

Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information*

Ittai Abraham [†], Susan Athey [‡], Moshe Babaioff [§], Michael D. Grubb [¶]

December 1, 2016

Abstract

We study the effect of ex-ante information asymmetries on revenue in common-value second-price auctions (SPA). The motivating application of our results is to online advertising auctions in the presence of “cookies,” which allow individual advertisers to recognize advertising opportunities (impressions) for users who, for example, are existing customers. Cookies create substantial information asymmetries both ex ante and at the interim stage, when advertisers form their beliefs. We distinguish information structures in which cookies identify “lemons” (low value impressions) from those in which cookies identify “peaches” (high value impressions). To make progress in a setting with multiple Nash equilibria, we first introduce a new refinement, “tremble robust equilibrium” (TRE). We then characterize the unique TRE in both first-price and second-price common-value auctions with two bidders who each receive binary signals. This generates two novel insights. First, common-value second-price auction revenues are vulnerable to ex ante asymmetry if relatively rare cookies identify lemons, but not if they identify peaches. Second, first-price auction revenues are substantially higher than second-price auction revenues under the same conditions. Two extensions show that these insights are robust in settings with more than two bidders and richer signal structures. Finally, we consider revenue maximization in a richer setting with a private component to valuations.

*First version January 3, 2011. The paper was accepted for presentation at ACM-EC 2013 and a one page abstract appeared in its proceedings. We thank Gleb Romanyuk, Vera Sharunova, and Juan Camilo Castillo for exceptional research assistance. We thank Nageeb Ali and Preston McAfee for careful reading and valuable suggestions. We also thank seminar audiences at Harvard, Stanford GSB, Northwestern, NBER, NYCE Day, UC Berkeley, University of Michigan, Cornell, Microsoft Research, Hebrew University, Bocconi, Bologna, AEA 2011 Meetings, and Tel Aviv International Workshop on Game Theory for useful feedback.

[†]VMware Research, ittaia@gmail.com

[‡]Stanford University and Microsoft Research, athey@stanford.edu

[§]Microsoft Research, moshe@microsoft.com

[¶]Boston College, michael.grubb@bc.edu

1 Introduction

This paper develops new theoretical results about the impact of the information structure on revenue in second-price common-value auctions, with comparison to first-price auctions. Our focus is on situations where bidders are not only asymmetrically informed at the interim stage (after observing their signals) but are also asymmetric at the ex ante stage (before observing their signals). In other words, situations where ex ante it is known that particular bidders are likely to be better informed than others. Our primary motivation is to better understand the market design problems presented by the use of cookies in online advertising markets.

Although there are a variety of mechanisms for selling display advertising, auctions are a leading method, especially for “remnant” inventory, and cookies play an important role in these markets. Cookies placed on users’ computers by specific web sites can be used to match a user with information such as the user’s purchasing history with an online retailer, their recent history of airline searches on a travel website, or their browsing and clicking behavior across a network of online publishers (such as publishers on the same advertising network). As shown by Shiller (2014), such information can be a good predictor of a web surfer’s value to an advertiser. As a result, advertisers increasingly use cookies to customize their bidding and target their advertising in display ad auctions (Helft and Vega, 2010).

For example, Google’s ad exchange is currently described as a second-price auction that takes place in real time: that is, at the moment an internet user views a page on an internet publisher, a call is made to the ad exchange, bidders on the exchange instantaneously view information provided by the exchange about the publisher and the user as well as any cookies they may have for the individual user, and based on that information, place a bid.

The cookie is only meaningful to the bidder if it was placed by the bidder (e.g., Amazon.com may have a cookie on the machines of regular customers), or if the bidder has purchased access to specific cookies from a third-party information broker. Cookie-based bidding potentially makes display auctions inherently asymmetric at both the ex ante and the interim stage. At the ex ante stage, bidders may vary greatly in their likelihood of holding informative cookies, both because popular websites have more opportunities to track visitors and because different sites vary in the sophistication of their tracking technologies. At the interim stage, for a particular impression, a bidder who has a cookie has a substantial information advantage relative to those who do not.

If cookies only provided advertisers with private-value information, then increasing sophistication in the prevalence and use of cookies by advertisers would present ad inventory sellers a two-way trade-off between better matching of advertisements with impressions and reduced competition in thinner markets (Levin and Milgrom, 2010). In such a private value setting, Board (2009) shows

that irrespective of such asymmetry, more cookies and more targeting always increase second-price auction revenue as long as the market is sufficiently thick. However, cookies undoubtedly also contain substantial common-value information. For instance, when one bidder has a cookie which identifies an impression as due to a web-bot rather than a human, the impression is of zero value to all bidders. Similarly, if a cookie identifies a high-income frequent online shopper, the impression is likely of high value to many bidders. As a result, the inherent asymmetry created by cookies can lead to cream skimming or lemons avoidance¹ by informationally advantaged bidders, with potentially dire consequences for seller revenues.

Thus, a designer of online advertising markets (or other markets with similar informational issues) faces an interesting set of market design problems. One question is whether the market should encourage or discourage the use of cookies, and how the performance of the market will be affected by increases in the prevalence of cookies. This is within the control of the market designer: in display advertising, it is up to the marketplace to determine how products are defined. All advertising opportunities from a given publisher can be grouped together, for example. Google's ad exchange reportedly does not support revealing all possible cookies. A second market design question concerns the allocation problem: if an auction is to be used, what format performs best? Both first-price and second-price auctions are used in the industry. There are a number of other design questions, as well, including whether reserve prices, entry fees, or other modifications to a basic auction should be considered. Our analysis begins to address these questions by focusing on comparing two commonly used mechanisms, first and second-price auctions, and identifying information structures in which cookies may be particularly costly. (Design of an optimal mechanism is left for future work.)

In order to understand the market design tradeoffs involved in an environment where some bidders are known *ex ante* to have better access to common-value information, the first part of our paper specifies a model of pure common-value second-price auctions. Perhaps surprisingly, the existing literature leaves a number of questions open. For example, while it is well known that the presence of an informationally-advantaged bidder will moderately reduce seller revenues in a sealed-bid first-price auction (FPA) for an item with common value (Wilson, 1967; Weverbergh, 1979; Milgrom and Weber, 1982b; Engelbrecht-Wiggans, Milgrom, and Weber, 1983; Hendricks and Porter, 1988), substantially less is known about the same issue in the context of second-price auctions. One of the main impediments to progress has been the well known multiplicity of Bayesian Nash Equilibria in second-price common-value auctions (Milgrom, 1981). As a consequence, little is

¹Cream skimming refers to buying up the best inventory, while lemons avoidance refers to avoiding the worst inventory.

known about what types of information structures lead to more or less severe reductions in revenue.

Solution Concept In order to address the multiplicity problem, we begin by suggesting a new refinement, *tremble robust equilibrium*. Tremble robust equilibrium (TRE) selects only Bayesian Nash Equilibria that are near to an equilibrium (in undominated bids) of a perturbed game in which a random bidder enters with vanishingly small probability ϵ and then bids smoothly over the support of valuations. In addition to capturing an aspect of the real-world uncertainty faced by bidders in the kinds of applications we are interested in, we argue that this refinement has a number of attractive properties. First, in all of the cases we analyze, this refinement selects a unique equilibrium. Second, when bidders are ex ante symmetric in the setting with discrete signals we study, TRE selects the analog of the symmetric equilibrium studied by Milgrom and Weber (1982a) in a setting with continuous signals. Third, it rules out intuitively unappealing equilibria in which uninformed bidders bid aggressively because they can rely on others to set fair prices. We provide additional motivation for our choice of refinement in Section 2, where we also discuss some standard refinements and explain why they do not adequately address the multiplicity problem in common-value SPAs.

Main Results We then proceed to analyze a number of special cases of common-value second-price auctions using the TRE refinement. We begin with our baseline model and main results: We characterize the unique TRE in the SPA for any monotonic domain with two bidders who receive binary signals. (By monotonic, we mean that the common-value is nondecreasing in each bidder’s signal.) For comparison, we also characterize the unique TRE with monotonic bidding strategies in the FPA in the same setting (with the additional assumption that signals are affiliated). We characterize seller revenues in each case, and highlight how the information structure affects the difference in revenue raised by the two auction formats.²

To connect this model to display advertising auctions, suppose that there are two bidders and that each bidder uses cookie tracking rather crudely—each can only determine the presence or absence of their own cookie. That is, each bidder receives a binary signal which either takes on the value {no-cookie} or {cookie}, but cannot observe whether the competing bidder has a cookie (though they know the overall information structure, including the probability of cookies). Our results characterize the unique TRE and revenue in this setting for both first-price and second-price auctions.

In a common-value auction, a seller does best when all bidders are equally uninformed, as she

²Murto and Välimäki (2015) also study first-price and second-price common-value auctions with binary signals, focusing on entry costs rather than ex ante asymmetry.

can sell the object at its expected value. When bidders have informative cookies, we expect them to earn information rents and revenues to be lower. The question remains, however, how much lower revenues will be. If cookies are rare there are at least two competing intuitions. On the one hand, if bidders have little information, we might expect that information rents would be low and revenue would be close to expected surplus. This is always true when bidders are symmetric *ex ante*. On the other hand, if one bidder has much better access to cookies than the other, we might expect the less-informed bidder to be a meek competitor in a SPA due to fears of adverse selection—leading to low revenue. We show that which intuition is correct depends importantly on the information structure.

To gain insight from our equilibrium characterization, we consider two scenarios: In the first scenario, cookies identify “peaches,” or high-value impressions. This is perhaps the most natural assumption—someone who has been to an advertiser’s website before is more likely to be an active internet shopper than a random web surfer. In the second scenario, cookies identify “lemons,” or low-value impressions. This might occur if a prior visit indicates the surfer is in fact a web-bot and not a real person. In both cases, one bidder may be *ex ante* more likely to receive a cookie than the other bidder.

In the unique TRE, if bidders are equally well informed *ex ante* then SPA revenue is close to the full surplus when cookies are rare. If bidders are informed about peaches, we find that this remains true regardless of the level of *ex ante* asymmetry. Even when only a single bidder has access to cookies, revenue remains close to expected surplus when cookies are rare. Thus the first intuition that little information leads to little revenue loss holds true. The finding is sharply different, however, if bidders are informed about lemons. In that case, revenues decline as bidders become more asymmetric *ex ante*, falling from full surplus in the case of *ex ante* symmetry to the value of a lemon if only one bidder has access to cookies. Thus the second intuition that adverse selection may undermine SPA revenue when bidders are asymmetric *ex ante* dominates. In short, our first main insight is that common-value SPA revenues are vulnerable to *ex ante* bidder asymmetry when informative cookies are rare and identify lemons, but not when they identify peaches.

For comparison, we also examine FPA revenue in the same settings under the additional assumption that signals are affiliated.³ When bidders are symmetric *ex ante*, FPA revenue coincides with SPA revenue in the unique TRE.⁴ When bidders are asymmetric, however, FPA revenue remains close to full surplus if cookies are rare regardless of whether cookies identify peaches or lemons. Thus, our second main insight is that common-value auction revenue is substantially higher in the

³Affiliation is a strong form of correlation introduced to the auction literature by Milgrom and Weber (1982a).

⁴This is consistent with Milgrom and Weber’s (1982a) result that revenue is equal or higher in the symmetric equilibrium of the SPA than in the FPA when bidders are symmetric *ex ante* and signals are affiliated.

FPA than the SPA when ex ante asymmetric bidders receive informative cookies rarely but those cookies identify lemons.

Extensions We then proceed to two extensions which explore how robust these insights are to settings with more than two bidders and richer signal structures. First, we focus on the special case of extreme ex ante asymmetry in which only a single bidder is informed, but allow for any number of uninformed bidders and any signal structure for the informed bidder. This extension provides an analysis of the SPA that is complementary to the existing work on common-value FPAs with a single informed bidder (Wilson, 1967; Weverbergh, 1979; Milgrom and Weber, 1982b; Engelbrecht-Wiggans et al., 1983; Hendricks and Porter, 1988). In the SPA, TRE predicts that uninformed bidders will essentially choose not to compete: they bid the minimum possible posterior valuation of the informed bidder, which is the seller’s revenue. As a result, our findings about the important distinction between private information about lemons and peaches are robust: if cookies are rare, SPA revenue only suffers substantially when cookies identify lemons. For that reason, while FPA revenues are always higher than SPA revenues when only one bidder is informed, the difference is substantial when cookies are rare and identify lemons, but negligible when they identify peaches.

Our second extension restricts attention to information structures satisfying what we dub the *strong-high-signal property*, but within this setting, allows for more than two bidders, signals with more than two values, and multiple informed bidders that are asymmetric ex ante. We characterize the unique TRE of the SPA (although not the FPA) in this setting. We again find that in common-value auctions with ex ante asymmetric bidders, SPA revenues are much lower when bidders are informed about lemons than when informed about peaches. An important qualitative difference under the strong-high-signal property relative to the two-bidder binary-signal setting is that, when bidders are informed about lemons, SPA revenues collapse with even slight ex ante bidder asymmetry rather than declining smoothly as ex ante bidder asymmetry grows. (Our first extension’s comparison of revenue with a single informed bidder to that with ex ante symmetric bidders is silent on the transition between the two extremes.)

So far, we have focused mainly on the costs of information asymmetry, while suppressing any benefit of cookies. In the last section of the paper, we extend the model beyond pure common values. We show that the problems created by information asymmetry remain and we suggest alternative mechanism designs that extract most of the possible revenue.

2 Solution Concepts: A Discussion

This paper seeks to understand how revenues in a common-value second-price auction depend on the structure of information held by bidders. A serious challenge to comparing revenues across different information structures is that for any given information structure there are typically many different equilibria with widely different revenues (Milgrom, 1981).

A common approach in the literature with ex ante symmetric bidders is to focus on the symmetric equilibrium. As shown by Milgrom and Weber (1982a) and Matthews (1984), this selects the equilibrium in which each bidder bids the object's expected value conditional on the highest signal of competing bidders being equal to her own. This excludes extreme equilibria such as one in which one bidder bids an object's maximum value and all other bidders bid zero. Unfortunately, it is not clear how the symmetry refinement can be extended to ex ante asymmetric environments of the type we are interested in.

Consider the following simple “peach or lemon” scenario. A common-value good is equally likely to be a peach (with value P) or a lemon (with value $L < P$). There are two bidders in a second-price auction. One is perfectly informed about the value of the good, while the other only knows the prior probability it is a peach. What bidding strategies and revenues should we expect?

Nash equilibrium provides no prediction about revenue beyond an upper bound of the full surplus. It is an equilibrium for the informed bidder to bid his value and the uninformed bidder to bid P , which results in full surplus extraction. However, it is also an equilibrium for the uninformed bidder to bid $10P$ and the informed bidder to bid $L/2$, earning revenue $L/2$. There are no symmetric equilibria to focus on.

A natural refinement is to restrict attention to Nash equilibria in which bidders only use undominated bids. For such strategies, bids are always between L and P . Notice that unlike in the private value model, agents do *not* necessarily have a dominant strategy in a common-value second-price auction. Indeed, in the scenario described above the informed agent has a dominant strategy (to bid the value given his signal), while the uninformed agent does not.⁵

Thus, ruling out dominated bids restricts the informed bidder to use her dominant strategy and bid her value. However, the only restriction placed on the uninformed bidder is that he not bid less than L or more than P . Revenue could be anywhere between L and the full surplus.

⁵To see that, observe that for any two bids b_1 and b_2 such that $P \geq b_1 > b_2 \geq L$ there exist two strategies of the informed agent such that for one strategy the utility from b_1 is higher, while for the other strategy the utility from b_2 is higher. Bidding b_1 is superior to bidding b_2 when the informed is bidding $(b_1 + b_2)/2$ when the value is P , and bidding L when the value is L . On the other hand bidding b_2 is superior to bidding b_1 when the informed is bidding $(b_1 + b_2)/2$ when the value is L , and bidding L when the value is P (handing out the good items to the other bidder).

The uninformed bidder faces a severe adverse selection problem: for any bid less than P she only wins lemons. Our intuition is that this adverse selection problem makes the equilibrium in which the uninformed bids L most plausible. The reason higher bids can be equilibrium strategies is that the informed bidder always sets the price. The uninformed can bid above L safe in the knowledge that the price will always be set fairly at the item’s value. The model implicitly ignores the fact that the real world is a risky and uncertain place and that bidding above L exposes the uninformed bidder to the possibility of overpaying for a lemon without any possible benefit of winning a cheap peach.

Now consider perturbing the game by adding, with some small probability $\epsilon > 0$, a non-strategic bidder who bids randomly between L and P using a “nice” distribution (having full support and continuous density between L and P).⁶ The purpose is to make the game “noisy” to eliminate unreasonable equilibria by ensuring that the underlying adverse selection problem in the game has consequences. Given that the informed bidder bids the value, the presence of a random bidder means that L is the only undominated bid for an uninformed bidder. The informed bidder ensures that the uninformed bidder can never win the object at a discount below value. However, the random bidder ensures that any bid above L risks overpaying for a low value object when the random bidder sets the price. Thus bidding above L leads to a negative payoff. We observe that adding noise yields *unique* predictions for equilibrium bidding and revenue.

Motivated by the this example, we want to consider only Nash equilibria that are nearby to Nash equilibria (in undominated bids) of games perturbed with an ϵ probability of an additional random bidder. In the spirit of other perturbation based refinements, such as trembling-hand perfection, we identify Nash equilibria that are nearby by considering the limit as ϵ goes to zero. Therefore, we define a *Tremble Robust Equilibrium (TRE)* to be a Nash equilibrium that is the limit, as ϵ goes to zero, of a series of Nash equilibria (using undominated bids) of each modification of the original game in which another “random” bidder is added with small probability ϵ . The random bidder bids a random value drawn from a distribution with continuous and positive density over the “relevant” values. Moreover, if there is a TRE with a profile of strategies that is a Nash equilibrium not just in the limit as ϵ goes to zero, but also away from the limit for sufficiently small $\epsilon > 0$, we call it a *strong Tremble Robust Equilibrium*. The formal definitions of these new refinements are presented in Section 3.2.

⁶In an analysis of the generalized second price (GSP) auction for sponsored search with independent valuations and complete information, Hashimoto (2013) proposes to refine the set of equilibria by adding a non-strategic random bidder that participates in the auction with small probability. Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007) have shown that GSP has an envy-free efficient equilibrium; the main result of Hashimoto (2013) is that this equilibrium does not survive the refinement.

In the preceding peach or lemon scenario with one informed and one uninformed bidder, the unique TRE is a strong TRE and predicts that the informed bidder bids the value while the uninformed bidder bids the value of a lemon.

2.1 Related Refinements

Perturbation Based Refinements It is natural to ask how TRE compares to Selten’s (1975) trembling-hand perfect equilibrium. The two refinements yield very different predictions in our preceding peach or lemon example with one informed and one uninformed bidder. In particular, in Appendix F, we show that two extensions by Simon and Stinchcombe (1995) of trembling-hand perfection to infinite action-space games (which we adjust to incomplete information) are too permissive: they make the same revenue prediction as Nash equilibrium. Revenues could be anywhere between the value of a lemon and the full surplus. On the other hand, in the same setting, if we restrict the tremble of the informed agent to be *independent* of his signal then the unique trembling-hand perfect equilibrium predicts that the uninformed agent bids the unconditional expected value of the item, contrary to our expectation.

In our simple example with one bidder who learns whether the item is a peach or a lemon, Milgrom and Mollner’s (2016) *test-set equilibrium* is more restrictive than trembling-hand perfection. Test-set equilibria are those in which the informed bidder bids the item’s value, and the uninformed bidder mixes between the two possible values, bidding P with some probability $p \in [0, 1]$ and L otherwise. Unfortunately, this still yields the same ambiguous revenue prediction as Nash equilibrium.

