



The winner's bliss in common-value auctions under horizontal differentiation

Jiawei Chen ^{a,*}, Anh Nguyen ^b , Matthew Shum ^c

^a UC Irvine, United States of America

^b Carnegie Mellon, United States of America

^c Caltech, United States of America

ARTICLE INFO

JEL classification:

D44

D82

G14

L86

Keywords:

Winner's curse

Auctions

Common values

Horizontal differentiation

Advantageous selection

ABSTRACT

We study common-value auctions in which bidders have horizontally differentiated preferences. In a specific two-bidder parameterization, winning conveys good news about the object's value to the winner, a phenomenon we call the “winner's bliss”, in contrast to the conventional winner's curse. Winning against an opposed rival reveals a signal that confirms, rather than contradicts, the winner's own. Additional implications also differ from the conventional analysis. When bidders' preferences are sufficiently opposed, seller revenue decreases with information disclosure, and advantageous selection sustains bilateral trade under asymmetric information.

1. Introduction

The winner's curse is among the most studied phenomena in auction theory. In a common-value auction, winning reveals that opponents had less optimistic signals, causing the winner to revise downward her valuation (Thaler, 1988; Milgrom and Weber, 1982). Recent work identifies circumstances where the opposite holds. Bergemann et al. (2020) show that revenue-optimal mechanisms can produce a “winner's blessing” by skewing allocation toward lower-signal bidders. Lauermann and Speit (2023) find a related effect when the number of competitors is uncertain. Both channels require specific features of the environment, namely seller optimization or market-thickness uncertainty.

We identify a different channel, namely *horizontal preference heterogeneity*. When bidders disagree on which state makes the object valuable, winning reveals that the state is favorable to the winner. We call this the “winner's bliss”. It arises in the equilibrium of a standard first-price auction without revenue-maximizing design or uncertainty about the number of bidders.

The winner's bliss intuition is straightforward in the context of online advertising.¹ Among advertisers, a cosmetics brand targets women, while a beer brand targets men. Each observes a noisy signal about an internet user's unknown gender (the state). If the cosmetics brand wins, she infers that the beer brand also received a signal indicating the user is female, confirming rather than contradicting her own signal.

We focus on the first-price format because it is the prevailing auction format in online advertising, the application motivating this paper. We introduce a simple two-bidder model permitting equilibrium characterization. In this model, bidder preferences are parameterized by $\mu \in [0, 1]$, the probability that the two bidders share the same type, with $\mu = 0$ denoting horizontally differentiated preferences, and $\mu = 1$ indicating horizontal alignment. Our main results are:

1. Winner's bliss holds with horizontally differentiated preferences ($\mu < 1/2$); the conventional winner's curse holds with aligned preferences ($\mu > 1/2$). (Proposition 2)
2. Revenue is decreasing in preference alignment, and revealing the state reduces revenue when preferences are sufficiently opposed but raises it when sufficiently aligned. This echoes Board's (2009) allocation effect and Ganuza's (2004) “ignorance promotes competition” result. (Propositions 3 and 4)
3. In bilateral trade, horizontally differentiated preferences produce advantageous selection that sustains exchange (at $8/9 \approx 89\%$ of first-best surplus in our parameterization). This contrasts with adverse selection that impedes trade under aligned preferences. (Proposition 5)

For $\mu \in (0, 1]$, the results do not depend on the auction format. Definition 1 and Proposition 2 condition only on the event of winning,

* Corresponding author.

E-mail address: jiaweic@uci.edu (J. Chen).

¹ See Alcobendas et al. (forthcoming) for empirical evidence and discussions on how cookies in online ad auctions communicate noisy information on customer characteristics.




Fig. 1. (a) Equilibrium bidding strategies for different μ . Dashed line: ex-ante value $\mathbb{E}[v_A | s]$. For $\mu < 0.5$, bids exceed ex-ante value at low signals. (b) Informational content of winning: positive values indicate winner's bliss ($\mu < 0.5$), negative values indicate winner's curse ($\mu > 0.5$).

and the revenue of Proposition 3 is also obtained in the type-symmetric monotone equilibria of the second-price and all-pay auctions. Appendix B gives the details.

