as Bid

True-Cost Pay as Bid is pay as bid with supply equal to the expected marginal social cost curve, $S = \mathbb{E}_\omega [C(\cdot; \omega)]$.

Theorem 1. [Efficiency of True-Cost Pay as Bid] *Given any optimal allocation correspondence $q^\star$, True-Cost Pay as Bid has an essentially unique bidding equilibrium and this equilibrium implements $q^\star$.*

In a deterministic pay-as-bid auction, equilibrium bids are flat, and the resulting allocations equate marginal value with marginal price. By allowing for heterogeneous bid subsidies, the seller can equate marginal value with marginal price, less the heterogeneous subsidy. It follows that any desired allocation is attainable in equilibrium.

Theorem 2. [True-Cost Pay as Bid Is Detail-Free] *True-Cost Pay as Bid does not depend on the distribution of bidders' values.*

[4] Tiebreaking is handled pro-rata on the margin. We assume here that the bid function b^i is locally invertible, but in analysis we handle technical details relating to locally non-invertible (flat) bids.

5 Asymmetric information

When bidders are asymmetrically informed our approach gives us bounds on efficiency.[5] We use Pycia and Woodward (2022a)'s concept of the concept of δ-small informational asymmetry. We assume that the information about others is (i) concentrated (or at least not to much fat-tailed), and (ii) there are sufficiently many polluters bidding in the auction. Our bounds on expected efficiency are then be given in terms of the variance of individual values and the number of (say symmetric) polluters; the bounds are increasing in the variance and (quickly) decreasing in the number of polluters.

Theorem 3. [Efficiency guarantee with small global variance] *For every $\varepsilon > 0$ there is $\delta > 0$ such that if, for each relevant q, the variance of $v(q, s_i)$ is bounded by δ, then the per-permit efficiency loss is bounded by ε.*

If, in the theorem, the slope of the true cost of pollution curve is bounded below by some $\gamma > 0$, then the total efficiency loss is bounded from above by ε. In addition, our proof provides an explicit bound on δ in terms of ε. A key role in the proof is played by our Theorem 5. In particular, we also obtain the following Theorem.

Theorem 4. [Efficiency guarantee with locally small support] *Let $s = (s_i)_{i \in I}$ be a reference signal profile and let $(q_i^\star(s))_{i=\in I}$ be the vector of efficient assignments for this reference signal profile. Suppose that there is $\delta \geq 0$ such that for all signal profiles s' and all bidders i, the permit value $v^i(q_i^\star(s); s_i')$ is within δ of $v^i(q_i^\star(s); s_i)$. Then the ex-post per-permit inefficiency of TCPAB is bounded above by 2δ. Furthermore, if $\gamma > 0$ is a lower bound on the slope of the true cost of pollution, then total welfare loss is bounded by $4\delta^2/\gamma$.*

Corollary 1. [Efficiency guarantee with small support] *If the asymmetry of information is δ-small, then the ex-post inefficiency of TCPAB per permit purchased is bounded by 2δ. Furthermore, if $\gamma > 0$ is a lower bound on the slope of the true cost of pollution, then the total welfare loss is bounded by $\frac{4\delta^2}{\gamma}$.*

6 Other policy dimensions

6.1 Heterogenous Technology and Innovation

The True-Cost Pay as Bid can be adapted to promote innovation by selling different permits for use with different technologies (permits not tradeable across technologies). The price

[5] Pycia and Woodward (2022a) established related bounds on revenue in terms of the size of the support of uncertainty; we provide the bounds for efficiency in terms of variance. In particular, we do not require the support to be bounded.

paid by polluters using a promising technology can then be subsidized.

6.2 Multilateral Bargaining

For local pollutants such as particulate matter or SO2, one regulator is in charge of the design. For global pollutants, such as CO2, a multilateral agreement needs to be reached. We show that our auction can be easily negotiated over.

7 Conclusion

We have proposed a simple auction mechanism to sell pollution permits. The auction—True-Cost Pay as Bid—achieves efficiency without the need for the designer to know the abatement costs. Thanks to this property the mechanism self-adjusts to changing abatement costs. Furthermore the auction works well whether the polluters are symmetric or not.

References

Ausubel, Lawrence M. et al. (2014). "Demand Reduction and Inefficiency in Multi-Unit Auctions". In: *The Review of Economic Studies*.

Blanchard, Olivier et al. (2021). *Major Future Economic Challenges*. Tech. rep.

Cramton, Peter, David JC MacKay, et al. (2017). *Global Carbon Pricing: The Path to Climate Cooperation*. The MIT Press.

Cramton, Peter and Steven Stoft (2005). "A capacity market that makes sense". In: *The Electricity Journal* 18.7, pp. 43–54.

Gates, Bill (2021). *How to avoid a climate disaster: the solutions we have and the breakthroughs we need*. Penguin UK.

Gillingham, Kenneth and James H Stock (2018). "The cost of reducing greenhouse gas emissions". In: *Journal of Economic Perspectives* 32.4, pp. 53–72.

Karp, Larry S. and Christian P. Traeger (2021). *Smart Cap*. Tech. rep. CESifo Working Paper No. 8917.

Klemperer, Paul (2010). "The product-mix auction: A new auction design for differentiated goods". In: *Journal of the European Economic Association* 8.2-3, pp. 526–536.

Klemperer, Paul D and Margaret A Meyer (1989). "Supply function equilibria in oligopoly under uncertainty". In: *Econometrica: Journal of the Econometric Society*, pp. 1243–1277.

LiCalzi, Marco and Alessandro Pavan (2005). "Tilting the supply schedule to enhance competition in uniform-price auctions". In: *European Economic Review* 49.1, pp. 227–250.

McAdams, David (2007). "Adjustable supply in uniform price auctions: Non-commitment as a strategic tool". In: *Economics Letters* 95.1, pp. 48–53.

Pycia, Marek and Kyle Woodward (2022a). "Auctions of Homogeneous Goods: A Case for Pay-as-Bid".

— (2022b). "Pay-as-Bid Auctions with Asymmetric Information".

Santikarn, Marissa et al. (2021). "State and trends of carbon pricing 2021". In.

Schmalensee, Richard et al. (Sept. 1998). "An Interim Evaluation of Sulfur Dioxide Emissions Trading". In: *Journal of Economic Perspectives* 12.3, pp. 53–68. DOI: 10.1257/jep.12.3.53. URL: https://www.aeaweb.org/articles?id=10.1257/jep.12.3.53.

Wang, James JD and Jaime F Zender (2002). "Auctioning divisible goods". In: *Economic Theory* 19.4, pp. 673–705.

Weitzman, Martin L (1974). "Prices vs. quantities". In: *The Review of Economic Studies* 41.4, pp. 477–491.

Wilson, Robert (1979). "Auctions of shares". In: *The Quarterly Journal of Economics*, pp. 675–689.

Woodward, Kyle (2016). *Strategic Ironing in Pay-as-Bid Auctions: Equilibrium Existence with Private Information*. Working paper.

A A Bound on Market Price in Pay as Bid

Our analysis of optimal bidding relies on the following key theorem; in this theorem we impose no restrictions on bidders' information and we allow mixed-strategy equilibria.

