The GBM Auction: Supplementary Materials

Supporting theoretical results and detailed analyses

Contents

1 The GBM Auction: Supplementary Materials

This document contains supplementary theoretical results, extended discussions, and detailed proofs that were cut from the main paper to maintain focus and length. All material is organised by topic and reference to the main paper's sections and theorem numbers.

Note on organisation: Each supplement section corresponds to material from a specific part of the main paper. Theorem and proposition numbering follows the original paper to enable cross-referencing. This supplementary is self-contained for the material it covers — all proofs and discussions needed to understand the results are included.

1.1 Supplement S1: Formal Game Definition, Distributional Robustness, and Stationarity (Section 2.8)

1.1.1 The GBM Ascending Auction as a Sequential Game

The GBM ascending auction is a sequential game of incomplete information $\Gamma = (N, V, W, K_{\max}, \mathcal{A}, u)$ defined as follows:

Players: $N = \{1, \ldots, n\}$ bidders plus a passive seller.

Types: Each bidder i draws a private value $v_i \in [0, \bar{v}]$ independently from F. The type space is $V = [0, \bar{v}]^n$. Each bidder has exogenous wealth $w_i \geq v_i$.

Information: At each decision point, bidder i observes the full public history $h_t = (b_0, b_1, \ldots, b_t)$ of bids placed so far (and thus the current standing bid b_t and the number of bids), but not the identities, values, or wealth of other bidders.

Actions: When prompted (bidders are prompted in an arbitrary order; the analysis is invariant to the prompting rule since only the current price matters for the threshold strategy), bidder i observes the standing bid $x = b_t$ and chooses either to *pass* (exit permanently) or to *bid* $y \in [\alpha x, w_i]$, where $\alpha = 1 + \text{step\_min}$. The action space at standing bid x is $\mathcal{A}_i(x) = \{\text{pass}\} \cup [\alpha x, w_i]$.

Termination: The auction ends when (a) all bidders except one have exited (passed), or (b) the round limit $K_{\max}$ is reached and all active bidders have been prompted under the anti-sniping rule. In case (a), the remaining bidder wins and pays their last bid b_K. In case (b), an efficient tiebreaking rule determines the winner among remaining active bidders.

Anti-sniping rule: After a bid is placed, each remaining active bidder is guaranteed the opportunity to respond before the auction ends. This ensures that no bidder is prevented from acting by the time limit. All deployed GBM auctions include this rule.

Efficient tiebreaking rule: If the auction reaches $K_{\max}$ with multiple active bidders, a sealed-bid second-price auction among remaining active bidders determines the winner. This rule is needed only when multiple bidders from the non-monotone region

$v \leq v^*$ remain active at expiry — an event with negligible probability for practical parameters (see Proposition 6.1 in the main paper).

Payoffs: The winner receives $v_i - b_K$ (plus any incentives collected in earlier rounds); each outbid bidder at position $j < K$ receives $g(b_{j-1}, b_j) \cdot b_j > 0$; the seller receives $b_K - \sum_{j=1}^{K-1} g(b_{j-1}, b_j) \cdot b_j$.

Strategy: A (pure) strategy for bidder i is a measurable function $\sigma_i : [0, \bar{v}] \times \mathcal{H} \to \mathcal{A}$ mapping the bidder's value and the public history to an action. A *threshold strategy* is characterised by a stopping threshold $T(v)$ and a bid-size rule $y(x, v)$: bid $y(x, v)$ if $x < T(v)$, pass otherwise.

1.1.2 Distributional Robustness

The vast majority of results in this paper hold for any continuous F with positive density on $[0, \bar{v}]$. This includes not only the structural results — solvency (Proposition 2.1), Sybil resistance (Theorem 7.1), collusion resistance (Theorem 7.3), the revenue floor (Theorem 5.2), the welfare advantage (Theorem 6.2), the entry advantage (Theorem 6.4), and the FOSD bidder preference (Theorems 6.3–6.3b) — but also the BNE threshold (Theorem 4.1), the overbidding bound (Theorem 4.3), uniqueness (Theorem 4.5), variance reduction (Proposition 8.1), and the risk-averse theorems (Theorems 8.2–8.3). All of these are stated and proved for general F; the threshold depends on F only through $F(v)$ (the quantile). The equilibrium revenue share $\rho = \text{inc\_min}/\text{step\_min}$ is independent of F.