Related literature. Our work relates to papers showing that private information need not exacerbate the winner's curse or adverse selection. Ganuza and Penalva (2010) show that signal structures affect auction outcomes through the dispersion of bidders' posterior valuations. de Meza and Webb (2001) study advantageous selection in insurance markets, and Fieseler et al. (2003) show that negatively interdependent valuations can facilitate efficient trade. Finally, Loertscher and Muir (2025) study revenue-maximizing mechanisms for horizontally differentiated goods in a private-value Hotelling framework. In contrast, our setting is common-value: bidders receive signals about an unknown state, and horizontal preference heterogeneity makes winning good news about the object's value.

2. Model

We introduce a simple two-bidder first-price auction model permitting equilibrium and analytical characterization. An object has unknown state $\theta \in \{A, B\}$ with $\Pr(\theta = A) = 1/2$. Two risk-neutral bidders each have private type $t_i \in \{A, B\}$; bidder i receives utility $u_i = 1$ if $\theta = t_i$ and $u_i = 0$ otherwise. Types satisfy $\Pr(t_1 = t_2) = \mu \in [0, 1]$, with each type equally likely. Each bidder knows her own type but not her opponent's.

Each bidder observes an independent signal $S_i \sim U[0, 1]$ with

$$\Pr(\theta = A | S_i = s_i, S_j = s_j) = \frac{s_i + s_j}{2}.$$

A type- A bidder's expected value for the object is $v_A(s_i, s_j) \equiv \mathbb{E}[u_i | s_i, s_j] = (s_i + s_j)/2$; a type- B bidder's is $v_B(s_i, s_j) \equiv \mathbb{E}[u_i | s_i, s_j] = 1 - (s_i + s_j)/2$. The model has common values (each bidder's value depends on both signals) but signals are independent.² At $\mu = 1$ this is a standard common-value auction; at $\mu = 0$ it captures pure horizontal differentiation.

While the parameterization above has been chosen for analytical tractability, the online advertising interpretation applies as the two bidders are two advertisers whose types represent the target gender for their products, and whose signals inform about the gender of the impression. μ represents whether the two advertisers have aligned ($\mu \rightarrow 1$) or misaligned ($\mu \rightarrow 0$) preferences.

² Common-value auctions with independent signals are studied in the wallet game (Klemperer, 1998) and in Hendricks et al. (2008); our model shares this structure. For $\mu > 0$, the independent-signal structure also permits revenue equivalence across the three formats considered here; see Appendix B.

2.1. Equilibrium and the winner's bliss

We look for a symmetric Bayes-Nash equilibrium in monotone bidding strategies. Let $\beta_A(s)$ and $\beta_B(s)$ denote the equilibrium bid functions for type- A and type- B bidders, respectively. By type-symmetry of the model, $\beta_B(s) = \beta_A(1 - s)$. Proofs of all propositions are collected in the Appendix.

Proposition 1. For $\mu \in (0, 1]$, there exists a type-symmetric monotone equilibrium:

$$\beta_A(s) = \frac{1 - \mu}{2} + \frac{\mu s}{2}, \quad \beta_B(s) = \frac{1 - \mu s}{2}. \quad (1)$$

The monotone equilibrium exists for $\mu \in (0, 1]$ but not at $\mu = 0$. When the opponent is known to have opposite preferences, a bidder with the lowest signal can profitably deviate upward, breaking monotonicity. This reflects the tractability of our parameterization; in more general settings with horizontal preferences, monotone equilibria need not exist.

Definition 1. Winner's bliss occurs at signal s if $\mathbb{E}[v_A | s, \text{win}] > \mathbb{E}[v_A | s]$; winner's curse if the inequality reverses.