Theorem 5. [A Bound on Market Price] *In any mixed-strategy equilibrium of the pay-as-bid auction, for any signal profile* $(s,(\theta_1,\ldots,\theta_n))$ *all realizations of the market clearing price for the realized market quantity* $Q^\star(s,\cdot)$ *are bounded between the smallest and largest marginal value at the largest quantity feasible in equilibrium,*

$$\min_i \operatorname*{ess\,inf}_{\tilde{\theta}_i} v^i\left(\overline{q}^i\left(s,\tilde{\theta}_i\right);s,\tilde{\theta}_i\right) \leq p\left(Q^\star\left(s,\theta\right);s,\theta\right) \leq \max_i \operatorname*{ess\,sup}_{\tilde{\theta}_i} v^i\left(\overline{q}^i\left(s,\tilde{\theta}_i\right);s,\tilde{\theta}_i\right),$$

where $\overline{q}^i\left(s,\tilde{\theta}_i\right)$ *is the maximum quantity purchased by bidder* i *with type* $\left(s,\tilde{\theta}_i\right)$.

The proof of Theorem 5 shows a slightly stronger claim: for any realization of (s,θ), the equilibrium bid for the maximum quantity bidder i, with type $s_i=(s,\theta_i)$, can obtain equals the bidder's marginal value for this quantity; or, bids equal values at the maximum feasible quantity. The intuition for this claim is that if a bidder has strictly positive margin at the maximum feasible quantity, they can slightly increase their bid and obtain a non-negligible

additional quantity at minimally higher price, which is a profitable deviation; the proof of Theorem 5 formalizes this intuition and takes care of technical complications related to tie-breaking, flat bids, and binding monotonicity constraints. Note that this intuition applies only to the maximum quantity at which the increased bid is paid only when it is marginal; at any lower quantity the increased bid would need to be paid also when inframarginal.

Because bids are decreasing in quantity, the equilibrium market-clearing price paid by bidder i is minimized when their realized quantity is maximized, $q^i(s,\theta) = \overline{q}^i(s,\theta)$. Thus, the theorem provides bounds on the minimum market price. In the special case where bidders are symmetric and have only common and no idiosyncratic information, then

$$\min_i \operatorname*{ess\,inf}_{\tilde{\theta}_i} v^i\left(\frac{1}{n}Q^\star(s)\,;s,\tilde{\theta}_i\right) = v\left(\frac{1}{n}Q^\star(s)\,;s\right) = \max_i \operatorname*{ess\,sup}_{\tilde{\theta}_i} v^i\left(\frac{1}{n}Q^\star(s)\,;s,\tilde{\theta}_i\right)$$

That is, the market price at the maximum quantity is exactly equal to each bidder's marginal value at the last unit they receive.

Because bids equal values at bidders' maximum feasible quantities, and these quantities are sold when a bidder's opponents submit weak bids, the equality of market price and a bidder's marginal value obtains only when bidders submit weak bids. For such a bid profile the market-clearing price will be minimized, but the lower bound of Theorem 5 remains valid irrespective of the realization of supply. The market-clearing price at other bid profiles can (and frequently does) rise above $\max_i \operatorname{ess\,sup}_{\tilde{\theta}_i} v^i(\overline{q}^i(s,\tilde{\theta}_i);s,\tilde{\theta}_i)$.

A.1 Proof of Theorem 5 (Bound on Market Price)

Our equilibrium analysis relies on the identification of the minimum equilibrium market clearing price. In this appendix we prove Theorem 5, which bounds this price. The arguments do not depend on the presence (or absence) of idiosyncratic private information or mixed strategies. We consolidate all bidder-known uncertainty into $\zeta_i = (s,\theta_i,\xi_i)$, where s is the signal observed by all bidders, θ_i is bidder i's idiosyncratic private information, and ξ_i is a term parameterizing bidder i's potentially-mixed strategy; thus bidder i's bid $b^i : [0,\overline{Q}] \times \operatorname{Supp}\zeta_i \to \mathbb{R}_+$.[6] Where useful, we consider $\zeta_i|s$ to hold fixed the common signal s while letting θ_i and ξ_i vary.

We also introduce notation for the (essential) minimum market clearing price $\underline{p}$ and

[6]For compactness we also write $v(\cdot;\zeta_i) = v(\cdot;s,\theta_i)$, but we do not imply that a bidder's marginal value may vary with her action selection from a mixed strategy.

(essential) maximum receivable quantity $\overline{q}^i$, conditional on strategy profile $(b^j)_{j=1}^n$,

$$\underline{p}(s) = \operatorname*{ess\,inf}_{\zeta|s} p\left(Q; \left(b^j(\cdot;\zeta_j)\right)_{j=1}^n\right);$$
$$\overline{q}^i(\zeta_i) = \operatorname*{ess\,sup}_{\zeta_{-i}|s} q^i\left(Q; b^i(\cdot;\zeta_i), b^{-i}(\cdot,\zeta_{-i})\right).$$

Thus, when the bidding strategy profile is $(b^j)_{j=1}^n$, the market clearing price is almost never below $\underline{p}(s)$ when the common signal is s, and bidder i's allocation is almost never above $\overline{q}^i(\zeta_i)$ when her type is ζ_i.

Lemma 2. *In any equilibrium, conditional on common signal s, at least $n-1$ bidders, with probability 1, bid their true value for their maximum receivable quantity. That is,*

$$\#\left\{i\colon \Pr\left(b^i\left(\overline{q}^i(\zeta);\zeta\right) = v\left(\overline{q}^i(\zeta);\zeta\right)\middle|s\right) = 1\middle|s\right\} \geq n-1.$$

Proof. For a given agent i, common signal s, and $\lambda > 0$, consider an alternative bidding strategy b^λ defined by

$$b^\lambda(q;\zeta_i) = \begin{cases} b^i(q;\zeta_i) & \text{if } b^i(q;\zeta_i) \geq b^i\left(\overline{q}^i(\zeta_i);\zeta_i\right) + \lambda, \\ \min\left\{b^i\left(\overline{q}^i(\zeta_i);\zeta_i\right) + \lambda, v(q;\zeta_i)\right\} & \text{otherwise.} \end{cases}$$

Since $b^i(\cdot;\zeta_i)$ is left-continuous, for small λ this deviation will award the agent all excess quantity above $\sum_{j\neq i}\varphi^j(b^i(\overline{q}^i(\zeta_i);\zeta_i)+\lambda;\zeta_j)$. Let $q^\star(\lambda;\zeta)$ be the quantity obtained under this deviation when, under the original strategy, $q^i(\zeta)$ units would be obtained. Explicitly,

$$q^\star(\lambda;\zeta) = Q - \sum_{j\neq i}\varphi^i\left(b^i\left(\overline{q}^i(\zeta_i);\zeta_i\right)+\lambda;\zeta_j\right) = Q - \sum_{j\neq i}\underline{q}^{ji}(\lambda;\zeta),$$

where $\underline{q}^{ji}(\lambda;\zeta) = \varphi^j(b^i(\overline{q}^i(\zeta_i);\zeta_i)+\lambda;\zeta_j)$ is the quantity bidder j receives when the aggregate signal profile is ζ and bidder i implements bid b^λ; note that $\underline{q}^{ii}(\lambda;\zeta)$ is the maximum quantity for which bidder i bids above $b^i(\overline{q}(\zeta_i);\zeta_i)+\lambda$, which does not depend on ζ_{-i}, and denote this quantity by $\underline{q}^i_\lambda(\zeta_i)$. We will use the quantity $q^\star(\lambda;\zeta)$ to analyze the additional quantity the deviation yields above baseline,

$$\Delta^i_L(\lambda;\zeta) = q^i(\zeta) - \underline{q}^{ii}(\lambda;\zeta), \quad \Delta^i_R(\lambda;\zeta) = q^\star(\lambda;\zeta) - q^i(\zeta),$$
$$\Delta^i(\lambda;\zeta) = \Delta^i_L(\lambda;\zeta) + \Delta^i_R(\lambda;\zeta).$$