The results that use $F = U[0, 1]$ are the quantitative specializations: the closed-form advantage factor $1 + \rho(n - 1)$ (Theorem 6.2), the entry dominance condition $m(m + 1) < 2/\rho$ (Theorem 6.4, part iii), the CE tables (Section 8.2), and the explicit variance bound (Proposition 8.1). Extending these to specific non-uniform families (e.g., power distributions $F(v) = v^\beta$) is straightforward and left for future work.

1.1.3 Stationarity of Threshold Strategies

The time limit does not affect the equilibrium characterisation (Theorem 4.1). Threshold strategies are stationary: the optimal stopping rule depends on the current price, not on elapsed time or the number of remaining rounds. This is because at each decision point, past incentive earnings are sunk (they cancel in the continue-vs-stop comparison), and the marginal payoff of one more bid depends only on the current price and the win probability. The anti-sniping rule ensures that every active bidder can respond to every bid, so no bidder is forced to exit by the clock. Any round limit $K_{\max}$ large enough for the auction to reach $\bar{v}$ (i.e., $K_{\max} \geq \log(\bar{v}/b_0)/\log \alpha$) yields the same equilibrium outcome.

1.2 Supplement S2: Extended Related Literature (Section 1.3)

Classical auction theory. The Revenue Equivalence Theorem (Myerson 1981; Riley and Samuelson 1981) establishes that standard auction formats yield the same expected revenue under symmetric independent private values (IPV) with risk-neutral bidders. For comprehensive treatments, see Krishna (2002) and Milgrom (2004). GBM violates the RET's boundary condition (Section 3.3 of the main paper), placing it outside this equivalence class.

Ascending auctions. The strategic equivalence between the English auction and the second-price sealed-bid auction (Vickrey 1961) relies on the dominant strategy of bidding one's value. GBM breaks this dominance (Section 3.2 of the main paper), making the ascending and sealed-bid formulations non-equivalent.

Optimal mechanism design. Myerson's (1981) revenue-maximising mechanism for IPV settings uses a reserve price to exclude low-value bidders; Engelbrecht-Wiggans (1993) extends this analysis to settings where the seller's objective function differs from pure revenue maximisation. GBM takes the opposite approach to Myerson: it includes low-value bidders by paying them to participate. The revenue comparison (Section 6.4 of the main paper) shows GBM at 90% of the Myerson benchmark, but the endogenous entry argument (Section 6.5) can reverse this ranking.

Entry and participation literature. Levin and Smith (1994) characterise equilibrium entry in auctions with endogenous bidder participation. Samuelson (1985) and Bulow and Klemperer (1996) show that more participants always increase revenue; GBM's entry advantage (Theorem 6.4) extends this via the indirect participation effect. Bikhchandani and Riley (1991) study optimal entry rules; GBM's self-financing property avoids the need for an external entry subsidy.

Collusion and bidder rings. Graham and Marshall (1987) and McAfee and McMillan (1992) characterise the stability of collusive rings. Pesendorfer (2000) studies collusion in ascending auctions. GBM's solvency constraint (Theorem 7.3) makes collusion provably non-profitable, a stronger result than in standard auctions.

Auctions with externalities. Jehiel and Moldovanu (2000) analyse auctions in which losing bidders experience allocative externalities from the winner's identity. GBM introduces a different kind of externality — a *direct pecuniary externality* imposed by the mechanism designer. The lose-payoff in GBM is not a strategic consequence of the allocation but a structural feature of the payment rule. While conceptually distinct from allocative externalities, the shared property is that losers' payoffs are non-zero, which similarly breaks the RET boundary condition.

Loser payoffs in mechanism design. The general principle that losers should earn positive payoffs is sometimes invoked in optimal mechanism design (Myerson and Satterthwaite 1983). GBM operationalizes this principle in ascending auctions with a specific self-financing structure.