Proposition 2. In the equilibrium of Proposition 1,

$$\mathbb{E}[v_A | s, \text{win}] - \mathbb{E}[v_A | s] = \frac{(2\mu - 1)(s - 1)}{4}.$$

Winner's bliss holds for all $s \in [0, 1)$ when $\mu < 1/2$; winner's curse when $\mu > 1/2$.

The threshold $\mu = 1/2$ is where the opposing-type channel (winning confirms the bidder's signal) exactly offsets the same-type channel (winning means weaker opponent signal). The effect is largest at low signals and vanishes as s approaches 1. Fig. 1 illustrates this pattern.

Proposition 3. Expected revenue is $R(\mu) = 1/2 - \mu/6$, strictly decreasing in μ .

Revenue is minimized at $\mu = 1$ (aligned preferences), where $R(1) = 1/3$, and approaches its supremum of $1/2$ as $\mu \rightarrow 0$ (opposing preferences). This supremum is not attained: as noted after Proposition 1, the monotone equilibrium underlying $R(\mu)$ does not exist at $\mu = 0$ itself, so $R(\mu)$ is defined on $\mu \in (0, 1]$.




Fig. 2. (a) Expected revenue $R(q)$ versus the disclosure probability q for several levels of preference alignment μ : decreasing when $\mu < 3/4$ and increasing when $\mu > 3/4$. (b) Revelation revenue $\mu/2$ versus concealment revenue $1/2 - \mu/6$; the two cross at $\mu = 3/4$, so disclosure reduces revenue for opposed preferences and raises it for aligned ones.

2.2. Information disclosure under the winner's bliss

In some settings such as online ad auctions, platforms control how much user information they disclose to the competing bidders (advertisers). With horizontally differentiated bidders, disclosure can reduce revenue.

To explore this, consider a platform that reveals the state θ publicly with probability q and conceals it otherwise. Under concealment, bidders rely on their private signals and revenue is $1/2 - \mu/6$ by Proposition 3. Under revelation, θ becomes common knowledge and the bidders compete in a private-value auction.

Proposition 4 (Allocation Effect). Under full revelation of the state, expected revenue is $\mu/2$. For $\mu \in (0, 1]$, under disclosure probability q ,

$$R(q) = q \frac{\mu}{2} + (1 - q) \left(\frac{1}{2} - \frac{\mu}{6} \right),$$

which is decreasing in q when $\mu < 3/4$ and increasing when $\mu > 3/4$.

Revealing the state thus reduces revenue when preferences are sufficiently opposed ($\mu < 3/4$) and raises it when sufficiently aligned ($\mu > 3/4$). This is the allocation effect of Board (2009), whereby revealing which bidder values the object can soften competition. The linkage principle, under which disclosure raises revenue, holds only in the aligned region; under opposing preferences the allocation effect dominates. This threshold need not coincide with the selection threshold $\mu = 1/2$ of Proposition 2, which marks where winning switches from good to bad news. Fig. 2 illustrates this result.

2.3. Advantageous selection in bilateral trade

We apply the model to exchange. Trader 1 (buyer, type A) values the asset at $(s_1 + s_2)/2$; Trader 2 (seller, type B) values it at $1 - (s_1 + s_2)/2$. Trade occurs at exogenous price p .

This setting has genuine gains from trade only when $s_1 + s_2 > 1$ (i.e., $v_A > v_B$), placing it outside the Milgrom and Stokey (1982) no-trade premise (which requires ex ante efficiency of the initial allocation). The question is whether information asymmetry facilitates or impedes trade. Under aligned preferences, the counterparty's willingness to trade is bad news (adverse selection). Under horizontal preferences, it is good news (advantageous selection). A related observation appears in Fieseler et al. (2003), who show that efficient

mechanisms are easier to construct when valuations are negatively interdependent.

Consider threshold strategies where the buyer purchases if $s_1 > a^*$, while the seller sells if $s_2 > b^*$.

Proposition 5. At price $p = 1/2$, the unique equilibrium in threshold strategies is $a^* = b^* = 1/3$:

- Trade probability: $4/9 \approx 44.4\%$. First-best: $1/2$.
- Realized surplus efficiency: $8/9 \approx 88.9\%$.