Incentive compatibility requires that this deviation cannot be profitable, hence the additional

costs must outweigh the additional benefits,

$$\begin{aligned}&\mathbb{E}_{\zeta|s}\left[\int_{\underline{q}^i_\lambda(\zeta_i)}^{q^i(\zeta)} b^\lambda\left(x;\zeta_i\right)-b^i\left(x;\zeta_i\right)dx\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right]\\&\quad\geq\mathbb{E}_{\zeta|s}\left[\int_{q^i(\zeta)}^{q^\star(\lambda;\zeta)} v\left(x;\zeta_i\right)-b^\lambda\left(x;\zeta_i\right)dx\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right].\end{aligned}$$

Importantly, this inequality must hold both ex ante and interim, unconditional on θ_i. Because bids are weakly decreasing, the left-hand expectation is bounded above by

$$\begin{aligned}&\mathbb{E}_{\zeta|s}\left[\int_{\underline{q}^i_\lambda(\zeta_i)}^{q^i(\zeta)} b^\lambda\left(x;\zeta_i\right)-b^i\left(x;\zeta_i\right)dx\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta\right)\right]\\&\leq\mathbb{E}_{\zeta|s}\left[\int_{\underline{q}^i_\lambda(\zeta_i)}^{q^i(\zeta)} b^i\left(\overline{q}^i\left(\zeta_i\right);\zeta_i\right)+\lambda-b^i\left(\overline{q}^i\left(\zeta_i\right);\zeta_i\right)\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right]\\&=\lambda\mathbb{E}_{\zeta_{-i}|s}\left[\Delta^i_L\left(\lambda;\zeta\right)\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right].\end{aligned}$$

As marginal values are Lipschitz in quantity and $b^i(\overline{q}^i(\zeta_i);\zeta_i)<v^i(\overline{q}^i(\zeta_i);\zeta_i)$ by assumption, the right-hand expectation is bounded above by (M is the Lipschitz modulus of v)

$$\begin{aligned}&\mathbb{E}_{\zeta|s}\left[\int_{q^i(\zeta)}^{q^\star(\lambda;\zeta)} v\left(x;\zeta_i\right)-b^\lambda\left(x;\zeta_i\right)dx\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right]\\&\geq\mathbb{E}_{\zeta|s}\left[\int_{q^i(\zeta)}^{q^\star(\lambda,\zeta)}\left(v^i\left(\overline{q}^i\left(\zeta_i\right);\zeta_i\right)-\left(x-q^i\left(\zeta\right)\right)M-\left(b^i\left(\overline{q}^i\left(\zeta_i\right);\zeta_i\right)+\lambda\right)\right)_+dx\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right]\\&\geq\mathbb{E}_{\zeta|s}\left[\frac{1}{2}\left(\mu\left(\zeta_i\right)-\lambda\right)\min\left\{\Delta^i_R\left(\lambda;\zeta\right),\frac{\mu\left(\zeta_i\right)-\lambda}{M}\right\}\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right],\end{aligned}$$

where $\mu(\zeta_i)=v^i(\overline{q}^i(\zeta_i);\zeta_i)-b^i(\overline{q}^i(\zeta_i);\zeta_i)$. If it is the case that $(\mu(\zeta_i)-\lambda)/M\leq\Delta^i_R(\lambda;\zeta)$ for all λ, then it is impossible that the overall inequality is satisfied for all λ (its left-hand side converges to zero in λ, while the right-hand side converges to a strictly positive value) and incentive compatibility is violated. Therefore we assume that the $\min\{\cdot,\cdot\}$ resolves to $\Delta^i_R(\lambda;\zeta)$. Then the overall inequality implies

$$\lambda\mathbb{E}_{\zeta|s}\left[\Delta^i_L\left(\lambda;\zeta\right)\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right]\geq\mathbb{E}_{Q,\zeta|s}\left[\frac{1}{2}\left(\mu\left(\zeta_i\right)-\lambda\right)\Delta^i_R\left(\lambda;\zeta\right)\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right].$$

Since $\Delta^i_R(\lambda;\zeta)$ is bounded, there is $m^i(\lambda)$ such that

$$\lambda\mathbb{E}_{\zeta|s}\left[\Delta^i_L\left(\lambda;\zeta\right)\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right]\geq\frac{1}{2}\left(m^i\left(\lambda\right)-\lambda\right)\mathbb{E}_{Q,\zeta|s}\left[\Delta^i_R\left(\lambda;\zeta\right)\middle|q_i\geq\underline{q}^i_\lambda\left(\zeta_i\right)\right].$$

For any i, any λ, and any $\kappa > 0$, there is $\Lambda^i(\lambda, \kappa) > 0$ such that

$$\Lambda^i(\lambda, \kappa) < \frac{1}{2}\left(m^i(\lambda) - \lambda\right)\kappa.$$

The term $m^i(\lambda)$ can be specified so that $m^i(\lambda) - \lambda$ is decreasing in λ, so if $\Lambda^i(\lambda;\kappa) < (m^i(\lambda) - \lambda)\kappa/2$, then $\Lambda^i(\lambda,\kappa) < (m^i(\lambda') - \lambda')\kappa/2$ for all $\lambda' > \lambda$. Then let $\overline{\Lambda} = \min\{\Lambda^i(\lambda,\kappa)\colon \Pr_{\zeta_i}(b^i(\overline{q}^i(\zeta_i);\zeta_i) > v^i(\overline{q}^i(\zeta_i);\zeta_i)|s) > 0\}$. For any such κ, $\overline{\Lambda}$, it must be that

$$\kappa \mathbb{E}_{\zeta|s}\left[\Delta_L^i\left(\overline{\Lambda};\zeta\right)\middle| q_i \geq \underline{q}_\lambda^i(\zeta_i)\right] \geq \mathbb{E}_{Q,\zeta|s}\left[\Delta_R^i\left(\overline{\Lambda};\zeta\right)\middle| q_i \geq \underline{q}_\lambda^i(\zeta_i)\right].$$