Contrast with penny auctions. A superficial resemblance exists with penny auctions (also called pay-to-bid auctions), which are also ascending formats with modified loser payoffs. The resemblance is misleading: in penny auctions, each bid costs a non-refundable fee (typically \$0.50–\$1.00), so every losing bidder leaves with *less* money than they started with — the loser payoff is strictly negative. GBM inverts this structure entirely: every outbid participant receives their full bid back *plus* a positive incentive, making the loser payoff strictly positive. Penny auctions generate revenue by extracting sunk costs from losers; GBM generates participation by *paying* losers. The strategic consequences are correspondingly opposite: penny auctions exhibit sunk-cost escalation dynamics and have been widely criticised as exploitative (Augenblick 2016), whereas GBM's positive lose-payoff is individually rational at every decision point (Theorem 4.1).

All-pay auctions with caps and rebates. Che and Gale (1998) and Kaplan and Wettstein (2006) study all-pay auctions with bid caps and partial rebates to losing bidders. In these mechanisms, all bidders pay their bids, but losers recover a fraction. GBM achieves a similar economic effect through a different channel: losers receive their full bid back *plus* an incentive, so there is no sunk cost at any point. This distinction is important for participation incentives — a mechanism where losing is net-positive (GBM) is qualitatively different from one where losing is merely "less negative" (partial rebate).

Endogenous entry (Levin-Smith framework). Levin and Smith (1994) establish that equilibrium entry in auctions depends on the entry cost, the number of potential entrants, and the auction format. GBM strictly increases equilibrium entry by improving the participation insurance property (Theorem 6.5), a non-obvious result that requires both the FOSD welfare advantage and the variance reduction from risk aversion.

Risk aversion in auctions. Cox, Robson, and Smith (1982), Maskin and Riley (2000), and others have studied risk-averse bidding in standard auctions. GBM's variance reduction (Proposition 8.1) amplifies the risk-aversion welfare advantage beyond standard auctions, and the risk-averse threshold (Theorem 8.3) shows that risk aversion reduces overbidding, a monotonicity property not standard in auction literature.

1.3 Supplement S3: Relationship to Other Participation-Encouraging Mechanisms (Section 3.5)

GBM belongs to a broader family of mechanisms that modify auction or contest rules to encourage participation. This section positions GBM relative to four related designs. The common thread is that each mechanism makes losing less painful — but the structural details differ substantially, and these differences drive important distinctions in incentive properties, budget balance, and practical deployment.

1.3.1 Contests with Participation Prizes

Moldovanu and Sela (2001) study contests in which a principal allocates prizes to incentivise effort, including participation prizes awarded to all entrants. Morgan (2000) analyses a related idea in a public goods setting: financing public goods via lotteries, where the prize structure encourages voluntary contributions that would otherwise be underprovided — a parallel to GBM's use of incentive payments to encourage auction participation that would otherwise be insufficient.

GBM's incentive payment is structurally similar to both: it rewards losers. The key difference is in funding. In contests, the prize budget is set exogenously by the principal. In GBM, the losing payment is endogenous — it is funded from the bid increment, and its magnitude depends on the bid sequence. This self-funding property means the seller does not need to commit a prize budget ex ante; the incentive cost is automatically bounded by ρ times the winning price (Theorem 5.1 of the main paper). The solvency constraint ensures budget balance bid-by-bid (Proposition 2.1), a property that contest models do not typically guarantee.

1.3.2 The Amsterdam Auction

Goeree and Offerman (2004) propose a two-stage mechanism — a first-price sealed bid followed by an ascending phase — designed to encourage entry by offering losing bidders a premium. The Amsterdam Auction and GBM share the goal of attracting participants through positive losing payoffs, but differ in three ways.

First, GBM is purely ascending (simpler to implement and more transparent to bidders), while the Amsterdam Auction requires a sealed-bid stage that introduces strategic complexity.