Corollary 1 (Advantageous Selection). Both parties revise valuations favorably upon trade: $\mathbb{E}[v_A | s_1, S_2 > 1/3] > \mathbb{E}[v_A | s_1]$ and $\mathbb{E}[v_B | s_2, S_1 > 1/3] < \mathbb{E}[v_B | s_2]$. The counterparty's willingness to trade is good news, producing a positive feedback loop that sustains exchange.

As an application, in financial markets with heterogeneous investor styles (value vs. growth, domestic vs. foreign), this predicts that style diversity facilitates trade. Fig. 3 displays the equilibrium trade regions and surplus.

3. Conclusion

Horizontal preference heterogeneity produces a winner's bliss in common-value auctions.³ The resulting sign reversal in selection propagates to information disclosure and bilateral trade. Companion experimental work (Nguyen et al., 2023) provides initial laboratory evidence. The same logic may extend to two-sided matching markets under incomplete information, where being chosen by a participant with different preferences is more informative about match quality than being chosen by a similar one.

Appendix A. Proofs

³ A related logic applies in committee voting (Feddersen and Pesendorfer, 1998). Agreement from an ideologically opposed member is more informative than agreement from a like-minded one, since the opposed member must have received a strong signal to overcome her lean.




Fig. 3. (a) Trade regions in (s_1, s_2) space. Dark region: equilibrium trade ($s_1 > 1/3$ and $s_2 > 1/3$); light region: efficient trade ($s_1 + s_2 > 1$). (b) Trade probability and surplus as functions of price p ; the symmetric price $p = 1/2$ maximizes surplus.

Proof of Proposition 1. A type-A bidder with signal s who bids b wins against a same-type opponent (probability μ) when that opponent's signal $s' < \beta_A^{-1}(b)$, and against an opposing-type opponent (probability $1 - \mu$) when $s' > 1 - \beta_A^{-1}(b)$. Her expected profit is

$$\Pi(b, s) = \mu \int_0^{\beta_A^{-1}(b)} \left[\frac{s + s'}{2} - b \right] ds' + (1 - \mu) \int_{1 - \beta_A^{-1}(b)}^1 \left[\frac{s + s'}{2} - b \right] ds'.$$

Setting $\partial \Pi / \partial b = 0$ using $d\beta_A^{-1}/db = 1/\beta'_A$ and evaluating at $b = \beta_A(s)$ (so that $\beta_A^{-1}(b) = s$) yields

$$\mu \left(\frac{s - \beta_A(s)}{\beta'_A(s)} \right) + (1 - \mu) \left(\frac{1/2 - \beta_A(s)}{\beta'_A(s)} \right) = s,$$

which simplifies to $[\mu s + (1 - \mu)/2 - \beta_A(s)] = \beta'_A(s)s$, i.e., $\frac{d}{ds}[\beta_A(s)] = \mu s + (1 - \mu)/2$. Integrating:

$$s\beta_A(s) = \frac{\mu s^2}{2} + \frac{(1 - \mu)s}{2} + C.$$

Boundedness as $s \rightarrow 0$ requires $C = 0$, giving $\beta_A(s) = (1 - \mu)/2 + \mu s/2$. For global optimality, a bidder with signal s who deviates to $\beta_A(s')$ earns payoff $U(s', s)$ with $\partial U / \partial s' \propto (s - s')$, which is positive for $s' < s$ and negative for $s' > s$, confirming $s' = s$ is optimal.⁴ $\square$

Proof of Proposition 2. In equilibrium, a type-A bidder wins against a same-type opponent when $s_j < s$ (probability s) and against an opposing-type opponent when $s_j > 1 - s$ (probability s). Thus $\Pr(\text{win} | s) = s$. The conditional expected value:

$$\mathbb{E}[v_A | s, \text{win}] = \frac{\mu \cdot s \cdot \frac{3s}{4} + (1 - \mu) \cdot s \cdot \frac{s+2}{4}}{s} = \frac{s(2\mu + 1)}{4} + \frac{1 - \mu}{2}.$$