Define bidder j with type ζ_j to be *relevant* given price p (and common signal s) if $b^j(\overline{q}^j(\zeta_j);\zeta_j) \leq p < v^j(\overline{q}^j(\zeta_j);\zeta_j)$. Fixing price p and summing the above incentive inequality over all relevant agents gives

$$\begin{aligned}
&\kappa \sum_{j \text{ relevant}} \mathbb{E}_{\zeta|s}\left[\Delta_L^j\left(\overline{\Lambda};\zeta\right)\middle| q_j \geq \underline{q}_\lambda^j(\zeta_j)\right] \\
&\geq \sum_{j \text{ relevant}} \mathbb{E}_{\zeta|s}\left[\Delta_R^j\left(\overline{\Lambda};\zeta\right)\middle| q_j \geq \underline{q}_\lambda^j(\zeta_j)\right] \\
&= \sum_{j \text{ relevant}} \mathbb{E}_{\zeta|s}\left[\Delta^j\left(\overline{\Lambda};\zeta\right)\middle| q_j \geq \underline{q}_\lambda^j(\zeta_j)\right] - \mathbb{E}_{Q,\zeta|s}\left[\Delta_L^j\left(\overline{\Lambda};\zeta\right)\middle| q_j \geq \underline{q}_\lambda^j(\zeta_j)\right].
\end{aligned}$$

Thus,

$$(\kappa+1)\sum_{j \text{ relevant}} \mathbb{E}_{\zeta|s}\left[\Delta_L^j\left(\overline{\Lambda};\zeta\right)\middle| q_j \geq \underline{q}_\lambda^j(\zeta_j)\right] \geq \sum_{j \text{ relevant}} \mathbb{E}_{Q,\zeta|s}\left[\Delta^j\left(\overline{\Lambda};\zeta\right)\middle| q_j \geq \underline{q}_\lambda^j(\zeta_j)\right].$$

By definition, $\Delta^j(\overline{\Lambda};\zeta) = Q - \underline{q}_\lambda^j(\zeta_j) - \sum_{k\neq j}\underline{q}^{kj}(\overline{\Lambda};\zeta) \equiv Q - \underline{Q}^j(\overline{\Lambda};\zeta)$ and $\Delta_L^j(\overline{\Lambda};\zeta) = q^j(\zeta) - \underline{q}_\lambda^j(\overline{\Lambda};\zeta)$. Furthermore,

$$\sum_{j \text{ relevant}} q^j(\zeta) - \underline{q}_\lambda^j(\zeta_j) \leq \sum_j q^j(\zeta) - \underline{q}_\lambda^j(\zeta_j) = Q - \underline{Q}\left(\underline{p} + \delta\right).$$

Then it follows that

$$\kappa + 1 \geq \#\{j \text{ relevant}\}.$$

Since $\kappa > 0$ may be arbitrarily small, it follows that there is at most one relevant bidder; i.e., there is at most a single bidder i such that $\Pr(b^i(\overline{q}^i(\zeta);\zeta) < v(\overline{q}^i(\zeta);\zeta)) < 1$. □

Lemma 3. *For all bidders i and all bidder-common signals s,*

$$\Pr\left(b^i\left(\overline{q}^i(\zeta_i);\zeta_i\right) = v\left(\overline{q}^i(\zeta_i);\zeta_i\right)\middle| s\right) = 1.$$

Proof. Fix a common signal s. Lemma 2 shows that at least $n-1$ bidders j are such that $b^j(\overline{q}^j(\zeta_j);\zeta_j) = v(\overline{q}^j(\zeta_j);\zeta_j)$ with probability 1. If all n bidders' bids satisfy this condition, the desired result follows immediately from market clearing. Otherwise, there is some bidder i such that $b^i(\overline{q}^i(\zeta_i);\zeta_i) < v(\overline{q}^i(\zeta_i);\zeta_i)$ with $\zeta_i|s$-strictly positive probability. We show that (i) this bidder's bid must be constant in a neighborhood of $\overline{q}^i(\zeta_i)$, (ii) with $\zeta_{-i}|s$-positive probability, opposing bidders' bids are asymptotically flat near $\overline{q}^j(\zeta_j)$, and (iii) this implies that bidder i has a strict incentive to increase her (flat) bid near $\overline{q}^i(\zeta_i)$.

Let bidder i and parameter ζ_i be such that $b^i(\overline{q}^i(\zeta_i);\zeta_i) = \underline{p} < v(\overline{q}^i(\zeta_i);\zeta_i)$, and assume that b^i is strictly decreasing in a neighborhood to the left of $\overline{q}^i(\zeta_i)$. For $\lambda > 0$, define an alternate bid b^λ,

$$b^\lambda(q) = \begin{cases} b^i(q;\zeta_i) & \text{if } b^i(q;\zeta_i) \geq \underline{p}+\lambda, \\ \underline{p}+\lambda & \text{otherwise.} \end{cases}$$

Since $b^i(\overline{q}^i(\zeta_i);\zeta_i) < v(\overline{q}^i(\zeta_i);\zeta_i)$ and we analyze small $\lambda > 0$, we may assume that λ is small enough that for any feasible quantity q, $b^\lambda(q) \leq v(q;\zeta_i)$. Then whenever the market clearing price would be $p < \underline{p}+\lambda$ if bidder i submitted bid b^i, the market clearing price will be $\underline{p}+\lambda$ if she submits bid b^λ instead. Further, bidder i receives the full residual supply,

$$q_i^\lambda = Q - \sum_{j\neq i} \overline{\varphi}^j\left(\underline{p}+\lambda;\zeta_j\right).$$

The utility gain associated with bid b^λ versus bid b^i is bounded below by

$$\mathbb{E}_{\zeta_{-i}}\left[\begin{array}{l} \displaystyle\int_q^{Q-\sum_{j\neq i}\overline{\varphi}^j\left(\underline{p}+\lambda;\zeta_j\right)} v\left(x;\zeta_i\right) - \left(\underline{p}+\lambda\right) dx \\ \displaystyle -\int_{\overline{\varphi}^i\left(\underline{p}+\lambda;\zeta_i\right)}^{q} \left(\underline{p}+\lambda\right) - b^i\left(x;\zeta_i\right) dx \end{array}\middle| q \geq \overline{\varphi}^i\left(\underline{p}+\lambda;\zeta_i\right)\right]. \tag{3}$$

Because bidder i's opponents all have $\Pr(b^j(\overline{q}^j(\zeta_j);\zeta_j) = v(\overline{q}^j(\zeta_j);\zeta_j)) = 1$, and bids are below values and values are Lipschitz continuous, there is $M > 0$ such that $\overline{q}^j(\zeta_j) - \overline{\varphi}^j(\underline{p}+\lambda;\zeta_j) > M\lambda$ with probability 1 for all $j \neq i$. Then, letting $\lambda < v(\overline{q}^i(\zeta_i);\zeta_i) - b^i(\overline{q}^i(\zeta_i);\zeta_i)$, the bound in 3 is in turn bounded below by

$$\mathbb{E}_{\zeta_{-i}|s}\left[\int_{q}^{\left(Q-\sum_{j\neq i}\overline{q}^{j}(\zeta_j)\right)+(n-1)M\lambda} v\left(x\right)-\left(\underline{p}+\lambda\right)dx-\left(q-\overline{\varphi}^{i}\left(\underline{p}+\lambda;\zeta_i\right)\right)\lambda\middle| q\geq\overline{\varphi}^{i}\left(\underline{p}+\lambda;\zeta_i\right)\right]$$
$$\geq\mathbb{E}_{\zeta_{-i}|s}\left[\left(\left[\left(Q-\sum_{j\neq i}\overline{q}^{j}\left(\zeta_j\right)\right)+\left(n-1\right)M\lambda\right]-q\right)\lambda-\left(\left[Q-\overline{Q}\right]+\left(n-1\right)M\lambda\right)\lambda\middle| q\geq\overline{\varphi}^{i}\left(\underline{p}+\lambda;\zeta_i\right)\right]$$
$$=\mathbb{E}_{\zeta_{-i}|s}\left[\overline{Q}-\sum_{j\neq i}\overline{q}^{j}\left(\zeta_j\right)-q\middle| q\geq\overline{\varphi}^{i}\left(\underline{p}+\lambda;\zeta_i\right)\right]\lambda>0.$$