Second, GBM's incentive rate is proportional to the bid and determined at bid time — it depends only on the bidder's own choice, not on the final outcome — whereas the Amsterdam premium depends on the sealed-bid stage outcome.

Third, the solvency constraint provides GBM with an automatic budget-balance guarantee; the Amsterdam Auction requires the designer to calibrate the premium schedule to avoid losses.

1.3.3 Penny Auctions

Augenblick (2016) analyses penny auctions, which share two features with GBM: ascending prices and positive losing payoffs (since each bid costs a small fixed fee but moves the price by only one penny, losers forfeit their bid fees but the winner acquires the item far below market value). Despite these surface similarities, the economics are fundamentally different.

In penny auctions, the bid cost is a fixed fee unrelated to the item's value, the final price is far below the market price, and the mechanism produces a war-of-attrition dynamic with heavy losses concentrated on late entrants. GBM's minimum step is proportional to the current price (not fixed), the winning price reflects competition among bidders (Theorem 4.3 of the main paper), losers receive a net gain (not a net loss), and the solvency constraint prevents the escalation dynamics that characterise penny auctions.

The distinction matters practically: penny auctions have faced regulatory scrutiny due to their gambling-like properties (Augenblick 2016), whereas GBM's incentive structure is designed to be self-correcting — every participant leaves with more than they committed.

1.3.4 All-Pay Auctions with Caps and Rebates

Che and Gale (1998) and Kaplan and Wettstein (2006) study all-pay auctions with bid caps and partial rebates to losing bidders. In these mechanisms, all bidders pay their bids, but losers recover a fraction.

GBM achieves a similar economic effect through a different channel: losers receive their full bid back *plus* an incentive, so there is no sunk cost at any point. This distinction is important for participation incentives — a mechanism where losing is net-positive (GBM) is qualitatively different from one where losing is merely "less negative" (partial rebate). The all-pay literature also focuses on sealed-bid formats, whereas GBM is ascending, which provides price discovery and transparency. The revenue and welfare properties differ accordingly: the all-pay rebate results rely on revenue equivalence arguments that do not apply to GBM (Observation 3.3 of the main paper).

1.3.5 Equilibrium Strategy Complexity: GBM vs Amsterdam

A notable structural difference between GBM and the Amsterdam Auction concerns the complexity of equilibrium bidding strategies. In the Amsterdam Auction's first stage (sealed-bid), no dominant strategy exists: bidders face the standard first-price sealed-bid trade-off between bidding aggressively (higher winning probability, lower surplus) and bidding conservatively (lower winning probability, higher surplus). The Goeree-Offerman (2004) equilibrium involves mixed strategies in the sealed-bid stage, which are notoriously difficult to compute and implement in practice.

In GBM, by contrast, the ascending format admits a pure-strategy BNE throughout (Theorem 4.1 of the main paper). The bidder's decision at each round reduces to a simple threshold comparison: continue if the current price is below $T(v)$, stop otherwise. This threshold is a closed-form function of the bidder's value, the number of bidders, and the distribution F. The strategic simplicity of the ascending format — price discovery through sequential revelation rather than simultaneous guessing — is preserved in GBM despite the additional incentive channel. This is a meaningful practical advantage: bidders in GBM need only evaluate a single closed-form expression, whereas bidders in the Amsterdam Auction must solve a more complex optimisation problem that depends on the entire distribution of opponent strategies.

1.4 Supplement S4: Extended Equilibrium Proofs (Sections 4.1–4.5)

This section contains the full proofs of equilibrium results that were condensed in the main paper.

1.4.1 Proof Extension for Theorem 4.1: Why the Static Win Probability is Exact

For bidders with $v > v^*$ (the monotone region of T), the threshold $T(v)$ is strictly increasing, so the mapping from values to dropout prices is injective. The order in which these bidders exit perfectly reflects the order of their values. A bidder at price $p < T(v)$ learns only that at least one opponent has value above the dropout threshold corresponding to p — but since the exit order is deterministic and monotone, this does not update the probability that the bidder's value is the highest. Formally, the event "I win" $= \{v_i > \max_{j \neq i} v_j\}$ has probability $F(v)^{n-1}$ regardless of the current price.