Work by Parreiras (2006), Cheng and Tan (2010), Larson (2009), and Syrgkanis, Kempe, and Tardos (2013) introduce perturbations to select a unique equilibrium in two-bidder auctions with continuously distributed signals. Parreiras (2006) and Syrgkanis et al. (2013) perturb the auction format by assuming that winning bidders pay their own bid rather than the second highest with probability ϵ (Parreiras (2006) focuses on the limit as ϵ goes to zero). Cheng and Tan (2010) and Larson (2009) introduce private value perturbations to the common-value environment and take the limit as these perturbations go to zero. In contrast to TRE, the equilibrium selected is sensitive to assumptions about the distributions of the vanishing perturbations.⁷

⁷Cheng and Tan (2010) assume private value perturbations are perfectly correlated with common-value signals and are symmetric across bidders. The symmetry of perturbations (across asymmetric bidders) selects a unique equilibrium. Larson (2009) allows for asymmetric perturbations which are assumed to be independent of common-value signals and shows that the equilibrium selected depends on the ratio of the standard deviations of the two bidders’ private value perturbations. More generally, Liu (2014) shows that any of the equilibria identified by Milgrom (1981) can be selected by an appropriate choice of the distribution of private value perturbations.

Our finding that sufficient ex ante asymmetry favors first-price auctions over second-price auctions (reversing Milgrom and Weber’s (1982a) result from the symmetric case) is similar to Cheng and Tan’s (2010) result that ex ante asymmetry favors first-price auctions but contrasts with Parreiras’ (2006) and Syrgkanis et al.’s (2013) findings that Milgrom and Weber’s (1982a) first-price and second-price auction revenue ranking result is robust to asymmetry.

Inspired by our work, Liu (2014) studies equilibria that are “robust to noisy bids”, a concept closely related to TRE. Like TRE, the robust-to-noisy-bids refinement considers perturbations in which an additional bidder enters with probability ϵ and bids randomly. Unlike TRE, however, the refinement does not impose Nash equilibrium upon the perturbations. As a result, the refinement is distinct from TRE, and while ruling out “discontinuous” equilibria, admits the entire continuum of equilibria identified by Milgrom (1981).

Iterated Deletion of Dominated Strategies An alternative approach taken in the literature that has been applied to auctions with more than two bidders is to select equilibria that survive iterated deletion of dominated strategies. Harstad and Levin (1985) consider the case in which the first order-statistic of bidders’ signals is a sufficient statistic for the object’s value in the Milgrom and Weber (1982a) setting with symmetric bidders and continuously distributed signals. For this case, Harstad and Levin (1985) shows that iterated deletion of dominated strategies uniquely selects the symmetric Milgrom and Weber (1982a) equilibrium. Einy, Haimanko, Orzach, and Sela (2002) consider the case of asymmetric bidders and discrete signals with finite support. They show that if the information structure is *connected* then iterated deletion of dominated strategies selects a set of *sophisticated equilibria* with a unique Pareto-dominant (from bidders’ perspective) equilibrium.

Malueg and Orzach (2009) apply Einy et al.’s (2002) refinement in two examples and Malueg and Orzach (2012) apply it to the special case of two-bidder auctions with *connected* and *overlapping* information partitions. For a particular one-parameter family of common-value distributions, Malueg and Orzach (2012) find that distributions with sufficiently thin left tails yield lower revenues in second-price auctions than in first-price auctions.

Einy et al.’s (2002) result applies to our lemon or peach example with one informed and one uninformed bidder, as this can be represented as a connected domain. Iterated deletion of dominated strategies is unhelpful on its own: the uninformed bidder may still bid anywhere between the value of a lemon and a peach. However, the Pareto dominant equilibrium for the bidders is that in which the uninformed bidder bids the value of a lemon. This is the equilibrium we believe to be natural and coincides with the unique TRE.

The primary drawback to Einy et al.’s (2002) approach is that the required assumptions on

the information structure are very restrictive. For instance, we show in Online Appendix G that Einy et al.'s (2002) connectedness property is strictly more restrictive than our *strong-high-signal* property. Moreover, connectedness rules out many interesting settings such as our model of two bidders with binary signals in which neither bidder is perfectly informed. In contrast, our TRE refinement selects a unique equilibrium in this setting.

3 The Model and Tremble Robust Equilibrium

3.1 The Model

An auctioneer is offering an indivisible good to a set N of n potential buyers. Let Ω be the set of states of the world (possibly infinite). There is a commonly known prior distribution $H \in \Delta(\Omega)$ over states of the world. Let $\omega \in \Omega$ be the realized state of the world, which is not observed by the buyers. The value of the item to agent i when the state of the world is ω is $v_i(\omega)$, which is bounded. (Our analysis in Section 4 assumes that agents share the same common value $v(\omega)$, but we define TRE in this more general setting.)

Each buyer i gets a signal about the state of the world ω from a finite set of signals S_i . For every state $\omega \in \Omega$ and buyer i , there is a commonly known distribution over signals $d_i(\omega) \in \Delta(S_i)$. Each buyer i gets a private signal $s_i \in S_i$, sampled from $d_i(\omega)$. Signal $s_i \in S_i$ for agent i is *feasible* if agent i receives signal s_i with positive probability, and the vector of signals $s = (s_1, s_2, \dots, s_n) \in S_1 \times S_2 \times \dots \times S_n$ is feasible if it is realized with positive probability. Without loss of generality, we assume that for every i , every signal $s_i \in S_i$ is feasible. We denote the set of feasible signal vectors by S . When buyer i realizes signal s_i , we denote his updated expected value of the good by $v_i(s_i) = E[v_i(\omega)|s_i]$. Similarly, we denote the posterior expected value given signal vector s by $v_i(s) = E[v_i(\omega)|s]$.

3.2 Tremble Robust Equilibrium

We define the TRE refinement in the context of any auction game and in Section 4 apply it to second-price and first-price common-value auctions. The refinement is based on the restriction to the closure⁸ of the set of undominated bids and the addition of a random bidder that bids according

⁸In the second price auctions we study, the set of undominated bids is closed so the distinction does not matter. However, for many games with continuous action spaces but discontinuous payoff functions, such as the first price auctions we study, restricting bidders to undominated strategies can lead to non-existence. Hence we follow the standard approach of allowing for all bids in the closure of the set of undominated bids. See Jackson and Swinkels (2005) for a discussion.

to a *standard distribution*. To define a standard distribution, let v_{min} and v_{max} be the infimum and supremum undominated bids for any bidder i and signal $s_i \in S_i$.⁹

Definition 1 *We say that a distribution R is standard if it is continuous, its support is $[v_{min}, v_{max}]$ (the “relevant” values), and on that support it is differentiable and increasing with density function r that is continuous and positive.*

Consider an auction and the game λ that is induced by the auction. We next define the game perturbed by the addition of a random bidder.

Definition 2 *For a standard distribution R and $\epsilon > 0$, define $\lambda(\epsilon, R)$ to be the game induced by λ with the following modification: with probability ϵ there is an additional bidder submitting a bid b sampled according to R . We call $\lambda(\epsilon, R)$ an (ϵ, R) -tremble of the game λ .*

Let μ_i be a strategy of agent i . A mixed strategy maps the signal of the agent to a distribution over bids. The strategy is a pure strategy if for every signal the mapping is to a single bid. Let μ be a profile of strategies, one for each agent.

Definition 3 *(i) A Nash equilibrium μ is a Tremble Robust Equilibrium (TRE) of the game λ if bidders only bid within the closure of the set of undominated bids and there exists a standard distribution R , a sequence of positive numbers $\{\epsilon_j\}_{j=1}^{\infty}$ that converge to 0, and a sequence of strategy profiles $\{\mu^{\epsilon_j}\}_{j=1}^{\infty}$ such that*

1. *For every ϵ_j , μ^{ϵ_j} is a Nash equilibrium of $\lambda(\epsilon_j, R)$, the (ϵ_j, R) -tremble of the game λ , in which bidders only bid within the closure of the set of undominated bids.*
2. *For each bidder $i \in N$ and signal $s_i \in S_i$, $\{\mu_i^{\epsilon_j}(s_i)\}_{j=1}^{\infty}$ converges in distribution to $\mu_i(s_i)$.*

(ii) μ is a strong Tremble Robust Equilibrium if it is a TRE and, in addition, for the sequence $\{\mu^{\epsilon_j}\}_{j=1}^{\infty}$ satisfying (1) and (2) above, there exists k such that for every $j > k$ it holds that $\mu^{\epsilon_j} = \mu$.

4 Second-Price and First-Price Common-Value Auctions

In this section we consider the restriction of the above model to the common-value case and study the second-price auction (SPA), with comparisons to the first price auction (FPA). When we talk about the *second-price auction game* we refer to the game induced by a second-price auction with a random tie breaking rule, and similarly for the *first price auction game*. In the common-value

⁹If $B_i(s_i)$ is the set of undominated bids for bidder i with signal s_i then $v_{min} = \min_{i \in \{1, \dots, N\}} \min_{s_i \in S_i} \inf B_i(s_i)$ and $v_{max} = \max_{i \in \{1, \dots, N\}} \max_{s_i \in S_i} \sup B_i(s_i)$.

model, the state of the world determines a common value of the good to all buyers such that $v_i(\omega) = v(\omega)$ for some function $v(\omega)$ and every bidder i .

We begin with our main results: We characterize the unique TRE in the SPA for any monotonic domain with two bidders who receive binary signals. (By monotonic, we mean that the expected common-value is nondecreasing in each bidder’s signal.) For comparison, we also characterize the unique TRE with monotonic bidding strategies in the FPA in the same setting (with the additional assumption that signals are affiliated). We characterize seller revenues for each auction format, and highlight how the information structure affects the difference in revenue raised by the two auction formats using two contrasting information structures: An information structure where signals identify “lemons”, or very low value items, and an information structure where signals identify “peaches”, or very high value items. When private information identifies lemons, we find that SPA revenue is much more vulnerable to ex ante bidder asymmetry than is FPA revenue. However, we find that this is not the case when private information identifies peaches. We then proceed to two extensions which explore how robust these insights are to settings with more than two bidders and richer signal structures.

4.1 Baseline Model: Two Agents, Each with a Binary Signal

In this section we begin by characterizing the unique TRE of the SPA for any monotonic domain with two bidders who receive binary signals. Let $\{L_i, H_i\}$ be the low and high signals, respectively, of agent $i \in \{1, 2\}$. With some abuse of notation we will also use H_i to denote the event that the signal of agent i was realized to H_i , and similarly for L_i . We denote $V_{LL} = v(L_1, L_2)$, $V_{HH} = v(H_1, H_2)$, $V_{HL} = v(H_1, L_2)$, $V_{LH} = v(L_1, H_2)$, and make the following assumption:

Assumption 1 *The domain is monotonic ($V_{LH}, V_{HL} \in [V_{LL}, V_{HH}]$), at least one signal is informative ($V_{LL} < V_{HH}$), and all four possible signal realizations arise with positive probability ($\Pr[L_1, L_2], \Pr[L_1, H_2], \Pr[H_1, L_2], \Pr[H_1, H_2] > 0$).*

Note that while Assumption 1 rules out the uninteresting case in which neither bidder’s signal is informative, it does allow for the special case in which one bidder is entirely uninformed. For instance, bidder 2 is entirely uninformed if $V_{LL} = V_{LH} < V_{HL} = V_{HH}$, and $\Pr[H_1|H_2] = \Pr[H_1|L_2]$. Without loss of generality, we label agents 1 and 2 such that:

$$\Pr[H_1, L_2](V_{HH} - V_{HL}) \leq \Pr[L_1, H_2](V_{HH} - V_{LH}). \tag{1}$$

As discussed following Theorem 1, this labeling turns out to identify bidder 1 as the more aggressive bidder conditional on receiving a high signal.

Theorem 1 states our first result—a characterization of bidding in the unique TRE of the second-price auction.

Theorem 1 *Consider any SPA game with two bidders that each receive a binary signal, satisfying Assumption 1. (1) There exists a unique TRE. (2) If bidders are labeled as in equation (1) then, in the unique TRE:*

- *Every bidder i bids V_{LL} when getting signal L_i .*
- *Bidder 1 with signal H_1 always bids V_{HH} .*
- *If $V_{HH} > V_{LH}$, bidder 2 with signal H_2 bids V_{HH} with probability $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}$ and bids V_{LH} with the remaining probability. Otherwise, bidder 2 with signal H_2 bids $V_{HH} = V_{LH}$ with probability 1.*

According to the theorem, both bidders bid conservatively at $b = V_{LL}$ and earn zero payoff when they receive a low signal. Both bidders bid more aggressively conditional on a high signal, but not equally so. While bidder 1 always bids her maximum possible value V_{HH} given a high signal, bidder 2 mixes between the same bid and the lower possible value V_{LH} . Thus equation (1) identifies bidder 1 as the more aggressive bidder conditional on receiving a high signal. The loose intuition is that bidder 1 bids more aggressively because the potential downside from bidding V_{HH} conditional on a high signal is smaller than for bidder 2. The potential downside to bidder 1 (in a tremble of the game) is overpaying for an item worth only V_{HL} . This possibility is less likely when $\Pr[H_1, L_2]$ (and hence $\Pr[L_2|H_1]$) is small, and less consequential when the difference between V_{HH} and V_{HL} is small. Nevertheless, as bidder 2 may receive a high signal more often, bidder 2 may earn a higher expected payoff. We provide intuition for Theorem 1 in our sketch of the proof in Section 4.2. First, however, we discuss important special cases of the theorem and investigate its implications for revenue.

Theorem 1 encompasses two important special cases: (1) ex ante symmetric bidders ($V_{HL} = V_{LH}$ and $\Pr[H_1, L_2] = \Pr[L_1, H_2]$), and (2) a single informed bidder ($V_{LL} = V_{LH} < V_{HL} = V_{HH}$, and $\Pr[H_1|H_2] = \Pr[H_1|L_2]$). If bidders are ex ante symmetric, then our TRE refinement selects the symmetric equilibrium studied by Milgrom and Weber (1982a) and others. In the unique TRE, both agents bid V_{HH} given a high signal but bid V_{LL} otherwise. If only bidder 1 is informed, however, then the setting corresponds to the example discussed in Section 2. In this case, the unique TRE predicts that bidder 1 bids V_{HH} given a high signal and V_{LL} otherwise, but that bidder 2 always bids V_{LL} . In other words, bidder 2 chooses not to compete.

Theorem 1 is not confined to these two special cases, but also spans all the intermediate cases in which both bidders are informed but are nonetheless asymmetric ex ante. Begin with the case in which bidder 1 is the only informed bidder, and consider what changes if bidder 2 becomes informed. Theorem 1 shows that if bidder 2’s signal is informative about bidder 1’s signal ($\Pr[H_1|H_2] \neq \Pr[H_1|L_2]$), but bidder 1’s signal remains a sufficient statistic for the value ($V_{LL} = V_{LH} < V_{HL} = V_{HH}$), then bidding strategies and payoffs are unaffected. However, as bidder 2 begins to acquire information about the item’s value for which bidder 1’s signal is not a sufficient statistic, and hence V_{LH} and V_{HL} begin to differ from V_{LL} and V_{HH} , respectively, then bidder 2 gradually becomes more aggressive in two respects. First, as V_{LH} increases above V_{LL} , bidder 2’s minimum bid increases. Second, as V_{HL} decreases below V_{HH} , bidder 2 begins to place positive weight on a bid of V_{HH} . Equilibrium bidding, and bidder 2’s aggressiveness, vary continuously with these parameters from one extreme (a single informed bidder) to the other (ex ante symmetric bidders).

Next, we investigate how revenue varies with the information structure. An immediate corollary of Theorem 1 is a prediction about seller revenue in the unique TRE of the game. (Proofs of all corollaries are in Appendix C.)

Corollary 1 *The seller’s expected revenue under the unique TRE predicted by Theorem 1 is*

$$R_{SPA} = V_{LL} + (V_{HH} - V_{HL}) \Pr[H_1, H_2] \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} + (V_{LH} - V_{LL}) \Pr[H_1, H_2]. \quad (2)$$

Denote the ex ante expected value of the item as \bar{V} . This is the social surplus and would be the seller’s revenue if there were no asymmetric information—either because both bidders were uninformed or because both bidders were fully informed. With asymmetric information, we expect informed bidders to earn information rents, and hence for revenue to be below \bar{V} . At the same time, revenue should always be at least the minimal possible value of V_{LL} . It is an interesting question, however, where between these bounds revenue will fall.

4.1.1 Symmetric Case

Intuition suggests that if bidders have little information, then information rents could be low so that revenue is close to expected surplus \bar{V} . This intuition turns out to be correct if bidders are symmetric ex ante and cookies are rare, as shown in Corollary 2. By cookies being rare, we mean that that both bidders almost always receive the default signal of {no-cookie}, whether that be the low signal (such that $\Pr[L_1, L_2]$ is near 1) or the high signal (such that $\Pr[H_1, H_2]$ is near 1).

Corollary 2 *Given Assumption 1: If bidders are symmetric ex ante (such that $V_{LH} = V_{HL}$ and $\Pr[H_1, L_2] = \Pr[L_1, H_2]$) then SPA revenue in the unique TRE is:*

$$R_{SPA}^{symmetric} = \bar{V} - (\Pr[H_1, L_2] + \Pr[L_1, H_2])(V_{HL} - V_{LL}). \quad (3)$$

As cookies become rare, and either $\Pr[L_1, L_2]$ or $\Pr[H_1, H_2]$ approaches 1, revenue approaches the full expected surplus:

$$\lim_{\Pr[L_1, L_2] \rightarrow 1} R_{SPA}^{symmetric} = \lim_{\Pr[H_1, H_2] \rightarrow 1} R_{SPA}^{symmetric} = \bar{V}.$$

4.1.2 Peaches and Lemons Cases

To illustrate the implications of Corollary 1 beyond the symmetric case, we consider the following setting. There are two possible qualities for the item, low (L for Lemon) and high (P for Peach), that is $\Omega = \{L, P\}$. A peach is more valuable than a lemon, such that $v(L) < v(P)$. The ex ante expected value is \bar{V} . We then define two special cases:

Definition 4 *Both bidders are informed about peaches if $V_{LH} = V_{HL} = V_{HH} = v(P)$.*

If bidders are informed about peaches, our interpretation is the following: A cookie corresponds to the high signal and precisely identifies an item as a peach. Absence of a cookie corresponds to the low signal. If neither bidder has a cookie (both receive low signals), then Definition 4 implies an expected value for the item of:

$$V_{LL} = \bar{V} - (v(P) - \bar{V}) \frac{1 - \Pr[L_1, L_2]}{\Pr[L_1, L_2]}. \quad (4)$$

Definition 5 *Both bidders are informed about lemons if $V_{LL} = V_{LH} = V_{HL} = v(L)$.*

If both bidders are informed about lemons, our interpretation is the following: A cookie corresponds to the low signal and precisely identifies an item as a lemon. Absence of a cookie corresponds to the high signal. If neither bidder has a cookie (both receive high signals), then Definition 5 implies an expected value for the item of:

$$V_{HH} = \bar{V} + (\bar{V} - v(L)) \frac{1 - \Pr[H_1, H_2]}{\Pr[H_1, H_2]}. \quad (5)$$

Notice that while the assumption that bidders are informed about peaches or lemons imposes a symmetric mapping between signals and values (as $V_{LH} = V_{HL}$), it does not impose symmetry between bidders ex ante. In particular, bidders may still have asymmetric probabilities of receiving each signal ($\Pr[L_1, H_2] \neq \Pr[H_1, L_2]$). (In both cases, equation (1) labels bidders such that $\Pr[L_1, H_2] \geq \Pr[H_1, L_2]$.)