In the above we rely on the fact that the minimum market clearing price is obtained when aggregate supply is maximized. Since b^{λ} yields higher expected utility than b^{i} when $\lambda>0$ is small, b^{i} is not a best response, and therefore any best response b^{i} must be constant in a neighborhood of $\overline{q}^{i}(\zeta_i)$, if $b^{i}(\overline{q}^{i}(\zeta_i);\zeta_i)<v(\overline{q}^{i}(\zeta_i);\zeta_i)$.

Define $\check{q}^{i}(\zeta_i)=\underline{\varphi}^{i}(\underline{p};\zeta_i)$ to be the left endpoint of the flat interval of bidder i's bid, containing $\overline{q}^{i}(\zeta_i)$. Without loss of generality, we may assume that $b^{i}(q;\zeta_i)=\underline{p}$ for all $q>\check{q}^{i}(\zeta_i)$ whenever $b^{i}(\overline{q}^{i}(\zeta_i);\zeta_i)<v(\overline{q}^{i}(\zeta_i);\zeta_i)$: extending the flat portion of the bid function either does not affect allocation, or (by market clearing) increases allocation to some q such that $v(q;\zeta_i)>\underline{p}$. Since $\Pr(b^{i}(\overline{q}^{i}(\zeta_i);\zeta_i)<v(\overline{q}^{i}(\zeta_i);\zeta_i)|s)>0$ and $\check{q}^{i}(\zeta_i)<\overline{q}^{i}(\zeta_i)$ for all ζ_i with $b^{i}(\overline{q}^{i}(\zeta_i);\zeta_i)<v(\overline{q}^{i}(\zeta_i);\zeta_i)$, it follows that $\Pr(p(Q,\zeta)=\underline{p}|s)>0$. Consider a bidder $j\neq i$ and type ζ_j such that $b^{j}(\overline{q}^{j}(\zeta_j);\zeta_j)=v(\overline{q}^{j}(\zeta_j);\zeta_j)=\underline{p}$; since $\Pr(p(Q,\zeta)=\underline{p}|s)>0$, it must be that $\Pr(q_j=\overline{q}^{j}(\zeta_j)|s)>0$. If the bid $b^{j}(\cdot;\zeta_j)$ is optimal, it must not be utility-improving to decrease the bid to $b^{\lambda\mu}$, where[7]

$$b^{\lambda\mu}\left(q\right)=\begin{cases}b^{j}\left(q;\zeta_j\right) & \text{if } q<\overline{q}^{j}\left(\zeta_j\right)-\lambda,\\ \underline{p}+\mu & \text{otherwise.}\end{cases}$$

The bid $b^{\lambda\mu}$ saves payment $\int_{\overline{q}^{j}(\zeta_j)-\lambda}^{\overline{q}^{j}(\zeta_j)}b^{j}(q;\zeta_j)-(\underline{p}+\mu)dq$ whenever $q_j=\overline{q}^{j}(\zeta_j)$, but potentially reduces quantity when $q_j\in(\overline{q}^{j}(\zeta_j)-\lambda,\overline{q}^{j}(\zeta_j))$. The change in utility from implementing bid $b^{\lambda\mu}$ instead of bid $b^{j}(\cdot;\zeta_j)$ is bounded below by

$$\int_{\overline{q}^{j}(\zeta_j)-\lambda}^{\overline{q}^{j}(\zeta_j)}b^{j}\left(q;\zeta_j\right)-\left(\underline{p}+\mu\right)dq\Pr\left(q_j=\overline{q}^{j}\left(\zeta_j\right)\middle|s\right)$$
$$-\int_{\overline{q}^{j}(\zeta_j)-\lambda}^{\overline{q}^{j}(\zeta_j)}\int_{\overline{q}^{j}(\zeta_j)-\lambda}^{q}v\left(x;\zeta_j\right)-b^{j}\left(x;\zeta_j\right)dxdG^{i}\left(q;b^{j}\right).$$

[7]The μ term ensures that bidder j wins ties against the flat portion of bidder i's bid; this term will be taken to zero and thus will have no marginal effect on utility.

The derivative of this expression with respect to λ must be weakly negative,

$$\begin{aligned}&\left(b^j\left(\overline{q}^j\left(\zeta_j\right)-\lambda;\zeta_j\right)-\left(\underline{p}+\mu\right)\right)\Pr\left(q_j=\overline{q}^j\left(\zeta_j\right)\middle|s\right)\\&-\left(v\left(\overline{q}^j\left(\zeta_j\right)-\lambda;\zeta_j\right)-b^j\left(\overline{q}^j\left(\zeta_j\right)-\lambda;\zeta_j\right)\right)\Pr\left(q_j\in\left(\overline{q}^j\left(\zeta_j\right)-\lambda,\overline{q}^j\left(\zeta_j\right)\right)\middle|s\right)\leq 0.\end{aligned}$$

This inequality holds for all $\mu>0$. Letting M be the Lipschitz modulus of v, substituting in for $b^j(\overline{q}^j(\zeta_j);\zeta_j)=\underline{p}$ means that the previous inequality implies

$$\begin{aligned}&\left(b^j\left(\overline{q}^j\left(\zeta_j\right)-\lambda;\zeta_j\right)-\underline{p}\right)\Pr\left(q_j=\overline{q}^j\left(\zeta_j\right)\middle|s\right)-M\lambda\Pr\left(q_j\in\left(\overline{q}^j\left(\zeta_j\right)-\lambda,\overline{q}^j\left(\zeta_j\right)\right)\middle|s\right)\leq 0\\\iff\quad&-\frac{b^j\left(\overline{q}^j\left(\zeta_j\right);\zeta_j\right)-b^j\left(\overline{q}^j\left(\zeta_j\right)-\lambda;\zeta_j\right)}{\lambda}\leq\frac{M\Pr\left(q_j\in\left(\overline{q}^j\left(\zeta_j\right)-\lambda,\overline{q}^j\left(\zeta_j\right)\right)\middle|s\right)}{\Pr\left(q_j=\overline{q}^j\left(\zeta_j\right)\middle|s\right)}.\end{aligned}$$

Taking the limit as $\lambda\searrow 0$, we obtain that $b^j_q(\overline{q}^j(\zeta_j);\zeta_j)=0$. Thus any bidder $j\neq i$ with type ζ_j such that $b^j(\overline{q}^j(\zeta_j);\zeta_j)=v(\overline{q}^j(\zeta_j);\zeta_j)$ and $\Pr(q_j=\overline{q}^j(\zeta_j)|s)>0$ is such that $b^j_q(\overline{q}^j(\zeta_j);\zeta_j)=0$.