For bidders with $v \leq v^*$ (the non-monotone region), $T(v) \to \infty$ and the static win probability $F(v)^{n-1}$ is no longer exact — the non-monotonicity breaks the injective mapping from values to dropout prices. In the timed model, these bidders are capped by the round limit or the efficient tiebreaker at expiry (Section 2.8 of the main paper). Since the non-monotone region has measure $F(v^*)$ (which equals 0.01 for HIGH at $n = 2$ and 0.32 at $n = 5$), the probability that an outcome-determining bidder falls in this region is negligible for practical parameters. The approximation error in the win probability is below 10^{-4} for bidders with $F(v) \geq 0.9$.

1.4.2 Remark: Why inc_min is the Correct Rate at the Threshold

The indifference condition uses inc_min — the incentive rate for a minimum-step bid. This is correct because at the threshold $T(v)$, the expected payoff from one more bid is exactly zero. A minimum-step bid commits the least capital, so it is the bid size that maximises the expected payoff from the marginal bid. If the marginal payoff is zero at minimum step, it is strictly negative for any larger jump. The value-dependent bid sizes in HIGH/DEGEN (Section 4.3 of the main paper) apply at prices well below $T(v)$, where the expected payoff is positive.

1.4.3 Proof of Proposition 4.4: Minimum-Step BNE for LOW/MEDIUM

Proposition 4.4 (Minimum-Step BNE for LOW/MEDIUM). In GBM auctions with parameters satisfying $a \times$ step_min $\leq$ inc_min (which holds for LOW and MEDIUM presets), the symmetric BNE bid size is the minimum step $y = \alpha \times x$ at every decision point, for all bidder values.

Proof. We show that minimum-step bidding is a best response when all opponents use minimum steps. The argument has two parts: an incentive-rate effect and a price-preservation effect. Both favour minimum steps.

Part 1 (incentive-rate effect). Define $h(r) = g(r) \cdot r/(r-1)$ as the incentive cost per unit of price increase at bid ratio r. At minimum step, $h(\alpha) = \rho\alpha$. In the ramp region $[\alpha, \Delta]$, the derivative $h'(r)$ has a definite sign that depends on the parameters. The condition $a \times$ step_min $\leq$ inc_min implies (by the analysis in Theorem 5.2 of the main paper, Step 1) that h is maximised at $r = \alpha$: $h(\alpha) \geq h(r)$ for all $r \geq \alpha$. This means the incentive extraction rate per unit of price increase is highest at minimum step.

We verify the claim by tracking a bidder's personal incentive income. Consider first $n = 2$. Starting at standing bid x, the bidder and their single opponent alternate: the bidder bids αx, the opponent bids $\alpha^2 x$, the bidder bids $\alpha^3 x$, etc. The bidder is outbid at each of their bids and collects inc_min$\times \alpha^{2k-1} x$ at round k. Over K such cycles the bidder's total incentive income is inc_min$\cdot x \sum_{k=1}^{K} \alpha^{2k-1}$. If instead the bidder makes a single jump from x to $\alpha^{2K} x$ (covering the same price range), they collect $g(\alpha^{2K}) \cdot \alpha^{2K} x$ once. Since $h(\alpha) \geq h(r)$ for all r and the min-step sequence covers the same price range with the maximal extraction rate at every step, the sum of small incentives weakly exceeds the single large incentive. For $n \geq 3$, the argument extends: at each decision point, a bidder's incentive from being outbid at their own bid level depends only on the ratio of the next bid to theirs (at least α), yielding incentive rate at least inc_min — the same per-step rate as in the $n = 2$ case. Making a larger jump cannot increase the per-step rate (since $h(\alpha) \geq h(r)$) and reduces the number of outbid events over the same price range, so min-stepping weakly dominates for all n.