One intuition suggests that if bidders have little information, then information rents could be low so that revenue is close to expected surplus \bar{V} . Corollary 2 shows this to be correct if bidders are symmetric ex ante and cookies are rare. When bidders are asymmetric ex ante, however, there is a competing intuition—that we might expect low revenue because the less-informed agent bids low for fear of adverse selection. Investigating this possibility reveals a sharp distinction between the seemingly similar cases of information about peaches and lemons.

Peaches: If both bidders are informed about peaches, then Theorem 1 predicts that bidders 1 and 2 both bid $V_{HH} = v(P)$ when receiving a cookie (a high signal), and $V_{LL} \in (v(L), \bar{V})$ otherwise. Importantly, neither bidder faces an adverse selection problem conditional on receiving a high signal (a cookie), and can bid equally aggressively in this case. Moreover, absent a cookie, a bid of V_{LL} is still relatively close to the expected surplus of \bar{V} if cookies (high signals) are rare. Thus, when cookies are rare, revenues are close to the expected surplus of \bar{V} :

Corollary 3 *Given Assumption 1: If both bidders are informed about peaches then SPA revenue in the unique TRE is:*

$$R_{SPA}^{peaches} = \bar{V} - (v(P) - \bar{V}) \frac{\Pr[L_1, H_2] + \Pr[H_1, L_2]}{\Pr[L_1, L_2]}. \quad (6)$$

As cookies become rare and $\Pr[L_1, L_2]$ approaches 1, revenue approaches expected surplus:

$$\lim_{\Pr[L_1, L_2] \rightarrow 1} R_{SPA}^{peaches} = \bar{V}.$$

Lemons: If both bidders are informed about lemons, then Theorem 1 predicts that bidder 1 bids $V_{LL} = v(L)$ given a cookie (a low signal) and bids $V_{HH} \in (\bar{V}, v(P))$ otherwise. Bidder 2 also bids $v(L)$ given a cookie, but absent a cookie mixes between bidding V_{HH} and $V_{LH} = v(L)$. Unlike the peaches case, V_{LL} and V_{LH} are never close to \bar{V} , but rather both equal $v(L)$. Thus the seller only receives revenue above $v(L)$ when both bidders aggressively bid V_{HH} . However, in contrast to the peaches case, bidders with the high signal (meaning no cookie) now face a substantial adverse selection problem. Winning might imply that the other bidder was avoiding a known lemon. Hence bidder 2 bids less aggressively than in the peaches case—bidding $V_{LH} = v(L)$ with probability $1 - \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]}$ given the high signal. Corollary 4 shows the implications for revenue.

Corollary 4 *Given Assumption 1: If both bidders are informed about lemons then SPA revenue in the unique TRE is:*

$$R_{SPA}^{lemons} = v(L) + (\bar{V} - v(L)) \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]}. \quad (7)$$

Corollary 4 shows that SPA revenue varies with the probability, $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]}$, that bidder 2 aggressively bids V_{HH} upon receiving a high signal (no cookie). If bidders are symmetric ex ante, such that $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} = 1$, then unaggressive bidding due to adverse selection is not an issue and revenue achieves the upper bound: $R_{SPA}^{lemons} = \bar{V}$. However, when bidders are very asymmetric ex ante, such that $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \rightarrow 0$, bidder 2 stops competing for peaches entirely and revenue collapses to the lower bound: $R_{SPA}^{lemons} \rightarrow v(L)$. Thus, as $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]}$ varies between 0 and 1 (recall that bidders are labeled such that $\Pr[H_1, L_2] \leq \Pr[L_1, H_2]$), revenue varies from the lower bound $v(L)$ to the full expected surplus \bar{V} .

4.1.3 Peaches and Lemons Example

Next, consider an example that further illustrates the contrast between information about peaches and lemons.

Example 1 *The commonly known prior is that with probability 1/2 the impression is a peach (P), and with probability 1/2 it is a lemon (L). Normalize $v(L) = 0$ and $v(P) = 2$ so that $\bar{V} = 1$. We consider two cases:*

- *Both bidders are informed about peaches: Conditional on the impression being a peach, each bidder receives a cookie with independent probability p_i which identifies the impression as a peach. Otherwise the bidder receives no cookie and is uncertain whether the impression is a peach or a lemon. In this case a cookie corresponds to the high signal and lack of a cookie corresponds to the low signal.*
- *Both bidders are informed about lemons: Conditional on the item being a lemon, each bidder receives a cookie with independent probability q_i which identifies the impression as a lemon. Otherwise the bidder receives no cookie and is uncertain whether the impression is a peach or a lemon. In this case a cookie corresponds to the low signal and lack of a cookie corresponds to the high signal.*

We can translate the example into the notation of Theorem 1 as follows. If bidders are informed about peaches, then $V_{LH} = V_{HL} = V_{HH} = 2$, Bayes rule implies that $V_{LL} = 2 \frac{(1-p_1)(1-p_2)}{1+(1-p_1)(1-p_2)}$, and we label bidders such that $p_1 \leq p_2$.¹⁰ Similarly, if bidders are informed about lemons, then $V_{LH} = V_{HL} = V_{LL} = 0$, Bayes rule implies that $V_{HH} = \frac{2}{1+(1-q_1)(1-q_2)}$, and we label bidders such that $q_1 \geq q_2$.¹¹

¹⁰ Also, $\Pr[H_1, H_2] = \frac{1}{2}p_1p_2$, $\Pr[L_1, H_2] = \frac{1}{2}(1-p_1)p_2$, $\Pr[H_1, L_2] = \frac{1}{2}p_1(1-p_2)$, $\Pr[L_1, L_2] = \frac{1}{2} + \frac{1}{2}(1-p_1)(1-p_2)$.

¹¹ Also, $\Pr[L_1, L_2] = \frac{1}{2}q_1q_2$, $\Pr[H_1, L_2] = \frac{1}{2}(1-q_1)q_2$, $\Pr[L_1, H_2] = \frac{1}{2}q_1(1-q_2)$, $\Pr[H_1, H_2] = \frac{1}{2} + \frac{1}{2}(1-q_1)(1-q_2)$.

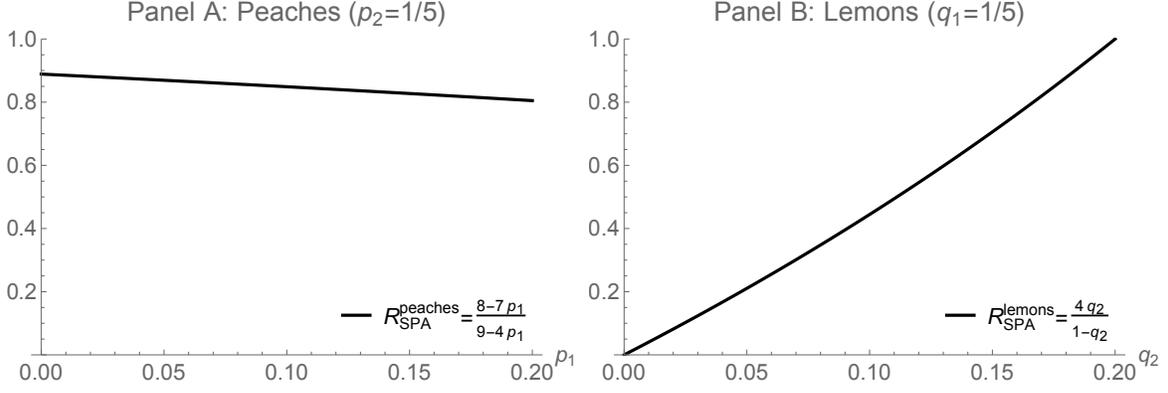


Figure 1: Panel A: SPA revenues in Example 1 when bidders are informed about peaches for $p_2 = 1/5$ and $p_1 \in [0, 1/5]$. Panel B: SPA revenues in Example 1 when bidders are informed about lemons for $q_1 = 1/5$ and $q_2 \in [0, 1/5]$.

Applying Corollary 1 yields predicted revenues of

$$R_{SPA}^{peaches} = \frac{2(1-p_1)(1-p_2) + p_1p_2}{1 + (1-p_1)(1-p_2)}, \quad (8)$$

$$R_{SPA}^{lemons} = \frac{q_2(1-q_1)}{q_1(1-q_2)}. \quad (9)$$

Revenues when bidders are informed about peaches are plotted as a function of $p_1 \in [0, 1/5]$ for the case of $p_2 = 1/5$ in Panel A of Figure 1. The right hand edge of the figure corresponds to ex ante symmetric bidders ($p_1 = p_2 = 1/5$), while the left-hand edge of the figure corresponds to only bidder 2 being informed ($p_1 = 0$ and $p_2 = 1/5$). The figure shows that expected revenues vary little with ex ante asymmetry of the bidders, and are always at least 80% of expected surplus—which holds for all values of $p_2 \leq 1/5$.

Revenues when bidders are informed about lemons are plotted as a function of $q_2 \in [0, 1/5]$ for the case of $q_1 = 1/5$ in Panel B of Figure 1. As both the figure and equation (9) make clear, revenues equal the full expected surplus of 1 when bidders are symmetric ex ante at the right-hand edge of the figure (where $q_1 = q_2 = 1/5$). However, moving left across the figure, as the bidders become increasingly asymmetric ex ante, revenues fall to zero as bidder 1 becomes the only informed bidder at the left-hand edge of the figure (where $q_1 = 1/5$ and $q_2 = 0$).

Comparing the peaches and lemons cases shows that the vulnerability of the second-price auction to adverse selection with ex ante asymmetric bidders varies greatly across the two types of information that cookies might contain. When cookies are relatively rare, revenues appear robust to the presence of bidders with ex ante better access to cookies that identify peaches, but revenues can collapse when cookies identify lemons.

4.2 Sketch of the Proof of Theorem 1

Next we sketch the proof of Theorem 1 and provide intuition for the result. For the complete proof see the formal outline in Appendix B and supporting details in Online Appendix H.

Fix any standard distribution R and $\epsilon > 0$ and let $\lambda(\epsilon, R)$ be the (ϵ, R) -tremble of the game. In the (ϵ, R) -tremble of the game the random bidder enters the auction with small probability $\epsilon > 0$ and is bidding according to a standard distribution R (its support is $[V_{LL}, V_{HH}]$). Denote the probability that the random bidder does not enter or enters but bids below x by $\hat{R}(x) = 1 - \epsilon + \epsilon \cdot R(x)$. Let its derivative, the density of random bids unconditional on entry, be $\hat{r}(x) = \epsilon \cdot r(x)$.

Note that in both the original game and any (ϵ, R) -tremble, the set of undominated bids is closed, so the requirement that bidders only bid within the closure of the set of undominated bids is equivalent to a requirement that bidders do not place dominated bids or that equilibrium is “in undominated bids”. The proof then relies on two results. (1) First, we show that in each of the games $\lambda(\epsilon, R)$ a mixed NE μ^ϵ in undominated bids exists (Lemma 3). (2) Second, we show that the limit of any sequence of NE μ^ϵ in undominated bids of the games $\{\lambda(\epsilon, R)\}_\epsilon$ must converge to μ as ϵ goes to zero (Lemma 4). As μ is a NE of the original game, these two results imply that it is the *unique* TRE.

We defer the first result to the appendix and next present the high level arguments for the second result. We first observe that if bidders never submit dominated bids, bidder $i \in \{1, 2\}$ that receives signal L_i must not bid outside the interval $[V_{LL}, v(L_i, H_j)]$, while bidder j that receives signal H_j must bid at least $v(L_i, H_j)$. As a result, a bidder i with signal L_i will never bid above V_{LL} because doing so means paying at least the item’s value (a lower bound for j ’s bid) and risks overpaying if the random bidder sets the price. Thus a bidder i with signal L_i always bids exactly V_{LL} .

Turning to bidder i ’s strategy given the high signal H_i , we first establish notation to describe the bidding strategies. Recall that μ^ϵ denotes a NE of the tremble $\lambda(\epsilon, R)$, and define $G_{Hi} = \mu_i^\epsilon(H_i)$ to be the cumulative distribution of bidder i ’s bids conditional on her receiving the signal H_i . When it exists, we denote the derivative of G_{Hi} by g_{Hi} .

If $V_{LH} = V_{HH}$, equation (1) implies that $V_{HL} = V_{HH}$, and hence it is a dominant strategy for each bidder i to bid V_{HH} conditional on receiving signal H_i . If $V_{HL} = V_{HH}$ but $V_{LH} < V_{HH}$, then bidder 1 has a dominant strategy to bid V_{HH} given signal H_1 . Bidder 2 must then bid V_{LH} given signal H_2 because all incremental wins from bidding higher would either be priced at their value (when bidder 1 sets the price at V_{HH}) or above their value (when the random bidder sets the price).

If $\max\{V_{LH}, V_{HL}\} < V_{HH}$, we show that bidding strategies conditional on high signals must fall into one of two cases. In both cases, the more cautious bidder (bidder 2) with signal H_2 bids an atom

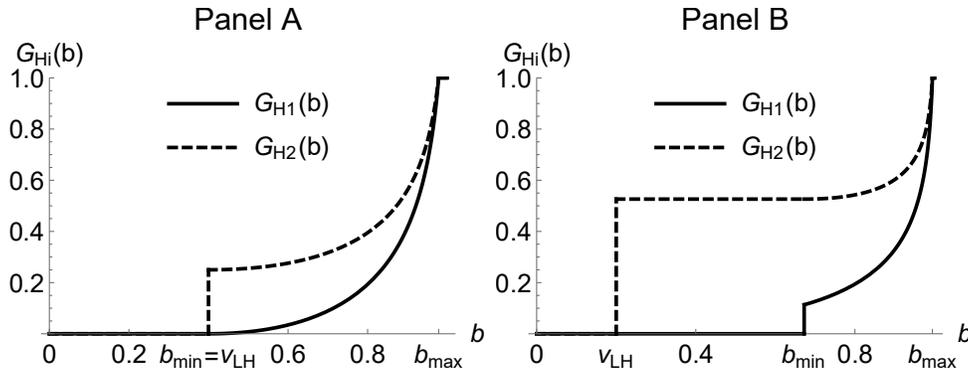


Figure 2: Two examples of the bidding CDFs for the two bidders when getting their high signals in the unique NE of the tremble $\lambda(\epsilon, R)$. Panel A: In this example, bidder 2 bids an atom at $V_{LH} = b_{\min}$, and both bidders mix over $(b_{\min}, b_{\max}]$. Panel B: In this example, bidder 2 bids an atom at V_{LH} , bidder 1 bids an atom at $b_{\min} > V_{LH}$, and both mix over $(b_{\min}, b_{\max}]$.

at V_{LH} (except in the special case of symmetric bidders in which there are no atoms). Moreover, in both cases, both bidders mix continuously over the interval (b_{\min}, b_{\max}) for some b_{\min} and b_{\max} satisfying $\max\{V_{LH}, V_{HL}\} \leq b_{\min} < b_{\max} < V_{HH}$ and there are no bids outside $[V_{LH}, b_{\max}]$. In the first case (illustrated in Figure 2 Panel A), there is no gap in bidding as $V_{LH} = b_{\min}$ and there are no atoms in the bid distribution of the aggressive bidder (bidder 1). In the second case (illustrated in Figure 2 Panel B), there is a gap in bidding between V_{LH} and an atom in the aggressive bidder's bid distribution at $b_{\min} > V_{LH}$.

In the second case, the aggressive bidder's atom serves to keep bidder 2 with signal H_2 indifferent between bidding V_{LH} and bidding just above b_{\min} . It is just the right size to provide a benefit for bidding above V_{LH} equal to the additional cost associated with overpaying due to a random bid falling between V_{LH} and b_{\min} when the aggressive bidder has the low signal L_1 . This cost goes to zero as ϵ goes to zero and the random bidder vanishes. Thus the aggressive bidder's atom at b_{\min} also vanishes as ϵ goes to zero, and is not part of the TRE.

The remainder of the result follows from considering bidder first-order conditions which apply over the interval (b_{\min}, b_{\max}) where both bidders mix continuously. In this interval, if bidder i has signal H_i , his bid b could be pivotal in one of three ways. First, a bid b could tie bidder j and beat the random bidder (an event with density $\Pr[H_j|H_i]g_{H_j}(b)\hat{R}(b)$), leading to a gain of $(V_{HH} - b)$. Second, a bid b could tie the random bidder and beat bidder j with a high signal H_j (an event with density $\Pr[H_j|H_i]\hat{r}(b)G_{H_j}(b)$), again leading to a gain of $(V_{HH} - b)$. Third, a bid b could tie the random bidder and beat bidder j with a low signal L_j (an event with density $\Pr[L_j|H_i]\hat{r}(b)G_{H_j}(b)$), leading to a loss from overpayment of $(b - E[V | H_i, L_j])$. The first-order

condition for b to be an optimal bid requires that these expected gains and losses from a slight bid change are equal so there is no benefit to raising or lowering the bid:

$$\Pr [H_j|H_i] \left(\hat{r}(b) G_{H_j}(b) + \hat{R}(b) g_{H_j}(b) \right) (V_{HH} - b) = \Pr [L_j|H_i] \hat{r}(b) (b - E[V | H_i, L_j]) \quad (10)$$

In the limit as ϵ goes to zero and the random bidder vanishes, a bid is only pivotal if it ties the strategic bidder. Thus the right-hand side of the first-order condition in equation (10) goes to zero and all bids in $(b_{\min}, b_{\max}]$ must approach V_{HH} . Note that this implies that, in the limit as ϵ goes to zero, the aggressive bidder 1 bids V_{HH} with probability 1. This follows because bidder 1's atom at b_{\min} vanishes so that in the limit all her bids fall in $(b_{\min}, b_{\max}]$. To determine the probability bidder 2 bids V_{HH} , we solve the differential equation given in equation (10) to find $G_{H_2}(b)$ for $\epsilon > 0$ and take the limit of $1 - G_{H_2}(V_{LH})$ as ϵ goes to zero (see Online Appendix H for details).

Next, we provide an informal intuition for the size of bidder 2's atom at V_{HH} . In the limit as ϵ tends to zero, all bidding mass in $(b_{\min}, b_{\max}]$ approaches V_{HH} . Thus, in the limit bidder j bids V_{HH} conditional on H_j with probability $\lim_{\epsilon \rightarrow 0} \int_{b_{\min}}^{b_{\max}} g_{H_j}(b) db$. Moreover, as bidder 1 has an atom of size 1, bidder 2's atom is equal to the ratio of the atoms:

$$\Pr(b_2 = V_{HH}|H_2) = \frac{\lim_{\epsilon \rightarrow 0} \int_{b_{\min}}^{b_{\max}} g_{H_2}(b) db}{\lim_{\epsilon \rightarrow 0} \int_{b_{\min}}^{b_{\max}} g_{H_1}(b) db} = \lim_{\epsilon \rightarrow 0} \frac{\int_{b_{\min}}^{b_{\max}} g_{H_2}(b) db}{\int_{b_{\min}}^{b_{\max}} g_{H_1}(b) db}. \quad (11)$$

The second equality above relies on the fact that $\lim_{\epsilon \rightarrow 0} \int_{b_{\min}}^{b_{\max}} g_{H_1}(b) db = 1 > 0$.