Now return to bidder i with type ζ_i such that $b^i(\overline{q}^i(\zeta_i);\zeta_i)<v(\overline{q}^i(\zeta_i);\zeta_i)$ and $\breve{q}^i(\zeta_i)<\overline{q}^i(\zeta_i)$, and consider the alternate bid function b^λ defined in the first portion of this proof. We now place a slightly different bound on the utility gained by implementing bid b^λ versus bid $b^i(\cdot;\zeta_i)$. Payments increase by at most $\overline{Q}\lambda$, with at most probability 1; and, whenever $q_i>\breve{q}^i(\zeta_i)$ under $b^i(\cdot;\zeta_i)$, bidder i receives the full residual quantity $Q-\sum_{j\neq i}\overline{\varphi}^j(\underline{p}+\lambda;\zeta_j)$. Then a lower bound on the utility improvement generated by the alternate bid b^λ (versus $b^i(\cdot;\zeta_i)$) is

$$\mathbb{E}_{\zeta_{-i}}\left[\int_q^{Q-\sum_{j\neq i}\overline{\varphi}^j\left(\underline{p}+\lambda;\zeta_j\right)}v\left(x;\zeta_i\right)-\underline{p}dx-\overline{Q}\lambda\middle|q\geq\breve{q}^i\left(\zeta_i\right)\right].$$

For b^λ to not be utility-improving, this expectation must be weakly negative. Dividing through by λ and taking the limit at $\lambda\searrow 0$ gives

$$\mathbb{E}_{\zeta_{-i}}\left[-\left(v\left(Q-\sum_{j\neq i}\overline{q}^j\left(\zeta_j\right);\zeta_i\right)-\underline{p}\right)\sum_{j\neq i}\overline{\varphi}^j_p\left(\underline{p};\zeta_j\right)-\overline{Q}\middle|q\geq\breve{q}^i\left(\zeta_i\right)\right]\leq 0.$$

By assumption, $v(Q-\sum_{j\neq i}\overline{q}^j(\zeta_j);\zeta_j)>\underline{p}$, and from the previous paragraph we have that $\overline{\varphi}^j_p(\underline{p};\zeta_j)=-\infty$ with strictly positive probability. Then the above inequality cannot be satisfied. It follows that there is no bidder i such that $\Pr(b^i(\overline{q}^i(\zeta_i);\zeta_i)<v(\overline{q}^i(\zeta_i);\zeta_i)|s)>0$. □

Corollary 2. [Bid pasting under symmetry] *When bidders are symmetric and have symmetric information, $(s,\theta_i)=(s,0)$ for all bidders i, the equilibrium market clearing price*

equals the marginal value for efficient quantity,

$$p^{\star}(s) = \hat{v}\left(nq^{\star}(s); s\right).$$

Corollary 3. [Efficient bid pasting at minimum price] *Let $q_i^{eff}(s)$ be the efficient allocation to bidder i when the signal is s. The minimum market-clearing price $\underline{p}$ is*

$$\underline{p} = \operatorname*{ess\,inf}_{i,s'} v^i\left(q_i^{eff}(s'); s_i'\right).$$

Proof. First we show that $\underline{p} \leq \operatorname{ess\,inf}_{i,s'} v^i(q_i^{\text{eff}}(s'); s_i')$. Otherwise, Theorem 5 implies that there is a bidder i and signal s such that $\overline{q}^i(s) < q_i^{\text{eff}}(s)$. Let $Q^{\text{eff}}(s) = \sum_{i\in I} q_i^{\text{eff}}(s)$. By assumption,

$$p^{\star}(s) > v^i\left(q_i^{\text{eff}}(s); s_i\right) = C\left(Q^{\text{eff}}(s)\right).$$

Then $Q(s) = \sum_{i\in I} q_i(s) \geq Q^{\text{eff}}(s)$, and there is a bidder j for whom $q_j(s) \geq q_j^{\text{eff}}(s)$. For this bidder,

$$v^j\left(q_j(s); s_j\right) \leq v^j\left(q_j^{\text{eff}}(s); s_j\right) = v^i\left(q_i^{\text{eff}}(s); s_i\right) < p^{\star}(s).$$

Thus bidder j is bidding above their marginal value, which cannot occur in equilibrium.

Second we show that $\underline{p} \geq \operatorname{ess\,inf}_{i,s'} v^i(q_i^{\text{eff}}(s'); s_i')$. Otherwise, there is a signal s such that $p^{\star}(s) < \operatorname{ess\,inf}_{i,s'} v^i(q_i^{\text{eff}}(s'); s_i')$; Theorem 5 further implies that there is a signal $\tilde{s}$ with $v^i(q_i^{\star}(\tilde{s}); \tilde{s}_i) \leq v^i(q_i^{\text{eff}}(\tilde{s}); \tilde{s}_i)$. Since marginal values are decreasing, it must be that $Q^{\star}(\tilde{s}) = \sum_{i\in I} q_i^{\star}(\tilde{s}) > \sum_{i\in I} q_i^{\text{eff}}(s') = Q^{\text{eff}}(s')$. Market clearing then implies that $p^{\star}(Q^{\star}(\tilde{s})) \geq p^{\star}(Q^{\text{eff}}(\tilde{s}))$, since the supply curve S is increasing, a contradiction. Thus $\underline{p} \geq \operatorname{ess\,inf}_{i,s'} v^i(q_i^{\text{eff}}(s'); s_i')$. □

B Proofs for Section 2

Proof of Lemma 1. Consider a supply curve S such that the uniform-price auction admits a bidding equilibrium $b^{\star}$ such that $q \circ b^{\star} = q^{\star}$ (if no such S exists, the claim is trivially true). By assumption, there is a level of aggregate supply $\hat{Q}$ such that S is locally nonconstant at $\hat{Q}$; without loss of generality, assume that $S(Q') < S(\hat{Q})$ for all $Q' < \hat{Q}$ (otherwise, the result goes through assuming that $S(\hat{Q}) < S(Q')$ for all $Q' > \hat{Q}$). Fix $\delta > 0$, and let $\varepsilon = S(\hat{Q}) - S(\hat{Q} - n\delta) > 0$. Without loss of generality we may assume that there is $\tilde{Q} \in [\hat{Q} - n\delta, \hat{Q})$ such that for all $Q' \in (\tilde{Q}, \hat{Q}]$, $S(Q') - S(\tilde{Q}) \geq (Q' - \tilde{Q})\varepsilon/\delta$.[8]

[8] An exception to this condition arises when S is discontinuous at $\hat{Q}$. In this case we may analyze $\hat{Q}'$ just above $\hat{Q}$, taking limits where appropriate.