Since $\mathbb{E}[v_A | s] = s/2 + 1/4$, subtraction gives the result. $\square$

Proof of Proposition 3. Revenue equals the expected winning bid. With probability $\mu/2$ both bidders are type A; the winner has the higher signal $s_{(2:2)}$ with density $2s$, paying $\beta_A(s_{(2:2)})$. With probability $\mu/2$ both are type B (winner has the lower signal). With probability $1 - \mu$ types differ; type A wins when $s_A + s_B > 1$. Integrating across cases yields $R = 1/2 - \mu/6$. $\square$

⁴ In companion experimental work (Nguyen et al., 2023), subjects' bid sensitivity to signals increases with μ , consistent with the equilibrium prediction.

Proof of Proposition 4. Condition on $\theta = A$. A type-A bidder has value 1 and a type-B bidder value 0; the latter bids 0. A value-1 bidder faces a rival who shares her value with probability μ and uses the symmetric mixed strategy with c.d.f. $G(b) = \frac{1-\mu}{\mu} \frac{b}{1-b}$ on $[0, \mu]$, equating her payoff $[(1 - \mu) + \mu G(b)](1 - b)$ to $1 - \mu$ over the support. Given $\theta = A$, both bidders are type A with probability $\mu/2$, exactly one is type A with probability $1 - \mu$, and neither is with probability $\mu/2$; computing the expected higher bid across these cases gives $\mu/2$. Combining with the concealment revenue from Proposition 3 gives $R(q)$, and $\partial R / \partial q = 2\mu/3 - 1/2$. $\square$

Proof of Proposition 5. Buyer purchases if $\mathbb{E}[v_A | s_1, s_2 > b^*] \geq p$, giving $a^* = 2p - (1 + b^*)/2$. Seller sells if $p \geq \mathbb{E}[v_B | s_2, s_1 > a^*]$, giving $b^* = 2(1 - p) - (1 + a^*)/2$. At $p = 1/2$: $a^* = (1 - b^*)/2$, $b^* = (1 - a^*)/2$, yielding $a^* = b^* = 1/3$. Expected surplus is $\int_{1/3}^1 \int_{1/3}^1 (s_1 + s_2 - 1) ds_2 ds_1 = 4/27$; first-best surplus is $\int \int_{s_1 + s_2 > 1} (s_1 + s_2 - 1) ds = 1/6$. The ratio is $(4/27)/(1/6) = 8/9$. $\square$

Appendix B. Robustness to the auction format

For $\mu \in (0, 1]$, the effective signals in our model are independent, and the value function below is linear in a bidder's own effective signal with a constant coefficient. Thus, any equilibrium that allocates the object to the bidder with the higher effective signal and gives the lowest effective-signal type zero expected utility yields the same interim payments and hence the same expected revenue (Myerson, 1981; Milgrom and Weber, 1982; Krishna and Morgan, 1997). We verify this conclusion directly for the first-price, second-price, and all-pay auctions.

It is convenient to work with the *effective signal* $z_i = s_i$ for a type-A bidder and $z_i = 1 - s_i$ for a type-B bidder. Because $S_i \sim U[0, 1]$ independently of type, z_i is uniform on $[0, 1]$, independent across bidders and independent of type. Letting $V(z_i, z_j)$ denote a bidder's expected value given both effective signals, with the expectation taken over the opponent's unobserved type, the same expression obtains for either type of bidder,

$$V(z_i, z_j) = \mu \frac{z_i + z_j}{2} + (1 - \mu) \frac{z_i + 1 - z_j}{2} = \frac{1}{2} \left[z_i + (1 - \mu) + (2\mu - 1)z_j \right], \quad (2)$$

so that $V(z, z) = \mu z + (1 - \mu)/2$ is increasing whenever $\mu > 0$. In these coordinates the auction is symmetric, and the equilibrium of Proposition 1 takes the common form $\beta(z) = (1 - \mu)/2 + \mu z/2$ for bidders of both types.