We solve the first-order condition from equation (10) for the bid density $g_{H_j}(b)$ and present the solution in equation (12). This characterizes the bid density of bidder j required for i with signal H_i to bid b :

$$g_{H_j}(b) = \frac{\Pr[L_j|H_i]}{\Pr[H_j|H_i]} \cdot \frac{\hat{r}(b)}{\hat{R}(b)} \cdot \frac{b - E[V | H_i, L_j]}{V_{HH} - b} - \frac{\hat{r}(b)}{\hat{R}(b)} \cdot G_{H_j}(b). \quad (12)$$

The first term in equation (12) is proportional the ratio of i 's potential loss from overpaying when bidder j has a low signal L_j to i 's potential gain from winning when bidder j has a high signal H_i . The second term is $\mathcal{O}(\epsilon)$ for all b , and hence it is unimportant for small ϵ . (In contrast the first term is large near b_{\max} as $\lim_{\epsilon \rightarrow 0} b_{\max} = V_{HH}$.) Substituting this expression into equation (11), while omitting the second term and cancelling $\hat{r}(b)/\hat{R}(b)$, gives bidder 2's atom at V_{HH} :

$$\Pr(b_2 = V_{HH}|H_2) = \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \lim_{\epsilon \rightarrow 0} \frac{\int_{b_{\min}}^{b_{\max}} \frac{b - V_{HL}}{V_{HH} - b} db}{\int_{b_{\min}}^{b_{\max}} \frac{b - V_{LH}}{V_{HH} - b} db} = \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}. \quad (13)$$

The second equality above follows from the fact that $\lim_{\epsilon \rightarrow 0} b_{\max} = V_{HH}$ and a result shown in Lemma 23 in the online appendix. Thus, bidder 2's atom at V_{HH} is proportional to the potential overpayment by bidder 1 bidding V_{HH} when bidder 2 has a low signal to the potential overpayment by bidder 2 bidding V_{HH} when bidder 1 has a low signal. Finally, bidder 2's atom at V_{LH} has complementary probability $1 - \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}$.

4.3 Comparison to FPA

For comparison, we now consider the FPA game in a slightly more restrictive setting than that studied in Section 4.1. In particular, in addition to Assumption 1, we assume that:

Assumption 2 *The signals are affiliated:* $\Pr[H_1, H_2] \Pr[L_1, L_2] \geq \Pr[L_1, H_2] \Pr[H_1, L_2]$.

In this setting, for the FPA game, we label bidders such that

$$\Pr[L_1, H_2] \geq \Pr[H_1, L_2]. \quad (14)$$

Note that this labeling of bidders coincides with our labeling in the SPA for Theorem 1 and Corollaries 1-4 if $V_{LH} = V_{HL}$ but not necessarily otherwise. Bidding strategies are monotone if bidder $i \in \{1, 2\}$ places higher bids given signal H_i than given signal L_i .¹²

Theorem 2 *Consider any FPA game with two bidders that each receive a binary signal that satisfies Assumptions 1-2, and label bidders as in equation (14). The unique Nash equilibrium with monotone bidding strategies is also the unique TRE with monotone bidding strategies and is characterized by equations (18)-(25) in Appendix A.*¹³

The equilibrium characterization in Appendix A shows that, conditional on signals received, equation (14) identifies bidder 1 as the more aggressive bidder and bidder 2 as the less aggressive bidder in the FPA game. This coincides with the finding in the SPA game if $V_{LH} = V_{HL}$, including the special cases in which bidders are informed either about peaches or about lemons, but not necessarily otherwise. (As in the SPA game, being the aggressive bidder does not necessarily mean being the bidder with a higher expected payoff.)

Next, the following corollary characterizes seller revenue in the unique TRE with monotonic strategies in the FPA game.

¹²To be precise, monotone bidding by i implies that if i bids more than b with positive probability given signal L_i then i must bid b or lower with zero probability given signal H_i . See also Definition 10 in Appendix A.

¹³Note that for the case $V_{LL} < \min\{V_{LH}, V_{HL}\} \leq \max\{V_{LH}, V_{HL}\} < V_{HH}$, Rodriguez's (2000) Proposition 1 implies that equilibrium bidding strategies must be monotone. If one bidder is entirely uninformed (e.g., bidder 2 is uninformed if $V_{LL} = V_{LH} < V_{HL} = V_{HH}$ and $\Pr[H_1|H_2] = \Pr[H_1|L_2]$) the uninformed bidder's signal realizations do not matter—only her unconditional bid distribution (Engelbrecht-Wiggans et al., 1983). While we are unaware of any Nash equilibria of the FPA game with non-monotone bidding strategies when both bidders are informed but $V_{LL} = \min\{V_{LH}, V_{HL}\}$ or $\max\{V_{LH}, V_{HL}\} = V_{HH}$, we have not ruled them out either. As a result, we simply choose to focus on the unique Nash equilibrium in monotone bidding strategies. Moreover, it should be noted that in the FPA game, TRE and Nash equilibrium coincide given monotone bidding strategies. Thus TRE is not helpful in refining the Nash prediction, but the fact that the unique Nash equilibrium in monotone bidding strategies is also a TRE ensures that we are comparing apples to apples when comparing to the TRE of the SPA game.

Corollary 5 *The seller's expected revenue under the equilibrium characterized by Theorem 2 is*

$$\begin{aligned}
R_{FPA} = & V_{LL} + \Pr[H_1, H_2] \frac{\Pr[H_1]}{\Pr[H_2]} (V_{HH} - V_{LH}) + \Pr[H_1, H_2] (V_{LH} - V_{LL}) \\
& + \frac{(\Pr[L_1, H_2])^2 - (\Pr[H_1, H_2] \Pr[L_1, L_2] - \Pr[H_1, L_2] \Pr[L_1, H_2])}{\Pr[L_1, H_2] \Pr[L_1] + (\Pr[H_1, H_2] \Pr[L_1, L_2] - \Pr[H_1, L_2] \Pr[L_1, H_2])} \\
& \cdot (\Pr[L_1, H_2] - \Pr[H_1, L_2]) (V_{LH} - V_{LL}) \quad (15)
\end{aligned}$$

As cookies become rare, and either $\Pr[L_1, L_2]$ or $\Pr[H_1, H_2]$ approaches 1, revenue approaches the full expected surplus:

$$\lim_{\Pr[L_1, L_2] \rightarrow 1} R_{FPA} = \lim_{\Pr[H_1, H_2] \rightarrow 1} R_{FPA} = \bar{V}.$$

The expression for revenue in the FPA is more cumbersome than that for revenue in the SPA. Nevertheless, it allows for an insightful comparison of revenue between the two auction formats. First, when bidders are ex ante symmetric, FPA revenue coincides with that in the SPA characterized in Corollary 2. This is consistent with Milgrom and Weber's (1982a) result that, given ex ante symmetric bidders and affiliated signals, revenue is equal or higher in the symmetric equilibrium of the SPA than in the FPA. Second, when cookies are rare, the finding that FPA revenue is always close to expected surplus can be compared with the characterization of SPA revenue when bidders are informed about peaches (Corollary 3) or lemons (Corollary 4).

Peaches Case: Consider the case in which both bidders are informed about peaches. In this case, Corollaries 3 and 5 show that both the SPA and the FPA perform similarly well when cookies are rare, both with revenues close to expected surplus.

Returning to Example 1, Panel A of Figure 3 illustrates this finding. The figure replicates plots of SPA revenue when bidders are informed about peaches (panel A) and lemons (panel B) first shown in Figure 1 with the addition of FPA revenue for comparison. As before, the left-hand side of each plot corresponds to the case in which only a single bidder is informed, while the right-hand side of each plot corresponds to ex ante symmetry.

Panel A shows that FPA and SPA revenue are so similar that they are hard to distinguish in the figure. Moreover, for the chosen parameter values at which cookies identify no more than 20% of peaches, both auction formats capture at least 80% of surplus as revenue.

Lemons Case: Next, consider the case in which both bidders are informed about lemons. In this case, comparing Corollaries 4 and 5 reveals an important difference between first-price and second-price auctions. If bidders are informed about lemons, FPA revenues are robust to bidder asymmetry when cookies are rare, always being close to expected surplus. In contrast, bidder asymmetry can

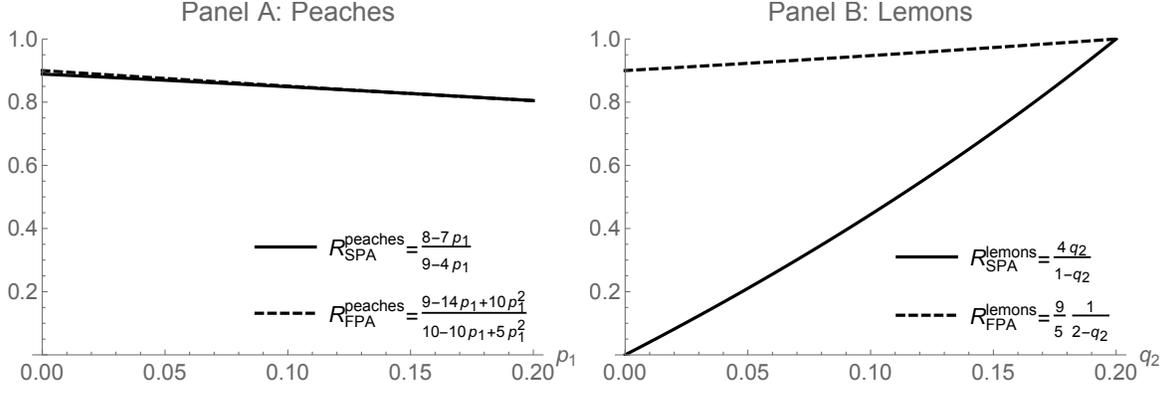


Figure 3: Panel A: SPA and FPA revenues in Example 1 when bidders are informed about peaches for $p_2 = 1/5$ and $p_1 \in [0, 1/5]$. Panel B: SPA and FPA revenues in Example 1 when bidders are informed about lemons for $q_1 = 1/5$ and $q_2 \in [0, 1/5]$.

be devastating to SPA revenue even when there is little private information because cookies are rare for both bidders. In fact, the two auction formats yield revenues at opposite bounds when bidders are informed about lemons, cookies are rare such that $\Pr[H_1, H_2]$ is large, and bidders are asymmetric such that $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]}$ is small. In such a setting, SPA revenues are close to the lower bound of $v(L)$, while FPA revenues are close to the full expected surplus of \bar{V} . In the context of Example 1, this is the scenario illustrated on the left-hand side of Panel B in Figure 3.

Summary: Figure 3 clearly summarizes our findings. Whether bidders are informed about peaches or about lemons, both SPA and FPA yield high revenues when bidders are symmetric ex ante. However, when cookies are rare (they arrive only 10% of the time for the better informed bidder in the example shown in Figure 3) bidder asymmetry has sharply different implications if bidders are informed about peaches than if they are informed about lemons. When bidders are informed about peaches, SPA revenues are nearly identical to FPA revenues and close to full surplus whether bidders are asymmetric or not. When bidders are informed about lemons, however, SPA revenues collapse to the value of a lemon, $v(L)$, as bidder asymmetry widens, while FPA revenues remain robust. As a result, whether online-advertising impression marketplaces are losing substantial revenue due to bidder asymmetry by running SPA rather than FPA depends importantly on what information is contained in bidders' cookies. If cookies identify peaches, then the loss may be minimal. If cookies identify lemons, however, the loss could be substantial.

4.4 Extension: Only One Informed Bidder

We now consider the special case in which only a single bidder is informed, but allow for multiple uninformed bidders and for the informed bidder to receive a signal with many possible realizations rather than only two. Our findings about the important distinction between private information about lemons and peaches are robust in this extension. (Proofs are in Appendix D.)

Suppose that only a single *informed buyer* I receives an informative signal $s_I \in S_I$ about the value, while all $n \geq 0$ others are *uninformed buyers* (always receiving a null signal). Thus we write $v(s_I) = E[v(\omega)|s_I]$ for the informed bidder's interim expected value conditional on s_I and denote the minimum and maximum such values by v_{min} and v_{max} , respectively. Throughout this section we assume that $0 \leq v_{min} < v_{max} < \infty$.

When only one bidder is informed, Theorem 1 implies that: (1) the informed bidder bids the expected value conditional on her signal (either $v(L)$ or $v(H)$); (2) the uninformed bidder bids to match the informed bidders' lowest bid, her minimum possible posterior valuation, $v(L)$; (3) revenues are equal to that minimum posterior valuation. Theorem 3 establishes that all three results apply more generally, for any number of uninformed bidders and any signal distribution of the informed bidder.

Theorem 3 *In any common-value domain with one informed buyer and $n \geq 0$ uninformed buyers, the unique TRE of the SPA game is a strong TRE in pure strategies in which:*

1. *the informed buyer bids $b_I(s_I) = v(s_I)$.*
2. *each of the uninformed buyers bids the informed bidders minimum possible expected value,*

$$b_U = \min_{\hat{s}_I \in S_I} v(\hat{s}_I) = v_{min}.$$

Moreover, revenue is $R_{SPA}^{1-informed} = v_{min}$.

Theorem 3 shows that the revenue of the SPA with only one informed bidder in the unique TRE is as low as it can be with undominated bids. When the informed bidder receives a signal that identifies a lemon with positive probability, this means that revenues can be very low. In the setting of the previous section, we found FPA revenues to be more robust. This insight can also be extended to allow for any number of uninformed bidders and a general signal distribution for the informed bidder. In particular, using the FPA revenue result in Theorem 4 of Engelbrecht-Wiggans et al. (1983), FPA revenues can be bounded below, independent of the signal distribution of the informed bidder.

Proposition 1 Consider any common-value domain with one informed buyer and $n \geq 0$ uninformed buyers that satisfies $0 \leq v_{min} < v_{max} < \infty$. There exists a TRE of the FPA. Letting $\bar{V} = E[v]$ and F be the cumulative distribution of $v(s_I)$, TRE implies that (1) $R_{FPA}^{1-informed} = \int_0^\infty (1 - F(v))^2 dv > R_{SPA}^{1-informed}$, and (2)

$$\bar{V} \geq R_{FPA}^{1-informed} \geq v_{min} + \frac{(\bar{V} - v_{min})^2}{v_{max} - v_{min}}.$$

Comparing Proposition 1 with Theorem 3 bounds the difference in FPA and SPA revenues when only one bidder is informed:

$$\bar{V} - v_{min} \geq R_{FPA}^{1-informed} - R_{SPA}^{1-informed} \geq \frac{(\bar{V} - v_{min})^2}{v_{max} - v_{min}}. \quad (16)$$

This difference can be large when the informed bidder receives a signal that identifies a lemon with positive probability and $v_{min} = v(L) < \bar{V}$. However, the difference is negligible when cookies are rare and always correspond to above average impressions. In this case, absence of a cookie is both the only negative signal and not very informative so that $v_{min} \approx \bar{V}$.

At the opposite extreme, we can consider the case of ex ante symmetric bidders with affiliated signals. For a SPA with two bidders and binary signals, we found in Section 4.1 that the TRE coincided with the symmetric equilibrium when bidders are ex ante symmetric. Under the conjecture that this is true with more than two bidders and richer information structures, Milgrom and Weber’s (1982a) result ranking second-price auction revenue equal or higher than first-price auction revenue applies:¹⁴

$$R_{FPA}^{symmetric} \leq R_{SPA}^{symmetric}. \quad (17)$$

Comparing equations (16)-(17) shows that from the seller’s perspective, while SPA perform well in symmetric settings, sufficient asymmetry leads the SPA to substantially underperform the FPA if an informed bidder sometimes receives signals that cause a low posterior valuation. Thus, this insight which was first shown in Section 4.3, appears to be much more general than the 2-bidder and binary signal case—but rather applies to any number of bidders with rich information structures.

4.4.1 Comparison to Akerlof (1970)

The prediction of TRE in a SPA with one informed bidder bears a striking similarity to Nash equilibrium in Akerlof’s (1970) lemons market: In both cases, uninformed buyers only buy “lemons”, paying no more than the value of a lemon, and similar adverse selection phenomena drive both results. Nevertheless, there is an important difference between the two cases. In Akerlof’s (1970)

¹⁴Milgrom and Weber’s (1982a) result is proved for continuous signals, but the authors point out in footnote 15 that it is true more generally.

market for lemons there is a single market price, which therefore cannot be commensurate with quality for all items. In stark contrast, in the SPA with a single informed bidder, information is revealed about the quality of the item during the auction. When the informed bidder sets the price, it reflects the fair value of the item given its quality. Hence, there are Nash equilibria in which uninformed buyers bid high prices and win high quality items at high prices as well as low quality items at low prices. We believe the fact that our refinement excludes such equilibria and selects one in the spirit of Akerlof's (1970) market for lemons is an attractive feature of the refinement. However, this was not a foregone conclusion that could be reached without the TRE refinement.

4.5 Extension: Many Agents, each with Finitely Many Signals

Characterizing TRE of the SPA with N bidders and an arbitrary information structure is beyond the scope of this paper. Instead, in this final extension, we characterize TRE in the SPA game with N bidders that each receive finitely many signals that jointly satisfy the *strong-high-signal property*. We therefore begin by recursively defining the strong-high-signal property, and providing examples in which it holds. Next, Theorem 4 characterizes the unique TRE when the strong-high-signal property holds. In particular, the profile of strategies in which each agent bids the posterior given his signal and the *worst* feasible combination of signals of the others is a strong TRE in pure strategies and the unique TRE. Finally, Propositions 2-3 apply the result to make revenue predictions for settings in which bidders are informed about peaches or bidders are informed about lemons. We find again that SPA revenues can be much lower when bidders are informed about lemons than when informed about peaches.

4.5.1 Strong-High-Signal Property

Recall that we denote the vector of all bidder signals by s . Extending notation from the previous section, $v_{min} = \min_{s \in S} \{v(s)\}$ and $v_{max} = \max_{s \in S} \{v(s)\}$ are the minimal and maximal possible values conditional on any feasible signal profile, respectively.¹⁵ In every domain satisfying the strong-high-signal property, there exists a signal s_i for some agent i such that his interim expected value $v(s_i) \equiv E[v(\omega)|s_i]$ is equal to v_{max} . Such a signal is strong in the sense that it is a sufficient statistic for the value. (Conditional on s_i , other signals $s_{j \neq i}$ are uninformative.) Such a signal is also high in the sense that no other information set could lead to a higher expected value. Moreover, if we condition on that signal *not* being realized and consider the domain with that restriction, that

¹⁵Thus v_{min} and v_{max} are the infimum and supremum undominated bids for any bidder i and signal $s_i \in S_i$, as defined in Section 3.2

domain also satisfies the condition.¹⁶ (Any domain in which all agents are uninformed satisfies the strong-high-signal property.)

Definition 6 Consider a common-value domain with n agents, each with finitely many signals. We say that such a domain satisfies the strong-high-signal property if: (1) For some agent i and signal $s_i \in S_i$ it holds that $v(s_i) = v_{max}$, and (2) if we consider the same domain but restricted to the case that agent i does not receive the signal s_i , if that restricted domain contains any feasible vector of signals then it also satisfies the strong-high-signal property.

A variety of information structures satisfy the strong-high-signal property. First, any domain with a single informed bidder (as in Section 4.4) satisfies the strong-high-signal property.¹⁷ Second, the case of two-bidders with binary signals (as in Section 4.1) satisfies the strong-high-signal property if and only if $V_{HL} = V_{HH}$ or $V_{LH} = V_{HH}$.¹⁸ (This includes the case of cookies identifying peaches from Section 4.1, but not the case of cookies identifying lemons if both bidders are informed.) Moving beyond these cases already covered by our earlier results, Online Appendix G shows that any *connected domain* satisfies the strong-high-signal property (but that the converse does not hold). In a connected domain, each bidder's information is a partition of the interval of common values $[\min_{\omega} v(\omega), \max_{\omega} v(\omega)]$. Formally, a connected domain is defined as follows:

Definition 7 Let each agent i have a partition Π_i of the set of states Ω and receive a signal that is the element of the partition that includes the realized state. The information partition Π_i of bidder i is connected (with respect to the common value v) if every $\pi_i \in \Pi_i$ has the property that, when $\omega_1, \omega_2 \in \pi_i$ and $v(\omega_1) \leq v(\omega_2)$ then every $\omega \in \Omega$ with $v(\omega_1) \leq v(\omega) \leq v(\omega_2)$ is necessarily in π_i . A common-value domain is connected (with respect to the common value) if the information partition Π_i is connected for every agent i .