Let the bidder signal $\hat{s}$ be such that optimal aggregate supply is equal to $\hat{Q}$, that is $Q^\star(\hat{s}) = \hat{Q}$. Define $\tilde{\delta} = (\hat{Q} - \tilde{Q})/n$, and consider a bid function $\hat{b}$ where

$$\hat{b}(q) = \begin{cases} v(0; \hat{s}) & \text{if } q < \frac{1}{n}\hat{Q} - \tilde{\delta} \\ 0 & \text{otherwise.} \end{cases}$$

If all bidders submit the bid function $\hat{b}$, aggregate supply is $\hat{Q} - n\tilde{\delta}$, the market-clearing price is $S(\hat{Q} - n\tilde{\delta})$, and each bidder receives quantity $(\hat{Q}/n) - \tilde{\delta}$. Because supply is deterministic, it is sufficient to show that there is no constant bid $\tilde{b}$ that yields any deviating bidder greater utility. By submitting the alternate bid $\tilde{b} \in (S(\hat{Q} - n\tilde{\delta}), S(\hat{Q})]$, the deviating bidder guarantees herself quantity $q = S^{-1}(\tilde{b}) - [(n-1)\hat{Q}/n] + (n-1)\tilde{\delta}$, and her utility increases by

$$\int_{\frac{1}{n}\hat{Q} - \tilde{\delta}}^{S^{-1}\left(\tilde{b}\right) - (n-1)\left[\frac{1}{n}\hat{Q} - \tilde{\delta}\right]} v(x; s) - S\left(\hat{Q} - n\tilde{\delta}\right) dx - \left[\tilde{b} - S\left(\hat{Q} - n\tilde{\delta}\right)\right]\left[S^{-1}\left(\tilde{b}\right) - (n-1)\left[\frac{1}{n}\hat{Q} - \tilde{\delta}\right]\right].$$

The left-hand term is the increase in gross utility, and the right-hand term is the loss from increased payments. By construction, $S^{-1}(\tilde{b})$ is bounded by

$$\hat{Q} - n\delta \leq S^{-1}\left(\tilde{b}\right) \leq \hat{Q} - n\tilde{\delta} + \left(\tilde{b} - S\left(\hat{Q} - n\tilde{\delta}\right)\right)\frac{\tilde{\delta}}{\varepsilon}.$$

It follows that the effect of the deviation $\hat{b}$ on the bidder's utility is bounded above by

$$\int_{\frac{1}{n}\hat{Q} - \tilde{\delta}}^{\frac{1}{n}\hat{Q} - \tilde{\delta} + \left(\tilde{b} - S\left(\hat{Q} - n\tilde{\delta}\right)\right)\frac{\delta}{\varepsilon}} v(x; s) - S\left(\hat{Q} - n\tilde{\delta}\right) dx - \left[\tilde{b} - S\left(\hat{Q} - n\tilde{\delta}\right)\right]\left[\frac{1}{n}\hat{Q} - \tilde{\delta}\right].$$

The derivative of this bound with respect to $\tilde{b}$ is

$$\left[v\left(\frac{1}{n}\hat{Q} - \tilde{\delta} + \left(\tilde{b} - S\left(\hat{Q} - n\tilde{\delta}\right)\right)\frac{\delta}{\varepsilon}\right) - S\left(\hat{Q} - n\tilde{\delta}\right)\right] - \left[\frac{1}{n}\hat{Q} - \tilde{\delta}\right].$$

Since $\tilde{\delta} \leq \delta$, when δ is small this derivative is strictly negative (recall that $v(\hat{Q}/n; s) = S(\hat{Q})$). Then there is no profitable deviation when the submitted bid profile is $(\hat{b})_{i=1}^n$, and $q(\hat{b}) \neq q^\star(s)$ by construction. □

C Proof of Theorem 1

Proof. Let S be a strictly increasing supply curve.[9] If S is optimal, there cannot be a quantity $\hat{q}$ and a perturbation $\varepsilon > 0$ such that S^{ε} yields greater efficiency, where

$$S^{\varepsilon}(q) = \begin{cases} S(q) & \text{if } q \notin [\hat{q}, \hat{q}+\varepsilon], \\ S(\hat{q}+\varepsilon) & \text{if } q \in [\hat{q}, \hat{q}+\varepsilon]. \end{cases}$$

The expected inefficiency induced by S is

$$\mathbb{E}_{\theta,\tilde{S}}\left[\int_{nq^{\star}(\theta)}^{nq^{\text{eff}}\left(\theta,\tilde{S}\right)} \hat{v}(x;\theta) - \tilde{S}(x)\, dx\right].$$

Here, $nq^{\star}(\theta)$ is the equilibrium (aggregate) quantity allocated when the bidders' signal is θ, and $nq^{\text{eff}}(\theta, \tilde{S})$ is the ex post efficient aggregate quantity given bidder signal θ and realized marginal damage curve $\tilde{S}$. Letting $nq^{\varepsilon}(\theta)$ be the equilibrium aggregate quantity under supply curve S^{ε}, the expected inefficiency induced by S^{ε} is

$$\mathbb{E}_{\theta,\tilde{S}}\left[\int_{nq^{\varepsilon}(\theta)}^{nq^{\text{eff}}\left(\theta,\tilde{S}\right)} \hat{v}(x;\theta) - \tilde{S}(x)\, dx\right].$$

The difference between the two is

$$\mathbb{E}_{\theta,\tilde{S}}\left[\int_{nq^{\varepsilon}(\theta)}^{nq^{\star}(\theta)} \hat{v}(x;\theta) - \tilde{S}(x)\, dx\right].$$

If S is optimal, it must be that the derivative of the above with respect to ε is zero:

$$\mathbb{E}_{\theta,\tilde{S}}\left[\left(\hat{v}\left(q^{\varepsilon}(\theta);\theta\right) - \tilde{S}\left(q^{\varepsilon}(\theta)\right)\right) n\frac{dq^{\varepsilon}(\theta)}{d\varepsilon}\right]\Bigg|_{\varepsilon=0} = 0.$$

At $\varepsilon = 0$, we have $q^{\varepsilon}(\theta) = \hat{q}$; and since we are examining pay as bid, we have $\hat{v}(\hat{q};\theta) = S(\hat{q})$. Then the above is

$$\mathbb{E}_{\theta,\tilde{S}}\left[\left(S(\hat{q}) - \tilde{S}(\hat{q})\right) n\frac{dq^{\varepsilon}(\theta)}{d\varepsilon}\,\middle|\, q^{\star}(\theta) = \hat{q}\right] = 0.$$

[9] The case of locally-constant supply curves may be handled by examining the endpoints of the locally-constant region. If the supply curve is globally constant, the analysis must be adjusted but the result does not change.