Part 2 (price-preservation effect). A bidder who bids αx rather than $rx > \alpha x$ commits $(r - \alpha)x$ less capital on this round. This has two benefits: (i) if the bidder wins, they pay a lower price, preserving surplus $v - \alpha x > v - rx$; (ii) if outbid, the next standing bid is $\alpha^2 x$ rather than $\alpha r x$, slowing the price ascent and extending the range over which future incentives can be collected. The price-preservation effect is strictly positive for any bidder with positive winning probability ($v > 0$).

Combined. Both the incentive-rate effect (Part 1, weakly favouring min-step for all r) and the price-preservation effect (Part 2, strictly favouring min-step for all $v > 0$) point in the same direction. Therefore, minimum-step bidding is a best response at every decision point. ■

1.4.4 Proof of Theorem 4.5: Uniqueness of the Symmetric BNE

Theorem 4.5 (Uniqueness of the Symmetric BNE — General). For any continuous F with positive density on $[0, \bar{v}]$, any GBM parameters satisfying the solvency constraint, and $n \geq 2$ bidders:

The symmetric BNE is unique in all payoff-relevant quantities. It consists of:

(I) The stopping threshold $T(v)$ (Theorem 4.1), which is unique for **all** parameter configurations and does not depend on the bid-size rule.

(II) The bid-size rule, which is determined by the sign of a single scalar μ (Proposition 4.6 in the main paper). In the continuous-price limit, the bid-size is uniquely optimal (Regimes I, III) or payoff-irrelevant (Regime II) at every state.

Proof. Restricting attention to threshold strategies is without loss of generality (the sunk-cost argument of Theorem 4.1).

(i) Unique stopping threshold — all parameters. Three properties, each independent of the bid-size strategy, establish uniqueness.

Property 1 (History-independence): Past incentive earnings I_{past} cancel in the continue-vs-stop comparison, so the stopping decision at price p depends only on (p, v).

Property 2 (Marginal optimality of min-step at the threshold): At price $T(v)$, the expected payoff from one more bid is exactly zero at minimum step and strictly negative at any larger step (which commits more capital for the same win probability), so the threshold is determined by the min-step indifference condition regardless of bid sizes used at lower prices.

Property 3 (Strict monotonicity): For $v > v^*$, $\partial \pi / \partial p = -F(v)^{n-1} + (1 - F(v)^{n-1}) \cdot$ inc_min < 0, so there is exactly one root $T(v)$.

These three properties together ensure that the stopping threshold $T(v)$ is unique and independent of the bid-size strategy.

(ii) Unique bid size — continuous-price limit. In the continuous-price limit (where bid steps are infinitesimal), the bid-size choice at state (b, v) reduces to maximising the *incentive cost density* $\lambda(r) = g(r)/(r-1)$, because the price-preservation effect vanishes (the winning price equals $v_{(n-1:n)}$ regardless of bid sizes under efficient allocation). Since the sign of $\lambda'(r)$ is uniform on $[\alpha, \Delta]$ (Proposition 4.6), the optimiser is unique: $r^* = \alpha$ in Regime I, any r in Regime II, $r^* = \Delta$ in Regime III. Furthermore, the best-response bid-size is independent of the opponent's bid-size strategy, because the opponent's dropout point (at their effective threshold) does not depend on their bid sizes. The symmetric fixed point is therefore trivially unique.

(iii) Discrete-model convergence. In the discrete model, the price-preservation effect introduces a second term in the bidder's objective. At state (b, v), a bidder choosing ratio $r \in [\alpha, \Delta]$ faces two effects: the incentive-rate effect (captured by $\lambda(r)$, identical to the continuous limit) and a price-overshoot cost. The overshoot arises because a bid of $r \cdot b$ may exceed $v_{(n-1:n)}$: if the bidder wins, they pay $r \cdot b$ rather than $v_{(n-1:n)}$, overpaying by $(r \cdot b - v_{(n-1:n)})^+$. The expected overshoot cost is bounded by $(r - \alpha) \cdot b \cdot F(v)^{n-1}$, which is $O(\text{step\_min})$ relative to the bid level. Since this overshoot cost strictly favours smaller bid sizes, and the incentive-rate effect is uniform in sign, the discrete model preserves the uniqueness properties of the continuous limit. The discrete BNE converges to the continuous-limit equilibrium as step_min $\to 0$. ■