4.5.2 Tremble Robust Equilibrium

Theorem 4 characterizes the unique TRE of the SPA given the strong-high-signal property; its proof is in Appendix E.

¹⁶By definition, v_{max} is a function of the domain. When we remove a high signal, its value falls.

¹⁷This follows because at each point one can take the signal with the highest posterior value for the informed agent and recursively remove it.

¹⁸If $V_{HL} = V_{HH}$ then $v(H_1) = v(H_1, H_2) = v_{max}$, meaning that H_1 is a strong high signal. To prove that the property holds we only need to consider the domain restricted to agent 1 receiving L_1 . But that domain clearly satisfies the property as it has at most one informed bidder (bidder 2). The case $V_{LH} = V_{HH}$ also satisfies the strong-high-signal property by symmetric logic.

Theorem 4 Consider a SPA in a common-value domain with n agents, each with finitely many signals, in which the strong-high-signal property holds. Let μ be the profile of strategies in which agent i with signal $s_i \in S_i$ bids the minimal value consistent with her signal s_i :

$$v_{\min}(s_i) \equiv \min\{v(s_i, s_{-i}) \mid s_{-i} \text{ such that } (s_i, s_{-i}) \in S\}.$$

Then μ is the unique TRE and moreover, μ is a strong TRE in pure strategies.

Theorem 4 has significant implications regarding the revenue of the seller in the unique TRE. In this TRE, each bidder bids the posterior given his signal and the *worst* feasible combination of signals of the others, which can be much lower than the interim valuation given only the bidder’s signal. We further explore implications for revenue below.

Notice that for the connected domains studied by Einy et al. (2002), Theorem 4 applies and the unique TRE coincides with the Nash equilibrium highlighted by Einy et al. (2002) as the single “sophisticated equilibrium” that Pareto-dominates the rest in terms of bidders resulting utilities. In other words, the two refinements coincide on connected domains. Importantly, in Online Appendix G we give an example of a simple domain that satisfies the strong-high-signal property but is not connected, and also is not equivalent to any connected domain. This shows that Theorem 4 applies in more settings than do Einy et al.’s (2002) results.

4.5.3 Lemons versus Peaches

In this section, Propositions 2 and 3 contrast revenue consequences of cookies that identify various quality peaches with those of cookies that identify various quality lemons. (Proofs are in Appendix E.) Our definitions of what it means for cookies to identify lemons or peaches are adapted for the setting of multiple agents, multiple signals, and the strong-high-signal property. Nevertheless, they remain similar in spirit to those used in previous sections.

Consider a common-value domain satisfying the strong-high-signal property with items of value in $[0, 1]$ and expected value of \bar{V} . Assume that there are n agents, each receiving a signal s_i from an ordered, finite set of signals S_i . Let L_i and H_i denote the lowest and highest signals of agent i , respectively. We also assume that the domain is monotonic (meaning that the common value is nondecreasing in each bidder’s signal):

Definition 8 Let $s \leq s'$ if for every i it holds that $s_i \leq s'_i$. A common-value domain is monotonic if, for every pair of feasible signal vectors, the comparison $s \leq s'$ implies that $v(s) \leq v(s')$.

We define an agent i to be slightly *informed about peaches* if his *non-peaches* signal L_i occurs with probability close to 1 (as the cookies that identify peaches are rare). Further, we define an

agent i to be slightly *informed about lemons* if (1) her *non-lemons* signal H_i occurs with probability close to 1 and (2) lemons signals indicate that the value is close to zero (meaning cookies identifying lemons are rare but informative). Formal definitions are as follows:

Definition 9 Fix any $\epsilon_i \geq 0$.

- Bidder i is ϵ_i -informed about peaches, if $\Pr[s_i \neq L_i] \leq \epsilon_i$.
- Bidder i is ϵ_i -informed about lemons, if (1) $0 < \Pr[s_i \neq H_i] < \epsilon_i$, and (2) for any $s_i \in S_i \setminus \{H_i\}$, if (s_i, s_{-i}) is feasible then $v(s_i, s_{-i}) < \epsilon_i$.

If all n agents are slightly informed about peaches, then SPA revenue in the unique TRE is close to social surplus \bar{V} .

Proposition 2 Fix any nonnegative constants $\epsilon_1, \epsilon_2, \dots, \epsilon_n$. Consider any monotonic domain for which (1) $v(\omega) \in [0, 1]$, (2) the strong-high-signal property holds, and (3) every agent $i \in \{1, 2, \dots, n\}$ is ϵ_i -informed about peaches. In the unique TRE, SPA revenue is at least

$$R_{SPA}^{\epsilon\text{-peaches}} \geq \bar{V} - \sum_{j=1}^n \epsilon_j$$

In contrast to the previous result, Proposition 3 implies (as a special case) that when one or more bidders are slightly informed about lemons and the rest are slightly informed about peaches, then revenue will be close to zero (as long as some non-degeneracy conditions are satisfied).

Proposition 3 Fix $n \geq i$ and positive constants $\epsilon_1, \epsilon_2, \dots, \epsilon_i$. Consider any monotonic domain with n bidders for which: (1) $v(\omega) \in [0, 1]$, (2) the strong-high-signal property holds, (3) each agent $j \in \{1, 2, \dots, i-1\}$ is ϵ_j -informed about peaches, (4) agent i is ϵ_i -informed about lemons, and (5) the following non-degeneracy conditions hold:

- For any $j < i$, the signal L_j does not imply H_i (alternatively, $(L_j, s_i, s_{-\{i,j\}})$ is feasible for some $s_i \neq H_i$ and some $s_{-\{i,j\}}$).
- For any $j > i$ and any signal $s_j \in S_j$, the signal s_j does not imply H_i (alternatively, $(s_j, s_i, s_{-\{i,j\}})$ is feasible for some $s_i \neq H_i$ and some $s_{-\{i,j\}}$).

Then SPA revenue in the unique TRE is at most

$$R_{SPA}^{\epsilon\text{-lemons}} \leq \epsilon_i + \sum_{j=1}^i \epsilon_j$$

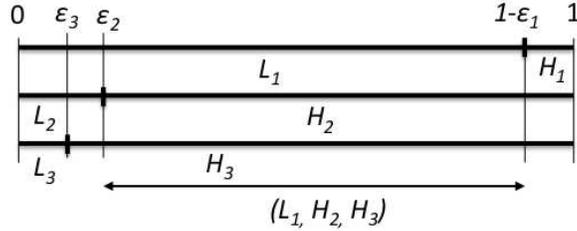


Figure 4: A simple example illustrating Proposition 3

The non-degeneracy conditions rule out the case of ex ante symmetric bidders. Thus comparing Propositions 2 and 3 yields a similar conclusion to that with two-bidders in Section 4.1. When cookies are relatively rare, revenues appear robust to the presence of bidders with ex ante better access to cookies that identify peaches, but revenues can collapse when cookies identify lemons. The difference is in the amount of ex ante asymmetry required for revenues to collapse. In the lemons case of the two-bidder model in Section 4.1, revenues decline smoothly as ex ante asymmetry widens. In contrast, in the current setting with the strong-high-signal property, even a slight asymmetry is sufficient to collapse revenues.

To illustrate Proposition 3, consider the domain illustrated in Figure 4 for which the proposition applies. The item's value v is sampled uniformly from $[0, 1]$. Each agent j has a different threshold t_j : he gets signal H_j if $v \geq t_j$ and signal L_j otherwise. It holds that $0 < t_3 = \epsilon_3 < t_2 = \epsilon_2 < t_1 = 1 - \epsilon_1 < 1$. Agent 1 is ϵ_1 informed about peaches, while agents 2 and 3 are ϵ_2 and ϵ_3 informed about lemons, respectively. It is easy to verify that the non-degeneracy conditions hold.¹⁹ Proposition 3 applies for $i = 2$ and implies that the revenue is at most $\epsilon_1 + 2\epsilon_2$ by the following argument. As illustrated in Figure 4, the signal profile (L_1, H_2, H_3) occurs if the value is between ϵ_2 and $1 - \epsilon_1$, an event that occurs with probability $1 - (\epsilon_1 + \epsilon_2)$. Therefore, a combination of signals other than (L_1, H_2, H_3) happen with probability $\epsilon_1 + \epsilon_2$ and as $v \leq 1$ it contributes at most $\epsilon_1 + \epsilon_2$ to the expected revenue. The signal combination (L_1, H_2, H_3) occurs with probability smaller than 1. While the bid of agent 2 in that case is high (almost $1/2$), both agent 1 and 3 bid at most ϵ_2 with signals L_1 and H_3 , respectively (as they never win when agent 2 gets signal H_2). The contribution to the expected revenue in this case is thus bounded by ϵ_2 . We conclude that the total revenue is at most $(\epsilon_1 + \epsilon_2) + \epsilon_2$.

The example in Figure 4 can be generalized to allow for many agents and many signals for each, as follows. The item's value v is sampled uniformly from $[0, 1]$. Each agent j has an increasing

¹⁹Non-degeneracy fails when $\epsilon_3 = \epsilon_2$ as the combination (L_1, L_2, H_3) becomes infeasible. It is easy to see that in this case the result fails as, on the likely profile (L_1, H_2, H_3) , both agent 2 and 3 are bidding high.

list of $k_j + 1$ thresholds satisfying $0 = t_j^0 < t_j^1 < t_j^2 < \dots < t_j^{k_j} = 1$, and his signal indicates the interval between two consecutive thresholds of his that includes the realized value. Fix $i \leq n$. The condition that every agent $j < i$ is ϵ_j -informed about peaches is satisfied when $t_j^1 > 1 - \epsilon_j$. The condition that agent i is ϵ_i -informed about lemons is satisfied when $t_i^{k_i-1} < \epsilon_i$. Every agent $j > i$ is also ϵ_i -informed about lemons when $t_i^{k_i-1} > t_j^{k_j-1}$. For such an agent j , the value conditional on his best signal is not as high as the value conditional on i 's best signal (this captures the second non-degeneracy condition). It is easy to verify that the first non-degeneracy condition is satisfied for any such a domain. Proposition 3 states that the revenue is at most $\epsilon_i + \sum_{j=1}^i \epsilon_j$. The seller's revenue is low even though with high probability (at least $1 - \sum_{j=1}^i \epsilon_j$) agent i gets signal H_i and is bidding relatively high (at least $(1 - \epsilon_i - \max_{j < i} \epsilon_j)/2$). The revenue is low as all other agents are bidding low (at most ϵ_i) and thus the second highest bid is also low.

5 Discussion: Mechanism Design

The previous section shows that in the common-value model the revenue of the SPA may be only a small fraction of total welfare. In this section we consider how to maximize the seller's revenue.

In the common-value model there is a trivial mechanism that is ex-ante individually rational and maximizes both welfare and revenue: Before signals are realized, make the first buyer a take-it-or-leave-it offer to buy the item for the price equal to the unconditional expected value of the item.

Unfortunately, this trivial mechanism does not extend to cases with a private component to the item's value. For example, in online advertising markets it is reasonable to assume that an informed buyer (advertiser) that has an accurate signal about the user (from a cookie on the user's machine) can tailor a specific advertisement to the specific user, generating some additional value over the common value created by placing a generic advertisement that is not user specific.

This motivates us to consider the following generalization of the model with a single informed bidder, in which the informed bidder is also *advantaged*. In this model there are n potential buyers. One random buyer i is informed about the state of the world (gets a signal $s_i \in S_i$), while the others are uninformed.²⁰ Signals are ordered by the expected common value to an uninformed bidder. Given the maximal signal s_{max} , the value for the informed bidder is larger than the common value

²⁰McAfee (2011) considers a related pure common-value model in which the probability of being informed is independent across bidders. Following an axiomatic approach, Arnosti, Beck, and Milgrom (2015) show that *modified second bid* auctions achieve highly efficient allocations in a related model with both common and idiosyncratic components to bidder values and a single *uninformed* bidder.

by a bonus $B > 0$. For other signals there is no bonus.²¹

Let E be the unconditional expected value of the item to an uninformed bidder, L be the expected value of the item conditional on the lowest signal s_{min} , and p_{max} be the probability of the highest signal s_{max} . The expected social welfare when the realized informed bidder always gets the item is $E + p_{max}B$. In this model selling the item ax-ante to a fixed agent at his expected value will generate revenue of $E + \frac{p_{max}B}{n}$, which can be significantly lower than the maximal social welfare.

The unique TRE of the second-price auction in this scenario is efficient. Yet, one can easily extend Theorem 3 to this model and see that for any realized informed bidder the unique TRE in this model is exactly the same as the one described by the theorem (with the adjustment that the informed bidder with signal s_{max} bids his value that includes the bonus). Thus revenues may fall far short of capturing total surplus.

Nevertheless, using an auction entry fee, we can build a mechanism that is ex-ante individually rational, is socially efficient, and extracts (almost) the entire welfare as revenue. This is true in the mechanism's *unique* TRE, as we explain below.

The mechanism has two stages. First bidders choose whether to pay an auction entry fee. Second, those who have paid the entry fee compete in a SPA. Theorem 3 (and its extension to this model) predicts a *unique* TRE in the SPA subgame. The payment in the SPA is always L . The SPA entry fee is set to be slightly less than the expected utility that an agent gets by participating, assuming all agents participate in the SPA and bid according to the *unique* TRE in that subgame. Thus the entry price is set to be slightly less than $(E + p_{max}B - L)/n$.

As TRE provides a *unique* prediction to the outcome of the second stage, agents have a *unique* rational decision when facing the entry decision, and they choose to pay the entry fee. Thus, in the *unique* subgame-perfect-equilibrium that uses the TRE refinement, agents will all choose to pay the entry fee and the SPA allocation will be socially efficient. Although the revenue in the SPA is low, essentially the entire expected utility an agent gets from this auction is charged as an entry fee. The revenue from entry would be (almost) $n(E + p_{max}B - L)/n = E + p_{max}B - L$, while the revenue in the SPA would be L . Thus the total revenue is (almost) the social welfare $E + p_{max}B$.

The above mechanism can only be used when both seller and agents can reasonably predict the outcome of the SPA that takes place at the second stage, for which the unique TRE prediction is potentially helpful. The mechanism can be extended to any other scenario in which a uniqueness result can be proven about the outcome of the SPA game under some solution concept.

²¹In this extension the advantaged bidder is not known ex ante. In contrast, the literature on *almost-common-value* auctions assumes that one bidder is known ex ante to value the object slightly more than other bidders (Bikhchandani, 1988; Avery and Kagel, 1997; Klemperer, 1998; Bulow, Huang, and Klemperer, 1999; Levin and Kagel, 2005).

Interim Individually Rational Mechanism

While the entry fee mechanism is ex ante individually rational, it is not interim individually rational once bidders have received their signals. We next design an interim individually rational mechanism for this setting, when the informed player has only two signals s_{min} and s_{max} . Our mechanism is dominant strategy incentive compatible. Let L be the value conditional on s_{min} and $P + B$ be the value of the advantaged bidder conditional on s_{max} .

While our model is not one of independent private values, it is sufficiently close that it seems useful to consider the optimal auction when each player’s value is sampled independently and identically from the following distribution: the value is L with probability $1 - 1/n$, and $P + B$ with probability $p_{max} = 1/n$. For this instance, Myerson’s optimal auction is to have some reserve price r and some floor price f . If some bidders bid at least r then we run a second-price auction with reserve r , otherwise we randomly choose a winner among those who bid at least f and charge the winner f .²²

In our advantaged bidder model, we propose using this mechanism with $f = L$ and $r = P + B - z$, where $z = (P + B - f)/n$ is the expected utility of agent i bidding f given signal s_{max} (conditional on every other agent j bidding f). The revenue obtained is $(1 - p_{max})f + p_{max}r$. Note that this is at least $(1 - 1/n)$ -fraction of the efficient social welfare which is $(1 - p_{max})L + p_{max}(P + B)$.

6 Conclusion

This paper analyzes the impact of ex ante information asymmetries in second-price and first-price common-value auctions, in an environment that captures key features of real world markets such as those for online advertising. In these environments, bidders may be asymmetrically informed at the interim stage—as some receive informative signals (called “cookies” in online advertising markets) while others do not. Moreover, bidders may be asymmetric ex ante, with some much more likely to receive a signal than others. The type of information contained in these signals may be qualitatively different across settings. For example, in some settings bidders may have access to cookies which occasionally reveal that a potential advertising viewer is a “robot” rather than a real person, or that the asset for sale is a “lemon” with no value. Alternatively, in other settings cookies might identify past customers who will be responsive to advertising, or that the asset for sale is a “peach” with high value.

In these settings, both Bayesian Nash Equilibrium and a number of standard refinements have

²²This mechanism is similar to the *buy-it-now or take-a-chance* mechanism proposed by Celis, Lewis, Mobius, and Nazerzadeh (2014).

limited predictive power for SPA revenues. We make progress by introducing the Tremble Robust Equilibrium refinement. This selects a Nash equilibrium that is robust to a vanishingly small probability that an additional bidder enters the auction and bids randomly over the support of valuations. Applying our refinement, we show that SPA revenues can be particularly vulnerable to ex ante asymmetry between bidders, even when those bidders are rarely informed. Whether this is true, however, depends on the type of information that signals contain—SPA revenues suffer substantially from ex ante asymmetry with respect to information about lemons, but not with respect to information about peaches. In contrast, if bidders are rarely informed, FPA revenues are close to expected surplus regardless of the details.

From a market design perspective, our findings suggest that auctioneers running second-price auctions should think carefully about enabling information structures that allow for some bidders to learn about lemons with substantially higher probability than others. For instance, if restricting access to cookies in an online advertising marketplace is unreasonable, a seller might consider identifying and publicly disclosing web robots and other lemons itself. Alternatively, if ex ante asymmetry with respect to information about lemons cannot be avoided, sellers may consider running first-price auctions rather than second-price auctions, as we predict they will yield substantially higher revenue in those circumstances. These insights may be particularly relevant for markets such as that for online advertising where second-price auctions are widely used and common practice allows for substantial ex ante informational asymmetry between bidders.

A Details of FPA Equilibrium

We denote i 's bidding distribution conditional on signal S_i by G_{S_i} .

Definition 10 Bidder i 's strategy is monotone if $G_{L_i}(b) < 1$ implies $G_{H_i}(b) = 0$.

Theorem 5 Consider any FPA game with two bidders that each receive a binary signal that satisfies Assumptions 1–2, and label bidders as in equation (14). For each of two cases, the following characterizes the unique Nash equilibrium with monotone bidding strategies, which is also the unique TRE with monotone bidding strategies.²³

(1) $V_{LH} > V_{LL}$ and $\Pr[L_1, H_2] > \Pr[H_1, L_2]$: Bidder 1 bids over the interval $[V_{LL}, b^*]$ with distribution $G_{L_1}(b)$ given signal L_1 and bids over the interval $[b^*, \bar{b}]$ with distribution $G_{H_1}(b)$ given signal H_1 . Bidder 2 bids V_{LL} given signal L_2 and bids over the interval $[V_{LL}, \bar{b}]$ with distribution $G_{H_2}(b)$ given signal H_2 . Critical value b^* and maximum bid \bar{b} satisfy $V_{LL} < b^* < \bar{b} < V_{HH}$ and

²³See footnote 13.