In the second-price auction, a bidder with effective signal z submits her expected value conditional on tying an opponent whose effective

signal is also z , $\beta_{\text{SP}}(z) = V(z, z) = \mu z + (1 - \mu)/2$. These bidding strategies constitute an equilibrium, since a bidder with effective signal z who instead submits the bid assigned to y obtains $\partial U / \partial y = (z - y)/2$, so that $y = z$ is optimal globally and not merely locally, exactly as in the proof of [Proposition 1](#). The analogous construction in the all-pay auction yields $\beta_{\text{AP}}(z) = \mu z^2/2 + (1 - \mu)z/2$. Both strategies award the object to the bidder with the higher effective signal, and both generate the revenue of [Proposition 3](#),

$$\mathbb{E}[\beta_{\text{FP}}(z_{(2:2)})] = \mathbb{E}[\beta_{\text{SP}}(z_{(1:2)})] = 2\mathbb{E}[\beta_{\text{AP}}(z)] = \frac{1}{2} - \frac{\mu}{6},$$

where $z_{(2:2)}$ and $z_{(1:2)}$ denote the higher and the lower of the two effective signals. [Definition 1](#) and [Proposition 2](#) apply unchanged in either format, because they condition only on the event of winning. The first-price auction therefore serves as a standard and empirically relevant format implementing an allocation whose informational and revenue properties are shared by the three standard formats considered here.

Data availability

No data was used for the research described in the article.

References

- Alcobendas, M., Kobayashi, S., Shi, K., Shum, M., forthcoming. The impact of privacy protection on online advertising markets. *Rev. Econ. Stud.*
- Bergemann, D., Brooks, B., Morris, S., 2020. Countering the winner's curse: Optimal auction design in a common value model. *Theor. Econ.* 15 (4), 1399–1434.
- Board, S., 2009. Revealing information in auctions: The allocation effect. *Econom. Theory* 38 (1), 125–135.
- de Meza, D., Webb, D.C., 2001. Advantageous selection in insurance markets. *RAND J. Econ.* 32 (2), 249–262.
- Feddersen, T., Pesendorfer, W., 1998. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *Am. Polit. Sci. Rev.* 92 (1), 23–35.
- Fieseler, K., Kittsteiner, T., Moldovanu, B., 2003. Partnerships, lemons and efficient trade. *J. Econom. Theory* 113 (2), 223–234.
- Ganuzi, J.J., 2004. Ignorance promotes competition: An auction model of endogenous private valuations. *RAND J. Econ.* 35 (3), 583–598.
- Ganuzi, J.J., Penalva, J., 2010. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica* 78 (3), 1007–1030.
- Hendricks, K., Porter, R.H., Tan, G., 2008. Bidding rings and the winner's curse. *RAND J. Econ.* 39 (4), 1018–1041.
- Klemperer, P., 1998. Auctions with almost common values: The ‘wallet game’ and its applications. *Eur. Econ. Rev.* 42 (3–5), 757–769.
- Krishna, V., Morgan, J., 1997. An analysis of the war of attrition and the all-pay auction. *J. Econom. Theory* 72 (2), 343–362.
- Lauermann, S., Speit, A., 2023. Bidding in common-value auctions with an unknown number of competitors. *Econometrica* 91 (2), 493–527.
- Loertscher, S., Muir, E.V., 2025. Optimal Hotelling auctions. Working Paper.
- Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. *J. Econom. Theory* 26 (1), 17–27.
- Milgrom, P.R., Weber, R.J., 1982. A theory of auctions and competitive bidding. *Econometrica* 50 (5), 1089–1122.
- Myerson, R.B., 1981. Optimal auction design. *Math. Oper. Res.* 6 (1), 58–73.
- Nguyen, A., Olivola, C., Shum, M., Zeidel, J., 2023. Winner's bliss, loser's curse. Working Paper.
- Thaler, R.H., 1988. Anomalies: The winner's curse. *J. Econ. Perspect.* 2 (1), 191–202.