Note that the left-hand multiplicand is independent of θ, and the right-hand multiplicand is independent of $\tilde{S}$. Thus we have

$$n\mathbb{E}_\theta\left[\left.\frac{dq^\varepsilon(\theta)}{d\varepsilon}\right| q^\star(\theta)=\hat{q}\right]\mathbb{E}_{\tilde{S}}\left[S(\hat{q})-\tilde{S}(\hat{q})\right]=0.$$

It follows that S is optimal only if $S(\hat{q})=\mathbb{E}_{\tilde{S}}[\tilde{S}(\hat{q})]$, for all $\hat{q}$. □

D Proofs for Section 5

Lemma 4. [Range of market price given small support] *Let $s=((s,\theta_i))_{i\in I}$ be a signal profile, and let $(q_i^\star(s))_{i\in I}$ be its associated efficient quantity allocation. If there is ε such that, for all signals s' and all bidders i, $|v^i(q_i^\star(s);s')-v^i(q_i^\star(s);s)|<\varepsilon$, thenfor all signals s' it is the case that in TCPAB, $|p^\star(s')-p^\star(s)|<\varepsilon$.*

Proof of Theorem 4. First we show that $p^\star(s')<p^\star(s)+\varepsilon$ for all s'. To establish a contradiction, suppose to the contrary that there exists a signal s' such that $p^\star(s')\geq p^\star(s)+\varepsilon$. Because aggregate supply is weakly increasing in market price, it must be that $\sum_{i\in I}q_i^\star(s)\leq\sum_{i\in I}q_i^\star(s')$, hence there is some bidder i such that $q_i^\star(s)<q_i^\star(s')$. By assumption,

$$v^i\left(q_i^\star\left(s'\right);s'\right)\leq v^i\left(q_i^\star\left(s\right);s'\right)<v^i\left(q_i^\star\left(s\right);s\right)+\varepsilon<p^\star\left(s'\right).$$

Thus this bidder is bidding above their marginal value, which is not feasible in equilibrium.

Second we show that $p^\star(s')>p^\star(s)-\varepsilon$ for all s'. To establish a contradiction, suppose to the contrary that there exists a signal s' such that $p^\star(s')\leq p^\star(s)-\varepsilon$. Since the minimum market-clearing price $\underline{p}$ is weakly below $p^\star(s')$, Corollary 3 implies that there is a signal s'' with $v^i(q_i^\star(s'');s'')<p^\star(s)-\varepsilon$ for all bidders i. By the assumption of the theorem, it must be that $\sum_{i\in I}q_i^\star(s'')>\sum_{i\in I}q^\star(s)$; since aggregate supply decreases when the market-clearing price falls, this contradicts the assumption that $q^\star(s'')$ is efficient.

Theorem 5 shows that the equilibrium market-clearing price is bounded below by

$$\underline{p}=\operatorname*{ess\,inf}_{i,s'}v^i\left(q_i^\star\left(s'\right);s_i'\right)\geq\operatorname*{ess\,inf}_{i,s'}v^i\left(q_i^\star\left(s\right);s_i'\right)-\delta.$$

Since marginal damages are increasing in quantity, Theorem 5 further implies that the equilibrium market-clearing price is bounded above by $\underline{p}+2\delta$. By Lemma 4, the efficient market-clearing price will be $p^\star(s')\in[\underline{p},\underline{p}+2\delta]$; this establishes the first claim of the theorem. The second claim, bounding efficiency loss, follows by analyzing the area under curve with height at most 2δ and width at most δ/γ. □

E Beyond efficiency

Finally, it is known from (Pycia and Woodward 2022b) that the revenue-maximizing pay as bid auction has deterministic supply and yields more revenue than all uniform-price auctions. If the regulator cares about revenue as well as efficiency—as when they reinvest auction proceeds into social projects—this dominance of pay as bid carries over into the current context.

Theorem 6. [Optimality of Elastic Pay as Bid] *Suppose that bidders are symmetric, and that the auctioneer's objective is a monotone combination of auction revenue and welfare:*

$$\pi\left(q;s,\omega\right)=f\left(R\left(q;s\right),C\left(\sum_{i=1}^{n}q_i;\omega\right)\right),$$

where f is increasing in its first argument and decreasing in its second argument. If values v are concave in quantity and monotone in signal, and costs C are convex, then there exists elastic supply such that pay as bid with this supply weakly outperforms all command-and-control schemes, and all tax schemes, all uniform-price auctions.

Proof. It is sufficient to show that properly designed Pay as Bid dominates uniform-price auctions.

We first show that the claim holds in the full-information context, where s is known. Consider an equilibrium $(b^i)_{i=1}^n$ of the full-information uniform-price auction with supply curve S. If the market-clearing price associated with this equilibrium is $p^\star$, at the equilibrium allocation $(q_i)_{i=1}^n$ it must be that $v^i(q_i;s)\leq p^\star$ for all i. Now consider the allocation $(q_i')_{i=1}^n$ where $\sum_{i=1}^n q_i=\sum_{i=1}^n q_i'$ and $v^i(q_i';s)=v^j(q_j';s)$ for all bidders i and j; by market clearing, it must be that $v^i(q_i';s)\equiv p'\geq p^\star$. Consider a pay-as-bid auction with supply curve $\hat{S}(Q')=p'$. By Theorem 5, the equilibrium in this allocation will have market-clearing price $p'\geq p^\star$ and allocation $(q_i')_{i=1}^n$, where $\sum_{i=1}^n q_i'=\sum_{i=1}^n q_i$. Then revenue in the pay-as-bid auction is weakly higher than in the uniform-price auction, and welfare is unchanged. It follows that the regulator's objective is weakly improved by implementing a pay-as-bid auction with supply curve S'.

As in Pycia and Woodward (2022a), if the efficient quantity and price are in monotone correspondence with one another—that is, if for all $(Q,p),(Q',p')\in\{(Q^\star(s),\mathbb{E}C(Q^\star(s);\omega):s\in\operatorname{Supp}s\}$ it is the case that $(Q-Q')(p-p')\geq 0$—we can trace out the optimal supply curve pointwise. Thus an adapted form of regularity is sufficient. Note that, fixing s, the

regulator's objective can be written as

$$\max_{Q} \mathbb{E}\left[\pi\left(Q; s, \omega\right)\right] = \lambda Q \hat{v}\left(Q; s\right) + \left(1-\lambda\right)\left(\int_0^Q \hat{v}\left(x; s\right) ds - \mathbb{E}\left[C\left(Q; \omega\right)\right]\right).$$

The first-order condition is

$$\hat{v}\left(Q; s\right) + \lambda Q \hat{v}_Q\left(Q; s\right) = \left(1-\lambda\right)\mathbb{E}\left[C_Q\left(Q; \omega\right)\right]. \tag{4}$$

When v is monotone in signal and concave in quantity, and C is convex in signal, the implied correspondence between $p^\star(s) = \hat{v}(Q^\star(s); s)$ and $Q^\star(s)$ will be monotone.[10] Following Pycia and Woodward (2022a), this establishes the desired result. □

[10]Suppose otherwise. Then without loss of generality there are signals s and s' with $s' > s$ and such that $p^\star(s') > p^\star(s)$ and $Q^\star(s') < Q^\star(s)$. Then $\mathbb{E}[C_Q(Q^\star(s'); \omega] < \mathbb{E}[C_Q(Q^\star(s); \omega]$, $\hat{v}(Q^\star(s'); s') > \hat{v}(Q^\star(s); s)$, and $Q^\star(s')\hat{v}_Q(Q^\star(s'); s')) \geq Q^\star(s)\hat{v}_Q(Q^\star(s); s)$ (recall that $\hat{v}_Q \leq 0$), establishing a contradiction.