$b^* < V_{LH}$. These values and the bidding distributions are described by equations (18)-(22).

$$b^* = \frac{V_{LH} \Pr[L_1, H_2] (\Pr[L_1, H_2] - \Pr[H_1, L_2]) + V_{LL} \Pr[L_1, L_2] \Pr[H_2]}{\Pr[L_1, H_2] (\Pr[L_1, H_2] - \Pr[H_1, L_2]) + \Pr[L_1, L_2] \Pr[H_2]} \quad (18)$$

$$\bar{b} = (1 - \Pr[H_1|H_2]) b^* + \Pr[H_1|H_2] V_{HH} \quad (19)$$

$$G_{L1}(b) = \frac{V_{LH} - b^*}{V_{LH} - b}, \quad b \in [V_{LL}, b^*] \quad (20)$$

$$G_{H1}(b) = \frac{\Pr[L_1, H_2]}{\Pr[H_1, H_2]} \frac{b - b^*}{V_{HH} - b}, \quad b \in [b^*, \bar{b}] \quad (21)$$

$$G_{H2}(b) = \begin{cases} \frac{\Pr[L_1, L_2]}{\Pr[L_1, H_2]} \frac{b - V_{LL}}{V_{LH} - b}, & b \in [V_{LL}, b^*] \\ \frac{\Pr[L_1, L_2]}{\Pr[L_1, H_2]} \frac{V_{HH} - b^*}{V_{LH} - b^*} \frac{b - V_{LL}}{V_{HH} - b} + \frac{\Pr[H_1, L_2]}{\Pr[H_1, H_2]} \frac{b - b^*}{V_{HH} - b}, & b \in [b^*, \bar{b}] \end{cases} \quad (22)$$

(2) $V_{LH} = V_{LL}$ or $\Pr[L_1, H_2] = \Pr[H_1, L_2]$: Bidder $i \in \{1, 2\}$ bids V_{LL} given signal L_i and bids over the interval $[V_{LL}, \bar{b}]$ with distribution $G_{Hi}(b)$ given signal H_i . Maximum bid $\bar{b} \in (V_{LL}, V_{HH})$ and bidding distributions $G_{Hi}(b)$ are described by equations (23)-(25).

$$\bar{b} = (1 - \Pr[H_1|H_2]) V_{LL} + \Pr[H_1|H_2] V_{HH} \quad (23)$$

$$G_{H1}(b) = \frac{\Pr[L_1, H_2]}{\Pr[H_1, H_2]} \frac{b - V_{LL}}{V_{HH} - b}, \quad b \in [V_{LL}, \bar{b}] \quad (24)$$

$$G_{H2}(b) = \frac{V_{HH} - \bar{b}}{V_{HH} - b} + \frac{\Pr[H_1, L_2]}{\Pr[H_1, H_2]} \frac{b - \bar{b}}{V_{HH} - b}, \quad b \in [V_{LL}, \bar{b}] \quad (25)$$

Figure 5 illustrates cumulative bid distributions in the FPA equilibrium characterized by Theorem 5 given the information structure of Example 1. Panel A shows equilibrium bidding when bidders are informed about peaches and $p_2 = 1/5$ and $p_1 = 1/10$, for which Case 1 of Theorem 5 applies. Panel B shows equilibrium bidding when bidders are informed about lemons and $q_1 = 1/5$ and $q_2 = 1/10$, for which Case 2 of Theorem 5 applies.

The proof of Theorem 5 is in Online Appendix I and proceeds in three parts. In Appendix I.1 we show that the conditions in the theorem are necessary for a Nash equilibrium in monotone bidding strategies. In Appendix I.2 we show that the same conditions are also sufficient. In other words, the described bidding strategies (which are monotone by inspection) do constitute a Nash equilibrium. Finally, in Appendix I.3 we show that the described bidding strategies constitute a TRE. Together, these three facts imply Theorem 5.

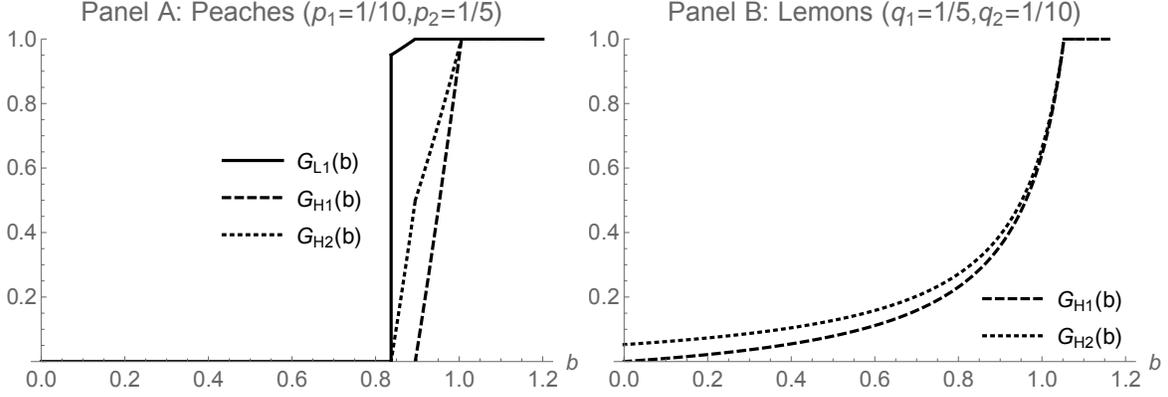


Figure 5: Equilibrium bidding in Example 1. Panel A: FPA cumulative bid distributions when bidders are informed about peaches for $p_2 = 1/5$ and $p_1 = 1/10$. Case 1 of Theorem 5 applies, and $V_{LL} \approx 0.837$, $b^* \approx 0.895$, $\bar{b} \approx 1.006$, and $V_{LH} = V_{HL} = V_{HH} = 1$. Panel B: FPA cumulative bid distributions when bidders are informed about lemons for $q_1 = 1/5$ and $q_2 = 1/10$. Case 2 of Theorem 5 applies and $V_{LL} = V_{LH} = V_{HL} = 0$, $\bar{b} \approx 1.053$, and $V_{HH} \approx 1.163$.

B Outline of the Proof of Theorem 1

If $V_{HH} = V_{LH}$ or $V_{HH} = V_{HL}$ then the strong-high-signal property holds,²⁴ and the result follows from Theorem 4, which is proven independently. We next present an outline of the proof of Theorem 1 for the case $V_{HH} > \max\{V_{HL}, V_{LH}\}$ along with four lemmas that we use to prove the result. The proof of these lemmas appears in Online Appendix H. Throughout, we maintain Assumption 1 and label bidders following equation (1).

Proof outline:

Fix any standard distribution R and $\epsilon > 0$ and let $\lambda(\epsilon, R)$ be the (ϵ, R) -tremble of the game. Let $\hat{R}(x) = 1 - \epsilon + \epsilon \cdot R(x)$ and $\hat{r}(x) = \epsilon \cdot r(x)$. Note that in both the original second price auction game and any (ϵ, R) -tremble, the set of undominated bids is closed, so the requirement that bidders only bid within the closure of the set of undominated bids is equivalent to a requirement that bidders do not place dominated bids or that equilibrium is “in undominated bids”. We therefore use this more succinct terminology throughout the proof.

To prove Theorem 1, we begin by developing a series of necessary conditions that any NE μ^ϵ in undominated bids of the tremble $\lambda(\epsilon, R)$ must satisfy. These are summarized in Lemmas 1 and 2 presented below. Next, we show that (for sufficiently small ϵ) a (mixed) NE of the tremble $\lambda(\epsilon, R)$ in undominated bids exists (Lemma 3). This existence result implies that for any standard distribution R , there exists a sequence of ϵ converging to zero and an associated sequence of NE $\{\mu^\epsilon\}$

²⁴See footnote 18.

in undominated bids corresponding to the trembles $\lambda(\epsilon, R)$. The final step is to use the necessary conditions developed in Lemmas 1 and 2 to show that the limit of any such sequence $\{\mu^\epsilon\}$ must converge to μ as ϵ goes to zero (Lemma 4). It then follows that μ is the *unique* TRE and the result in Theorem 1 holds.

Below, we present the four Lemmas 1-4. To simplify the notation we denote $v_1 = V_{HL}$, $v_2 = V_{LH}$, and $v_i = v(H_i, L_j)$, and (without loss of generality) normalize $V_{LL} = 0$ and $V_{HH} = 1$. Let $v(H_i) = E[v|H_i] = \Pr[H_j|H_i] + \Pr[L_j|H_i]v_i$. Moreover, for a given μ^ϵ , we define the following notation. First, for agent $i \in \{1, 2\}$, let $G_{H_i}(b)$ denote i 's bidding distribution conditional on signal H_i . Then define i 's infimum and supremum bids given signal H_i as $\underline{b}_i = \inf\{b : G_{H_i}(b) > 0\}$ and $\bar{b}_i = \sup\{b : G_{H_i}(b) < 1\}$. Finally, define $\underline{b} = \min\{\underline{b}_1, \underline{b}_2\}$, $b_{min} = \max\{\underline{b}_1, \underline{b}_2\}$, and $b_{max} = \max\{\bar{b}_1, \bar{b}_2\}$. Note that as bidders never submit dominated bids (by assumption) it holds that $1 \geq b_{max} \geq b_{min} \geq \underline{b} \geq 0$.

We start with some necessary conditions that any NE μ^ϵ in a fixed $\lambda(\epsilon, R)$ must satisfy.

Lemma 1 *Let Assumption 1 and $\max\{v_1, v_2\} < 1$ hold. For any standard distribution R and $\epsilon > 0$, let μ^ϵ be a Nash equilibrium in undominated bids of the game $\lambda(\epsilon, R)$. At μ^ϵ for some $i \in \{1, 2\}$, $j \neq i$, and b_{min} and b_{max} that satisfy $\max\{v_1, v_2\} \leq b_{min} \leq b_{max} \leq 1$ it holds that:*

1. Bidder i 's infimum bid is $\underline{b}_i = b_{min} \geq v_i$ and $G_{H_i}(b)$ is continuous for all $b \notin \{b_{min}, 1\}$. Bidder j 's infimum bid is $\underline{b}_j = v_j = \underline{b} \leq b_{min}$ and $G_{H_j}(b)$ is continuous for all $b \notin \{v_j, 1\}$.
2. $G_{H_i}(b_{max}) = G_{H_j}(b_{max}) = 1$. Moreover, if $b_{max} > b_{min}$ then $b_{max} = \bar{b}_i = \bar{b}_j$ and both G_{H_i} and G_{H_j} are increasing on (b_{min}, b_{max}) .
3. $G_{H_i}(b) = 0$ for every $b \in [0, b_{min})$. $G_{H_j}(b) = 0$ for every $b \in [0, v_j)$ and $G_{H_j}(b) = G_{H_j}(v_j)$ for every $b \in [v_j, b_{min}]$.
4. If $b_{min} = \underline{b} = v_j$ then $v_j \geq v_i$. If $v_j = v_i$ then $b_{min} = \underline{b} = v_j = v_i$ and no bidder has an atom below 1. If $v_j > v_i$ then j has an atom at $b_{min} = \underline{b} = v_j > v_i$ while i has no atom below 1.
5. If $b_{min} > \underline{b} = v_j$ then: (i) bidder j has an atom at v_j ; (ii) bidder i has an atom at

$$b_{min} = \frac{\Pr[H_j|H_i]G_{H_j}(v_j) + v_i \Pr[L_j|H_i]}{\Pr[H_j|H_i]G_{H_j}(v_j) + \Pr[L_j|H_i]} > \max\{v_i, v_j\}; \quad (26)$$

(iii) b_{min} satisfies $b_{min} \leq v(H_i)$; and (iv) $b_{min} = v(H_i)$ if and only if $G_{H_j}(v_j) = 1$.

It also holds that either

- $b_{max} = b_{min}$, in this case $G_{H_i}(b_{min}) = 1$ and $G_{H_j}(v_j) = 1$ (i always bids b_{min} and j always bids v_j). Or

- $b_{max} > b_{min}$, $G_{H_i}(b_{min}) > 0$ and

$$G_{H_i}(b_{min}) = \frac{\Pr[L_i|H_j] \int_{v_j}^{b_{min}} (x - v_j) \hat{r}(x) dx}{\Pr[H_i|H_j] \hat{R}(b_{min}) (1 - b_{min})}. \quad (27)$$

$$6. 0 \leq \max\{v_1, v_2\} \leq b_{min} \leq \max\{v(H_1), v(H_2)\} < 1.$$

Building on the preceding necessary conditions that apply for all ϵ , the next result gives tighter necessary conditions for NE in undominated bids in the tremble $\lambda(\epsilon, R)$ when ϵ is sufficiently small. To develop the result we first apply the first-order conditions for optimal bidding over the interval (b_{min}, b_{max}) to characterize bid distributions above b_{min} . Next, we show that for sufficiently small ϵ it holds that $b_{min} < b_{max} < 1$ (ruling out the cases $b_{max} = b_{min}$ or $b_{max} = 1$ allowed for in Lemma 1). Finally, we complete the proof by more tightly characterizing the size and placement of atoms at the bottom of bidders' bid distributions, and identifying bidders i and j from Lemma 1 as $i = 1$ and $j = 2$.

When equation (1) holds with equality, so does not distinguish the bidders, we label bidders such that $v_1 \geq v_2$. That is, we label bidders according to equation (1) and equation (28):

$$\Pr[H_1, L_2](1 - v_1) = \Pr[L_1, H_2](1 - v_2) \rightarrow v_1 \geq v_2. \quad (28)$$

Lemma 2 *Let Assumption 1, equations (1) and (28), and $\max\{v_1, v_2\} < 1$ hold. Let R be a standard distribution, $\epsilon > 0$, and μ^ϵ be a Nash equilibrium in undominated bids of the game $\lambda(\epsilon, R)$. If ϵ is small enough then at μ^ϵ there exist b_{min} and b_{max} such that $1 > b_{max} > b_{min} \geq 0$ and it always holds that:*

$$G_{H_1}(b_{min}) = \frac{\Pr[L_1|H_2] \int_{v_2}^{b_{min}} (x - v_2) \hat{r}(x) dx}{\Pr[H_1|H_2] \hat{R}(b_{min}) (1 - b_{min})} \quad (29)$$

$$G_{H_2}(v_2) = \frac{\hat{R}(b_{max})}{\hat{R}(b_{min})} - \left(\frac{\hat{R}(b_{max})}{\hat{R}(b_{min})} - G_{H_1}(b_{min}) \right) \cdot \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{\int_{b_{min}}^{b_{max}} \frac{x-v_1}{1-x} r(x) dx}{\int_{b_{min}}^{b_{max}} \frac{x-v_2}{1-x} r(x) dx} \quad (30)$$

$$G_{H_1}(b) = \begin{cases} 0 & \text{if } 0 \leq b < b_{min}; \\ \frac{\Pr[L_1|H_2]}{\Pr[H_1|H_2]} \cdot \frac{\epsilon}{\hat{R}(b)} \cdot \int_{b_{min}}^b \frac{x-v_2}{1-x} r(x) dx + G_{H_1}(b_{min}) \cdot \frac{\hat{R}(b_{min})}{\hat{R}(b)} & \text{if } b_{min} \leq b \leq b_{max}; \\ 1 & \text{if } b_{max} \leq b. \end{cases} \quad (31)$$

and

$$G_{H_2}(b) = \begin{cases} 0 & \text{if } 0 \leq b < v_2; \\ G_{H_2}(v_2) & \text{if } v_2 \leq b \leq b_{min}; \\ \frac{\Pr[L_2|H_1]}{\Pr[H_2|H_1]} \cdot \frac{\epsilon}{\hat{R}(b)} \int_{b_{min}}^b \frac{x-v_1}{1-x} \cdot r(x) dx + G_{H_2}(v_2) \cdot \frac{\hat{R}(b_{min})}{\hat{R}(b)} & \text{if } b_{min} \leq b \leq b_{max}; \\ 1 & \text{if } b_{max} \leq b. \end{cases} \quad (32)$$

Moreover, one of three cases will hold:

1. No atom case: $b_{min} = v_1 = v_2$ and $G_{H_1}(b_{min}) = G_{H_2}(b_{min}) = 0$ if and only if the two bidders are symmetric ($\Pr[H_1, L_2] = \Pr[L_1, H_2]$ and $v_1 = v_2$).
2. One atom case: $b_{min} = v_2 \geq v_1$, bidder 1 has no atom ($G_{H_1}(b_{min}) = 0$) and bidder 2 has an atom at $v_2 \geq v_1$ ($G_{H_2}(v_2) > 0$).
3. Two atom case: $b_{min} > v_2$, bidder 1 has an atom at

$$b_{min} = \frac{\Pr[H_2|H_1]G_{H_2}(v_2) + v_1 \Pr[L_2|H_1]}{\Pr[H_2|H_1]G_{H_2}(v_2) + \Pr[L_2|H_1]} > \max\{v_1, v_2\}, \quad (33)$$

($G_{H_1}(b_{min}) > 0$) and bidder 2 has an atom at v_2 ($G_{H_2}(v_2) > 0$).

If $\Pr[H_1, L_2](1 - v_1) = \Pr[L_1, H_2](1 - v_2)$ but the bidders are not symmetric, and it holds that $v_1 > v_2$ and $\Pr[H_1, L_2] < \Pr[L_1, H_2]$, then Case 3 (two atoms) holds. If $\Pr[H_1, L_2](1 - v_1) < \Pr[L_1, H_2](1 - v_2)$ then either Case 2 (one atom) or Case 3 (two atoms) holds.

We next show that, fixing any standard distribution R (such as the uniform distribution), for sufficiently small ϵ there exists a NE in undominated bids of the tremble $\lambda(\epsilon, R)$ satisfying the necessary conditions identified in Lemma 2. We prove existence separately for three sets of parameter values. For symmetric bidders, we show the existence of an equilibrium with no atoms (Case 1). For asymmetric bidders we show the existence of either a one-atom (Case 2) or a two-atom (Case 3) equilibrium.

In each case, the proof involves three steps. First we show existence of parameters b_{min} , b_{max} , $G_{H_1}(b_{min})$, and $G_{H_2}(v_2)$ that satisfy the necessary conditions in Lemma 2. Second, we show that, for the chosen parameters, G_{H_1} and G_{H_2} are well defined distributions (nondecreasing, and satisfying $G_{H_1}(0) = G_{H_2}(0) = 0$ and $G_{H_1}(1) = G_{H_2}(1) = 1$). Third we show that the constructed bid distributions are best responses. By construction, bidder $i \in \{1, 2\}$ is indifferent to all bids in the support of his bid distribution and we show that every bid outside the support gives equal or lower utility.

Lemma 3 *Let Assumption 1 and $\max\{v_1, v_2\} < 1$ hold. Fix any standard distribution R . For every small enough $\epsilon > 0$ there exists a mixed NE in undominated bids μ^ϵ of the game $\lambda(\epsilon, R)$.*

The final step is to show that any sequence of NE in undominated bids $\{\mu^\epsilon\}$ of the trembles $\lambda(\epsilon, R)$ converges to μ as ϵ goes to zero. The result, stated in Lemma 4, is proved by considering the implication of the necessary conditions identified in Lemma 2 as ϵ goes to zero. In particular, we prove a sequence of four claims about bid distributions in the limit as ϵ goes to zero given conditions from Lemma 2. (1) We show that $\lim_{\epsilon \rightarrow 0} b_{max} = 1$ by evaluating equations (31)-(32) at

$b = b_{max}$ and imposing $G_{H1}(b_{max}) = G_{H2}(b_{max}) = 1$. (2) From equation (29), we show that bidder 1's atom at b_{min} (if it exists at all) vanishes as ϵ goes to zero. (3) From equation (30), we show that bidder 2's atom at v_2 goes to $1 - \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}$. (4) Finally, we use equations (31)-(32) to show that all the bidding mass above each bidder's infimum bid goes to 1. Thus, in the limit, bidder 1 is bidding 1 with probability 1, while bidder 2 is bidding 1 with probability $\frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}$, as we need to show.

Lemma 4 *Let Assumption 1, equation (1), and $\max\{v_1, v_2\} < 1$ hold. Fix a standard distribution R and a sequence of ϵ converging to zero. The associated sequence of NE in undominated bids $\{\mu^\epsilon\}$ in the trembles $\lambda(\epsilon, R)$ converges to the NE μ in the original game λ .*

C Proofs of Corollaries 1-5

C.1 Proof of Corollary 1

Revenue is V_{LL} with probability $(1 - \Pr[H_1, H_2])$, V_{LH} with probability $\Pr[H_1, H_2] \cdot (1 - \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}})$, and V_{HH} with probability $\Pr[H_1, H_2] \cdot \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}$. Thus

$$R_{SPA} = V_{LL} + (V_{LH} - V_{LL}) \Pr[H_1, H_2] + (V_{HH} - V_{LH}) \Pr[H_1, H_2] \cdot \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]} \cdot \frac{V_{HH} - V_{HL}}{V_{HH} - V_{LH}}$$

Cancelling $(V_{HH} - V_{LH})$ from the third term yields equation (2).

C.2 Proof of Corollary 2

Part 1: Substituting $V_{HL} = V_{LH}$ and $\Pr[H_1, L_2] / \Pr[L_1, H_2] = 1$ into equation (2) yields $R_{SPA}^{symmetric} = V_{LL} + \Pr[H_1, H_2](V_{HH} - V_{LL})$. Adding \bar{V} , subtracting $\bar{V} = \Pr[L_1, L_2]V_{LL} + (\Pr[L_1, H_2] + \Pr[H_1, L_2])V_{HL} + \Pr[H_1, H_2]V_{HH}$, and cancelling terms yields equation (3). Part 2: As cookies become rare, and either $\Pr[L_1, L_2]$ or $\Pr[H_1, H_2]$ approaches 1, $\Pr[H_1, L_2]$ and $\Pr[L_1, H_2]$ must approach zero, and so the expression in equation (3) approaches \bar{V} .

C.3 Proof of Corollary 3

Part 1: Substituting $V_{LH} = V_{HL} = V_{HH} = v(P)$ into equation (2) yields $R_{SPA}^{peaches} = V_{LL} + \Pr[H_1, H_2](v(P) - V_{LL})$. Substituting equation (4) for V_{LL} then yields equation (6). Part 2: As $\Pr[L_1, L_2]$ approaches 1, $\Pr[H_1, L_2]$ and $\Pr[L_1, H_2]$ must approach zero, and so the expression in equation (6) approaches \bar{V} .

C.4 Proof of Corollary 4

Substituting $V_{LH} = V_{HL} = V_{LL} = v(L)$ into equation (2) yields $R_{SPA}^{lemons} = v(L) + (V_{HH} - v(L)) \Pr[H_1, H_2] \frac{\Pr[H_1, L_2]}{\Pr[L_1, H_2]}$. Substituting equation (5) for V_{HH} then yields equation (7).

C.5 Proof of Corollary 5

Part 1: Revenue is total expected surplus,

$$S = \Pr[H_1 H_2] V_{HH} + \Pr[L_1 H_2] V_{LH} + \Pr[H_1 L_2] V_{HL} + \Pr[L_1 L_2] V_{LL}, \quad (34)$$

less expected payoffs to each bidder. Notice that both bidders earn zero expected payoff conditional on receiving a low signal, as bidding V_{LL} is a best best response for each that yields either zero from losing, or zero from winning an item of value V_{LL} at the same price. Thus expected revenues are

$$R_{FPA} = S - \Pr[H_1] \Pi_1(b_{max} | H_1) - \Pr[H_2] \Pi_2(b_{max} | H_2). \quad (35)$$

Expected payoffs conditional on receiving a high signal are most easily calculated by considering the expected payoff from placing the maximum bid b_{max} , winning with probability one and paying b_{max} for an item with expected value $E[v(H_i, S_j) | H_i]$:

$$\Pi_1(b_{max} | H_1) = \Pr[H_2 | H_1] V_{HH} + \Pr[L_2 | H_1] V_{HL} - b_{max}, \quad (36)$$

$$\Pi_2(b_{max} | H_2) = \Pr[H_1 | H_2] V_{HH} + \Pr[L_1 | H_2] V_{LH} - b_{max}. \quad (37)$$

Revenues are therefore computed by first substituting equations (34) and (36)-(37) into equation (35) and then substituting equations (18)-(19) for b_{max} . Note that equations (18)-(19) can be substituted for b_{max} for both Case 1 and Case 2 of Theorem 5. This follows because the expression for b^* in equation (18) reduces to $b^* = V_{LL}$ under Case 2, in which case equations (19) and (23) coincide. Making the described substitutions yields

$$\begin{aligned} R_{FPA} = & \Pr[H_1, H_2] V_{HH} + \Pr[L_1, H_2] V_{LH} + \Pr[H_1, L_2] V_{HL} + \Pr[L_1, L_2] V_{LL} \\ & - \Pr[H_1] (\Pr[H_2 | H_1] V_{HH} + \Pr[L_2 | H_1] V_{HL}) \\ & - \Pr[H_2] (\Pr[H_1 | H_2] V_{HH} + \Pr[L_1 | H_2] V_{LH}) \\ & + (\Pr[H_1] + \Pr[H_2]) \Pr[H_1 | H_2] V_{HH} + (\Pr[H_1] + \Pr[H_2]) \cdot \\ & (1 - \Pr[H_1, H_2]) \frac{V_{LH} \Pr[L_1, H_2] (\Pr[L_1, H_2] - \Pr[H_1, L_2]) + V_{LL} \Pr[L_1, L_2] \Pr[H_2]}{\Pr[L_1, H_2] (\Pr[L_1, H_2] - \Pr[H_1, L_2]) + \Pr[L_1, L_2] \Pr[H_2]}, \end{aligned} \quad (38)$$

which upon rearranging terms and simplifying expressions coincides with equation (15).

Part 2: As cookies become rare and either $\Pr[L_1, L_2]$ or $\Pr[H_1, H_2]$ approaches 1 while $\Pr[L_1, H_2]$ and $\Pr[H_1, L_2]$ approach 0, the fourth term in equation 15 goes to zero. Moreover, the ratio $\Pr[H_1]/\Pr[H_2]$ approaches 1 and the sum of the second and third terms approaches $\Pr[H_1, H_2](V_{HH} - V_{LL})$. Thus $\lim_{\Pr[L_1, L_2] \rightarrow 1} R_{FPA} = \lim_{\Pr[H_1, H_2] \rightarrow 1} R_{FPA} = V_{LL} + \Pr[H_1, H_2](V_{HH} - V_{LL})$, which coincides with \bar{V} when $\Pr[L_1, H_2] = \Pr[H_1, L_2] = 0$.

D One Informed Agent

D.1 Proof of Theorem 3

Note that in the SPA game, the set of undominated bids is closed. Theorem 3 then follows from three observations:

1. Consider the SPA game or any (ϵ, R) -tremble of the game: The strategy of the informed buyer is a dominant strategy, being a best response to any possible strategies of the uninformed buyers. Moreover, it is the unique strategy in undominated bids (even among mixed strategies) as for any signal its bid is the unique bid that dominates any other bid.
2. Consider the SPA game or any (ϵ, R) -tremble of the game: For any uninformed buyer, bidding v_{min} is a best response to the strategies of the other buyers as it gives 0 utility and no strategy gives positive utility. Moreover, bidding v_{min} is undominated, while bidding less than v_{min} is dominated by bidding v_{min} .
3. Consider any (ϵ, R) -tremble of the game: For any uninformed buyer, bidding above $b_U = v_{min}$ cannot be a best response to the informed buyer's strategy of bidding $b_I = v(s_I)$. Bidding above v_{min} generates a negative expected payoff because the uninformed buyer would sometimes win and overpay at a price above $v(s_I)$ set by the random bidder, but would never pay less than the fair value $v(s_I)$ bid by the informed buyer.

Observations (1) and (2) are sufficient to show that μ is a NE in undominated bids in the SPA game and any (ϵ, R) -tremble of the game. Thus μ is a strong TRE. Observations (1) and (2) also ensure that the informed buyer bids as in μ and uninformed buyers bid at least v_{min} in any TRE. Observation (3) rules out the possibility of an uninformed buyer bidding above v_{min} in any NE in undominated bids of an (ϵ, R) -tremble of the game. Thus μ is also the unique TRE.

D.2 Proof of Proposition 1

D.2.1 A TRE Exists

In a FPA with a nonnegative²⁵ common value, an informed bidder, and m uninformed bidders, Theorem 1 of Engelbrecht-Wiggans et al. (1983) characterizes the set of Nash equilibria. The characterization describes the informed bidder's unique equilibrium bidding strategy β . Further, it describes the unique equilibrium distribution of the maximum uninformed bid, which is the product of the bid distributions for each uninformed bidder:

$$G(b) = \prod_{i \in 1 \dots m} G_i(b).$$

The characterization implies that the informed bidder bids over the interval $[v_{min}, \bar{V}]$ with no gaps and at most one atom at v_{min} . It specifies that $G(b) = \Pr(\beta(s_I) \leq b)$, which implies that $G(b)$ is continuous and increasing over $b \in [v_{min}, \bar{V}]$.

In one such NE (β, G_1, \dots, G_m) , uninformed bidder 1 bids with distribution $G_1(b) = G(b)$ and the remaining $m - 1$ uninformed bidders bid v_{min} with probability 1. To show that it is a TRE, first fix any standard distribution R . Next, define

$$\epsilon_{max} = \min_{b \in [v_{min}, \bar{V}]} \frac{g(b)}{r(b) + g(b)},$$

which is positive because $G(b)$ is continuous and increasing over $b \in [v_{min}, \bar{V}]$. Then fix any $\epsilon \in (0, \epsilon_{max})$, and let $\lambda(\epsilon, R)$ be the (ϵ, R) -tremble of the FPA game.

Theorem 1 of Engelbrecht-Wiggans et al. (1983) implies that $(\beta_\epsilon, G_{1,\epsilon}, \dots, G_{m,\epsilon})$ is a NE of the tremble if the informed bidder follows the same strategy as in the NE of the original game, $\beta_\epsilon = \beta$, and the distribution of the maximum of the uninformed *and* random bids coincides with $G(b)$ in the NE of the original game. This requires that

$$\prod_{i \in 1 \dots m} G_{i,\epsilon}(b) = G(b)/\hat{R}(b),$$

where

$$\hat{R}(b) = 1 - \epsilon + \epsilon \cdot R(x)$$

is the probability that the random bidder does not enter or enters but bids below b . Thus the following is a NE of the tremble: $\beta_\epsilon = \beta$, $G_{1,\epsilon}(b) = G(b)/\hat{R}(b)$, and the remaining $m - 1$ uninformed bidders bid v_{min} with probability 1. Note that $G_{1,\epsilon} = G(b)/\hat{R}(b)$ is a valid distribution for uninformed bidder 1's mixed strategy if it is nondecreasing, which holds if $g(b)/G(b) \geq \hat{r}(b)/\hat{R}(b)$,

²⁵Hence our assumption that $v_{min} \geq 0$.

where $\hat{r}(b) = \epsilon \cdot r(b)$ is the derivative of $\hat{R}(b)$. This is implied by $\epsilon < \epsilon_{max}$ because $\epsilon < \epsilon_{max}$ implies $g(b) > \epsilon r(b)/(1 - \epsilon)$ and:

$$\frac{g(b)}{G(b)} \geq g(b) > \frac{\epsilon r(b)}{1 - \epsilon} \geq \frac{\epsilon r(b)}{1 - \epsilon(1 - R(b))} = \frac{\hat{r}(b)}{\hat{R}(b)}.$$

Notice, that in the limit as ϵ goes to zero, $G_{1,\epsilon}(b) = G(b)/\hat{R}(b)$ approaches $G(b)$ because $\hat{R}(b)$ approaches 1. Moreover, the closure of the set of undominated bids is $b \leq v_{max}$ for uninformed bidders and $b \leq v(s_I)$ for the informed bidder. Thus bidders only bid within the closure of the set of undominated bids, both in the original game and in the sequence of trembles. Thus the original NE under consideration is a TRE.

D.2.2 FPA Revenue Results

Recall that $v(s_I) = E[v(\omega)|s_I]$ is the informed bidder's interim expected value conditional on s_I and that v_{min} and v_{max} are the minimum and maximum such values, respectively. Let $x = v(s_I)$ and F be the cumulative distribution function of x . According to Theorem 3, SPA revenue equals v_{min} . According to Theorem 4 of Engelbrecht-Wiggans et al. (1983), FPA revenue is

$$\int_0^\infty (1 - F(x))^2 dx,$$

which can be re-written as $v_{min} + \int_{v_{min}}^\infty (1 - F(x))^2 dx$. For an informed bidder, $F(v_{min}) < 1$ so this is clearly more than v_{min} . Thus $R_{FPA}^{1-informed} > R_{SPA}^{1-informed}$.

According to Theorem 4 of Engelbrecht-Wiggans et al. (1983), the informed agent's expected payoff is

$$\int_0^\infty F(x)(1 - F(x)) dx$$

Note that the revenue and the informed agent's profit sum up to $\bar{V} = E[v(\omega)]$, the expected value of the item (and social welfare). To bound the revenue from below we bound the informed agent's profit from above.

First, we temporarily normalize values such that the informed bidder's posterior valuations lie between 0 and 1. We denote its normalized distribution as $\hat{F}(x) = F(v_{min} + (v_{max} - v_{min})x)$. Further, we denote expected revenue as $\hat{R}_{FPA}^{1-informed} = \frac{R_{FPA}^{1-informed} - v_{min}}{v_{max} - v_{min}}$ and expected surplus as $\hat{V} = \frac{\bar{V} - v_{min}}{v_{max} - v_{min}}$. Then, we use the following result due to Ahlswede and Daykin (1978).

Lemma 5 *If, for 4 nonnegative functions g_1, g_2, g_3, g_4 mapping $\mathbb{R} \rightarrow \mathbb{R}$, the following holds:*

$$\text{for all } x, y \in \mathbb{R}, \quad g_1(\max(x, y)) \cdot g_2(\min(x, y)) \geq g_3(x) \cdot g_4(y),$$

then it follows that

$$\int_a^b g_1(t)dt \cdot \int_a^b g_2(t)dt \geq \int_a^b g_3(t)dt \int_a^b g_4(t)dt.$$

We apply this lemma by setting

$$g_1(t) = \hat{F}(x), \quad g_2(x) = 1 - \hat{F}(x), \quad g_3(x) = \hat{F}(x) \cdot (1 - \hat{F}(x)), \quad g_4(x) = 1.$$

Monotonicity of \hat{F} implies that the conditions of the lemma hold. Indeed, if $x'' > x'$ then

$$\hat{F}(x'') \cdot (1 - \hat{F}(x')) \geq \hat{F}(x'') \cdot (1 - \hat{F}(x''))$$

and

$$\hat{F}(x'') \cdot (1 - \hat{F}(x')) \geq \hat{F}(x') \cdot (1 - \hat{F}(x')).$$

Then, it follows that

$$(1 - \hat{V}) \cdot \hat{V} = \int_0^1 \hat{F}(t)dt \cdot \int_0^1 (1 - \hat{F}(t))dt \geq \int_0^1 \hat{F}(t)(1 - \hat{F}(t))dt.$$

As the revenue is equal to the total welfare \hat{V} minus the informed agent's profit we conclude that the revenue is bounded from below by \hat{V}^2 :

$$\hat{R}_{FPA}^{1-informed} = \hat{V} - \int_0^1 \hat{F}(t)(1 - \hat{F}(t))dt \geq \hat{V}^2$$

We finish by undoing our temporary renormalization. Substituting $\frac{R_{FPA}^{1-informed} - v_{min}}{v_{max} - v_{min}}$ in place of $\hat{R}_{FPA}^{1-informed}$ and $\frac{\bar{V} - v_{min}}{v_{max} - v_{min}}$ in place of \hat{V} in the preceding expression and rearranging terms yields:

$$R_{FPA}^{1-informed} \geq v_{min} + \frac{(\bar{V} - v_{min})^2}{v_{max} - v_{min}}.$$

E Many Agents, each with Finitely Many Signals

E.1 Proof of Theorem 4

We prove Theorem 4 by induction at the end of this section. First, however, we develop the results (Lemma 7 and Observation 1) that are applied at each induction step. Also note that, as the set of undominated bids is closed in the SPA, we can restrict attention to undominated bids.

Preliminaries: We define a natural binary relation between signals using the relation between the lower bounds they place on the expected value. We say that signal s_i of bidder i is *weakly higher* than signal s_j of bidder j if $v_{min}(s_i) \geq v_{min}(s_j)$, and is *strictly higher* than signal s_j of bidder j if $v_{min}(s_i) > v_{min}(s_j)$. An implication of the SHSP is:

Lemma 6 *Given the SHSP holds: If s is the realized signal vector and s_i is the weakly highest realized signal (such that $v_{min}(s_j) \leq v_{min}(s_i)$ for all j) then $v(s) = v_{min}(s_i)$.*

Proof. Suppose we identify a strong high signal, condition on it not occurring, identify a strong high signal of the restricted domain, and continue recursively. At each step, the SHSP guarantees that we can find a strong high signal s_i that must satisfy $v_{min}(s_i) = \max_{s_{-i}} \{v(s_i, s_{-i}) : (s_i, s_{-i}) \text{ is feasible given all signals removed in prior rounds do not occur}\}$. This means that at every round, the strong high signal is the one with the highest remaining $v_{min}(s_i)$. As a result, this process will eventually reach the domain in which any s_j with $v_{min}(s_j) > v_{min}(s_i)$ is excluded so that s_i has the highest remaining $v_{min}(s_i)$. Moreover, s_i will be a strong high signal within this restricted domain, and as a result, $v(s) = v_{min}(s_i)$ for any feasible s in the restricted domain (that excludes all signals s_i with $v_{min}(s_j) > v_{min}(s_j)$). The result then follows. ■

A Lemma: The next lemma is a major step in showing that bidder i with signal s_i does not bid above $v_{min}(s_i)$.

Lemma 7 *Fix a signal s_j received by bidder j and any strategy profile η in which every bidder i with signal s_i strictly higher than s_j ($v_{min}(s_i) > v_{min}(s_j)$) bids $v_{min}(s_i)$ with probability 1.*

1. *If η is a NE of the tremble $\lambda(\epsilon, R)$ with $\epsilon > 0$, then no bidder i with signal s_i (including bidder j with signal s_j) weakly lower than s_j bids strictly above $v_{min}(s_j)$.*
2. *Given η , in both the original game λ and in any tremble $\lambda(\epsilon, R)$ for $\epsilon > 0$, the utility of bidder j with signal s_j from bidding $v_{min}(s_j)$ is at least as high as his utility from any higher bid.*

Proof. Proof of part (1): Let $\bar{b}_i(s_i)$ be the supremum bid by bidder i with signal s_i under η . Let \bar{b} be the maximum supremum bid among signals weakly lower than s_j :

$$\bar{b} \equiv \max_{i \in N, s_i \in S_i} \{\bar{b}_i(s_i) : v_{min}(s_i) \leq v_{min}(s_j)\}.$$

Suppose η is a NE of the tremble $\lambda(\epsilon, R)$ but $\bar{b} > v_{min}(s_j)$. Let $\delta > 0$ be sufficiently small such that (1) $v_{min}(s_j) < \bar{b} - \delta$ and (2) for any bidder i and signal s_i , $\bar{b}_i(s_i) < \bar{b}$ implies $\bar{b}_i(s_i) < \bar{b} - \delta$. With positive probability, no signal strictly higher than s_j is realized²⁶ and the highest realized bid falls in the interval $(\bar{b} - \delta, \bar{b}]$. Therefore at least one bidder k with a signal s_k that satisfies $\bar{b}_k(s_k) = \bar{b}$ and $v_{min}(s_k) \leq v_{min}(s_j)$ (possibly $k = j$ and $s_k = s_j$) wins with positive probability with a bid in the interval $(\bar{b} - \delta, \bar{b}]$.

²⁶By definition of $v_{min}(s_j)$, there exists a feasible signal vector s such that $v(s) = v_{min}(s_j)$. Each element of s , s_k , satisfies $v_{min}(s_k) \leq v_{min}(s_j)$, and is therefore weakly lower than s_j .

Fix any bid $b \in (\bar{b} - \delta, \bar{b}]$ that wins with positive probability conditional on being placed by bidder k with signal s_k . We show below that for bidder k with signal s_k , bidding $v_{min}(s_j)$ is strictly more profitable than bidding b . Because bidder k with signal s_k makes such bids with positive probability, this contradicts η being a NE. The argument follows below.

Consider a particular realization in which bidder k receives signal s_k . Let b_{-k}^{max} be the highest realized bid of bidders other than k (including the random bidder). Further, let bidder i be the bidder who has the highest realized signal s_i and his bid be b_i . (If there are multiple bidders whose signals tie for the highest then choose any i from the set.) Lemma 6 implies that $v(s) = v_{min}(s_i)$.

Now compare k 's outcome from bidding b rather than $v_{min}(s_j)$. If $b_{-k}^{max} < v_{min}(s_j)$ or $b_{-k}^{max} > b$ then k 's outcome and payoff are unchanged by bidding b rather than $v_{min}(s_j)$. However, if $b_{-k}^{max} \in [v_{min}(s_j), b]$ then k wins and pays b_{-k}^{max} by bidding b at some cases were he was losing by bidding $v_{min}(s_j)$. Consider three cases. **First**, suppose that some bidder with a signal strictly higher than s_j is bidding b_{-k}^{max} . Then by assumption $b_{-k}^{max} = b_i = v_{min}(s_i)$ and by the strong-high-signal property (Lemma 6) $v(s) = v_{min}(s_i)$. Thus the additional win is priced at its value and does not change k 's payoff. **Second**, suppose $b_{-k}^{max} = v_{min}(s_j)$. Then by SHSP (Lemma 6) $v(s) = v_{min}(s_j)$ and the additional win is priced at its value and does not change k 's payoff. **Third**, suppose $b_{-k}^{max} \in (v_{min}(s_j), b]$ and it is not the bid of a bidder with a strictly higher signal. If no signal strictly higher than s_j is realized, then by SHSP (Lemma 6) $v(s) = v_{min}(s_j)$. If at least one signal strictly higher than s_j is realized, then by assumption $b_{-k}^{max} > b_i = v_{min}(s_i)$ and by SHSP (Lemma 6) $v(s) = v_{min}(s_i)$. In either case, the additional win must be priced strictly above its value ($b_{-k}^{max} > v(s)$) and strictly reduces k 's payoff.

The preceding paragraph shows that for any realization, bidding b yields an equal or lower payoff for k than bidding $v_{min}(s_j)$ and in the third case yields a lower payoff. The third case occurs with positive probability in any tremble $\lambda(\epsilon, R)$ with $\epsilon > 0$. Therefore bidding b rather than $v_{min}(s_j)$ reduces k 's expected payoff ex ante.

Proof of part (2): In the proof of part (1) above, we showed that bidding $v_{min}(s_j)$ is strictly better than bidding $b > v_{min}(s_j)$ for bidder j with signal s_j given η . Almost the same argument can be repeated to prove part (2). The only difference is that we cannot claim a strict payoff ranking because the third case in the proof of part (1) need not occur with positive probability in the original game λ without a random bidder. ■

An Observation: We next observe that bidder i with signal s_i that only submits undominated bids never bids below $v_{min}(s_i)$.

Observation 1 *In the original game λ and in any tremble $\lambda(\epsilon, R)$ for $\epsilon > 0$, for bidder i with signal s_i bidding $v_{min}(s_i)$ weakly dominates bidding any amount $b_i < v_{min}(s_i)$.²⁷*

The Proof: We combine Observation 1 with Lemma 7 to prove Theorem 4 by induction.

Proof. (of Theorem 4) Fix any strict linear order over the set of all bidders' feasible signals that is consistent with the order of lower bounds they place on the expected value. That is, fix an arbitrary order such that for every s_i and s_j in the set $\bigcup_{k=1}^N S_k$, if $v_{min}(s_i) > v_{min}(s_j)$ then s_i is ranked higher than s_j .

The proof proceeds by induction. The base case considers the highest signal according to the fixed order. Suppose the highest signal is bidder i 's signal s_i . By Observation 1, bidding $v_{min}(s_i)$ dominates any lower bid for bidder i with signal s_i . Moreover, SHSP implies that for the highest signal s_i , for any $s_{-i} \in S_{-i}$ such that (s_i, s_{-i}) is feasible, it holds that $v_{min}(s_i) = v(s_i, s_{-i})$. Thus $v_{min}(s_i) = v_{max}(s_i)$ and therefore in any tremble $\lambda(\epsilon, R)$ in which the bid of agent i with signal s_i belongs to $[v_{min}(s_i), v_{max}(s_i)]$ it holds that the bid must be $v_{min}(s_i) = v_{max}(s_i)$. Moreover, bidding $v_{min}(s_i)$ is a dominant strategy for bidder i with signal s_i in the original game λ and any tremble $\lambda(\epsilon, R)$.

We move to the induction step. Consider the l^{th} highest signal, which is s_j received by bidder j . Assume that every bidder i with strictly higher signal s_i (that is, $v_{min}(s_i) > v_{min}(s_j)$) bids $v_{min}(s_i)$ with probability 1. Observation 1 and claim (2) of Lemma 7 imply that it is a best response for bidder j with signal s_j to bid $v_{min}(s_j)$ in the original game λ and the tremble $\lambda(\epsilon, R)$. Moreover, Observation 1 and claim (1) of Lemma 7 imply that this is the *unique* best response in any NE in undominated bids of any tremble $\lambda(\epsilon, R)$, for $\epsilon > 0$.

Proceeding by induction through all signals shows that the pure strategy profile μ is a Nash equilibrium both in the original game λ and in any tremble $\lambda(\epsilon, R)$ with $\epsilon > 0$. Moreover, it is the unique Nash equilibrium in undominated bids in any tremble $\lambda(\epsilon, R)$ with $\epsilon > 0$. The theorem follows directly. ■

E.2 Proof of Proposition 2

If $\sum_{i=1}^n \epsilon_i \geq 1$ the claim follows trivially. Items have values in $[0, 1]$ and thus $E[v(\omega)] \leq 1$, this implies that $E[v(\omega)] - \sum_{j=1}^n \epsilon_j \leq 0$, and the claim about the revenue clearly holds as revenue is nonnegative since every bid is nonnegative. We next assume that $\sum_{i=1}^n \epsilon_i < 1$.

Let $L = (L_1, L_2, \dots, L_n)$ be the combination of signals in which each agent i gets signals L_i .

²⁷It is trivial to come up with strategies for the other bidders for which bid $v_{min}(s_i)$ gives strictly higher utility than bid $b_i < v_{min}(s_i)$.

Observe that $Pr[not L] \leq \sum_{i=1}^n Pr[s_i \neq L_i] \leq \sum_{i=1}^n \epsilon_i$ as every agent i is ϵ_i -informed about peaches, thus $Pr[L] \geq 1 - \sum_{j=1}^n \epsilon_j > 0$ which means that L is feasible. As the domain is monotonic and L_i is the lowest signal for agent i , for every feasible s it holds that $v(L) \leq v(s)$. This implies that $v_{min}(s_i) \geq v(L)$ for every agent i and signal $s_i \in S_i$.

As all bids are at least $v(L)$, the revenue is at least $v(L)$, thus it is sufficient to show that $v(L) \geq E[v(\omega)] - \sum_{j=1}^n \epsilon_j$.

Observe that

$$E[v(\omega)] = v(L) \cdot Pr[L] + v(not L) \cdot Pr[not L]$$

Which implies that

$$v(L) = \frac{E[v(\omega)] - v(not L) \cdot Pr[not L]}{Pr[L]} \geq E[v(\omega)] - Pr[not L] \geq E[v(\omega)] - \sum_{i=1}^n \epsilon_i$$

since $0 < Pr[L] \leq 1$, $v(not L) \leq 1$ (as for any ω it holds that $v(\omega) \in [0, 1]$), and $Pr[not L] < \sum_{i=1}^n \epsilon_i$.

E.3 Proof of Proposition 3

If $\sum_{i=1}^n \epsilon_i \geq 1$ the claim follows trivially. Items have values in $[0, 1]$ and thus all bids are at most 1, which implies that the revenue is at most 1. We next assume that $\sum_{i=1}^n \epsilon_i < 1$.

Since each $j < i$ is ϵ_j -informed about peaches it holds that

$$Pr[L_1, L_2, \dots, L_{i-1}] \geq 1 - \sum_{j=1}^{i-1} \epsilon_j$$

Now, since i is ϵ_i -informed about lemons it holds that $Pr[H_i] \geq 1 - \epsilon_i$, and thus

$$Pr[L_1, L_2, \dots, L_{i-1}, H_i] \geq Pr[L_1, L_2, \dots, L_{i-1}] + Pr[H_i] - 1 \geq 1 - \sum_{j=1}^i \epsilon_j > 0$$

The revenue obtained when the signals of agents $1, 2, \dots, i$ are *not* realized to $(L_1, L_2, \dots, L_{i-1}, H_i)$ is at most the maximal value of any item, which is 1, and that happens with probability at most $\sum_{j=1}^i \epsilon_j$. Thus this case contributes at most $\sum_{j=1}^i \epsilon_j$ to the expected revenue.

We next bound the revenue obtained when the signals of agents $1, 2, \dots, i$ are realized to $(L_1, L_2, \dots, L_{i-1}, H_i)$, an event that happens with probability at most 1. To prove the claim it is sufficient to show that the maximum bid of all agents other than i is at most ϵ_i , since this is an upper bound on revenue in this case. We first bound the bid $v_{min}(L_j)$ of any agent $j < i$ when getting signal L_j . By the first non-degeneracy assumption $(L_j, s_i, s_{-\{i,j\}})$ is feasible for some $s_i \neq H_i$ and some $s_{-\{i,j\}}$. As agent i is ϵ_i -informed about lemons it holds that $v(L_j, s_i, s_{-\{i,j\}}) \leq \epsilon_i$ and thus $v_{min}(L_j) \leq \epsilon_i$ for all $j < i$.

We next bound the bid $v_{min}(s_j)$ of any agent $j > i$ when getting any signal $s_j \in S_j$. By the second non-degeneracy assumption $(s_j, s_i, s_{-\{i,j\}})$ is feasible for some $s_i \neq H_i$ and some $s_{-\{i,j\}}$. As agent i is ϵ_i -informed about lemons it holds that $v(s_j, s_i, s_{-\{i,j\}}) \leq \epsilon_i$ and thus $v_{min}(s_j) \leq \epsilon_i$ for all $j > i$. We have shown that when the signals of agents $1, 2, \dots, i$ are realized to $(L_1, L_2, \dots, L_{i-1}, H_i)$ the maximum bid of all agents other than i is at most ϵ_i , thus the revenue in this case is bounded by ϵ_i , and the claim follows.

References

- Ahlsvede, R. and D. E. Daykin (1978). An Inequality for the Weights of Two Families of Sets, Their Unions and Intersections. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 43(3), 183–185. doi:10.1007/bf00536201.
- Akerlof, G. A. (1970). The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84(3), 488–500.
- Arnosti, N., M. Beck, and P. R. Milgrom (2015). Adverse Selection and Auction Design for Internet Display Advertising. doi:10.2139/ssrn.2603336.
- Avery, C. and J. H. Kagel (1997). Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation. *Journal of Economics and Management Strategy* 6(3), 573–603. doi:10.1111/j.1430-9134.1997.00573.x.
- Bikhchandani, S. (1988). Reputation in Repeated Second-Price Auctions. *Journal of Economic Theory* 46(1), 97–119. doi:10.1016/0022-0531(88)90151-2.
- Billingsley, P. (1995). *Probability and Measure* (Third ed.). New York: Wiley-Interscience.
- Board, S. (2009). Revealing Information in Auctions: The Allocation Effect. *Journal of Economic Theory* 38(1), 125–135. doi:10.1007/s00199-006-0177-4.
- Boyce, W. E. and R. C. DiPrima (1986). *Elementary Differential Equations and Boundary Value Problems* (4 ed.). Wiley.
- Bulow, J., M. Huang, and P. Klemperer (1999). Toeholds and Takeovers. *Journal of Political Economy* 107(3), 427–454. doi:10.1086/250068.
- Celis, L. E., G. Lewis, M. Mobius, and H. Nazerzadeh (2014). Buy-It-Now or Take-a-Chance: Price Discrimination Through Randomized Auctions. *Management Science* 60(12), 2927–2948. doi:10.1287/mnsc.2014.2009.
- Cheng, H. H. and G. Tan (2010). Asymmetric Common-Value Auctions with Applications to Private-Value Auctions with Resale. *Economic Theory* 45(1), 253–290. doi:10.1007/s00199-009-0505-6.
- Edelman, B., M. Ostrovsky, and M. Schwarz (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97(1), 242–259. doi:10.1257/aer.97.1.242.

- Einy, E., O. Haimanko, R. Orzach, and A. Sela (2002). Dominance Solvability of Second-Price Auctions with Differential Information. *Journal of Mathematical Economics* 37(3), 247–258. doi:10.1016/S0304-4068(02)00018-6.
- Engelbrecht-Wiggans, R., P. R. Milgrom, and R. J. Weber (1983). Competitive Bidding and Proprietary Information. *Journal of Mathematical Economics* 11(2), 161–169. doi:10.1016/0304-4068(83)90034-4.
- Harstad, R. M. and D. Levin (1985). A Class of Dominance Solvable Common-Value Auctions. *The Review of Economic Studies* 52(3), 525–528.
- Hashimoto, T. (2013). Equilibrium Selection, Inefficiency, and Instability in Internet Advertising Auctions. doi:10.2139/ssrn.1483531.
- Helft, M. and T. Vega (2010, August 30). Retargeting Ads Follow Surfers to Other Sites. *The New York Times* (A1). <http://www.nytimes.com/2010/08/30/technology/30adstalk.html>.
- Hendricks, K. and R. H. Porter (1988). An Empirical Study of an Auction with Asymmetric Information. *The American Economic Review* 78(5), 865–883.
- Jackson, M. O. and J. M. Swinkels (2005). Existence of Equilibrium in Single and Double Private Value Auctions. *Econometrica* 73(1), 93–139. doi:10.1111/j.1468-0262.2005.00566.x.
- Klemperer, P. (1998). Auctions with Almost Common Values: The ‘Wallet Game’ and its Applications. *European Economic Review* 42(3-5), 757–769. doi:10.1016/s0014-2921(97)00123-2.
- Larson, N. (2009). Private Value Perturbations and Informational Advantage in Common Value Auctions. *Games and Economic Behavior* 65(2), 430–460. doi:10.1016/j.geb.2007.07.010.
- Levin, D. and J. H. Kagel (2005). Almost Common Values Auctions Revisited. *European Economic Review* 49(5), 1125–1136. doi:10.1016/j.euroecorev.2003.06.001.
- Levin, J. and P. Milgrom (2010). Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review* 100(2), 603–07. doi:10.1257/aer.100.2.603.
- Liu, H. (2014). Equilibrium Selection in Common-Value Second-Price Auctions. *Games and Economic Behavior* 84, 1–6. doi:10.1016/j.geb.2013.12.003.
- Malueg, D. A. and R. Orzach (2009). Revenue Comparison in Common-Value Auctions: Two Examples. *Economics Letters* 105(2), 177–180. doi:10.1016/j.econlet.2009.07.018.

- Malueg, D. A. and R. Orzach (2012). Equilibrium and Revenue in a Family of Common-Value First-Price Auctions with Differential Information. *International Journal of Game Theory* 41(2), 219–254. doi:10.1007/s00182-011-0282-x.
- Matthews, S. A. (1984). Information Acquisition in Discriminatory Auctions. In M. Boyer and R. Kihlstrom (Eds.), *Bayesian Models in Economic Theory*, pp. 181–207. Amsterdam: North-Holland.
- McAfee, R. P. (2011). The Design of Advertising Exchanges. *Review of Industrial Organization* 39(3), 169–185. doi:10.1007/s11151-011-9300-1.
- Milgrom, P. and J. Mollner (2016). Test-Set Equilibrium with Auction Applications. Available at Mollner’s webpage at <https://sites.google.com/site/joshuamollner/>.
- Milgrom, P. R. (1981). Rational Expectations, Information Acquisition, and Competitive Bidding. *Econometrica* 49(4), 921–943. doi:10.2307/1912511.
- Milgrom, P. R. and R. J. Weber (1982a). A Theory of Auctions and Competitive Bidding. *Econometrica* 50(5), 1089–1122. doi:10.2307/1911865.
- Milgrom, P. R. and R. J. Weber (1982b). The Value of Information in a Sealed-Bid Auction. *Journal of Mathematical Economics* 10(1), 105–114. doi:10.1016/0304-4068(82)90008-8.
- Murto, P. and J. Välimäki (2015). Common Value Auctions with Costly Entry.
- Parreiras, S. O. (2006). Affiliated Common Value Auctions with Differential Information: The Two Bidder Case. *Contributions to Theoretical Economics* 6(1), Article 12. doi:10.2202/1534-5971.1304.
- Rodriguez, E. G. (2000). First Price Auctions: Monotonicity and Uniqueness. *International Journal of Game Theory* 29(3), 413–432. doi:10.1007/s001820000049.
- Selten, R. (1975). Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* 4(1), 25–55. doi:10.1007/bf01766400.
- Shiller, B. R. (2014). First-Degree Price Discrimination Using Big Data. Available at from Ben Shiller’s homepage at http://benjaminshiller.com/images/First_Degree_PD_Using_Big_Data_Apr_8,_2014.pdf.
- Simon, L. K. and M. B. Stinchcombe (1995). Equilibrium Refinement for Infinite Normal-Form Games. *Econometrica* 63(6), 1421–1443. doi:10.2307/2171776.

- Syrgekakis, V., D. Kempe, and E. Tardos (2013). Information Asymmetries in Common-Value Auctions with Discrete Signals. doi:10.2139/ssrn.2255268.
- Varian, H. R. (2007). Position Auctions. *International Journal of Industrial Organization* 25(6), 1163–1178. doi:10.1016/j.ijindorg.2006.10.002.
- Weverbergh, M. (1979). Competitive Bidding with Asymmetric Information Reanalyzed. *Management Science* 25(3), 291–294. doi:10.1287/mnsc.25.3.291.
- Wilson, R. B. (1967). Competitive Bidding with Asymmetric Information. *Management Science* 13(11), 816–820.