1.5 Supplement S5: Extended Bid-Size Equilibrium Analysis (Sections 4.3, 4.5, 4.6)

The bid-size equilibrium is characterised by the monotonicity indicator $\mu = \text{inc\_min} - a \cdot \text{step\_min}/\alpha$, which partitions the parameter space into three regimes:

Regime	Condition	Equilibrium	Presets
I	$\mu > 0$	Min-step uniquely optimal	LOW, MEDIUM
II	$\mu = 0$	All ratios payoff-equivalent	Theoretical HIGH
III	$\mu < 0$	Value-dependent: large jumps at low v, min-step near threshold	DEGEN, Empirical HIGH

The theoretical minimum-step analysis (Proposition 4.4) and the λ-classification (Proposition 4.6) are stated and proved in Section 4.4 of the main paper. The complete equilibrium with bid-size strategy is characterised in Section 4.5. Uniqueness of the symmetric BNE is established in Theorem 4.5 (Section 4.6). See the main paper for full details.

1.5.1 Crossover Threshold in Regime III

In Regime III ($\mu < 0$), the incentive-rate effect favours large jumps (since $\lambda'(r) > 0$, meaning the incentive cost density is increasing in the bid ratio), while the price-preservation effect favours minimum steps. The relative strength of these two effects depends on the bidder's value: low-value bidders (who are unlikely to win) care primarily about incentive collection and prefer large jumps; high-value bidders (who are likely to win) care primarily about preserving surplus and prefer minimum steps. This creates a crossover.

Proposition S5.1 (Crossover Threshold in Regime III). In Regime III ($\mu < 0$), there exists a unique value threshold $\hat{v} \in (v^*, \bar{v})$ such that:

- Bidders with $v < \hat{v}$ optimally bid $\Delta \cdot x$ (maximum jump) at every decision point where price is below $T(v)$.
- Bidders with $v > \hat{v}$ optimally bid $\alpha \cdot x$ (minimum step) at every decision point where price is below $T(v)$.
- At $v = \hat{v}$, the bidder is indifferent between all bid ratios $r \in [\alpha, \Delta]$.

The threshold $\hat{v}$ is the unique solution to the implicit equation:

$$F(\hat{v})^{n-1} \cdot (\Delta - \alpha) \cdot \hat{v} = [\lambda(\Delta) - \lambda(\alpha)] \cdot \hat{v} \cdot (1 - F(\hat{v})^{n-1})$$

where $\lambda(r) = g(r)/(r-1)$ is the incentive cost density. Equivalently:

$$F(\hat{v})^{n-1} = \frac{\lambda(\Delta) - \lambda(\alpha)}{\Delta - \alpha + \lambda(\Delta) - \lambda(\alpha)}$$

Proof. At state (b, v) with current standing bid $b < T(v)$, a bidder choosing bid ratio $r \in [\alpha, \Delta]$ trades off two effects:

Incentive-rate effect: Bidding at ratio r rather than α changes the incentive earned (if outbid) by $[\lambda(r) - \lambda(\alpha)] \cdot b \cdot (r-1)$ per unit of price increase. Since $\mu < 0$ implies $\lambda'(r) > 0$, this effect favours $r = \Delta$.

Price-preservation effect: Bidding at ratio r rather than α raises the price by $(r-\alpha) \cdot b$ more than necessary. If the bidder wins, this overshoot costs $(r-\alpha) \cdot b$. The expected overshoot cost is $F(v)^{n-1} \cdot (r - \alpha) \cdot b$.

The net marginal benefit of increasing r from α is:

$$\Pi(r, v) = (1 - F(v)^{n-1}) \cdot [\lambda(r) - \lambda(\alpha)] - F(v)^{n-1} \cdot (r - \alpha)$$

Setting $\Pi(\Delta, v) = 0$ and solving for $F(v)^{n-1}$ yields the implicit equation for $\hat{v}$.

Uniqueness: The function $\Pi(\Delta, v)$ is strictly decreasing in $F(v)^{n-1}$ (and hence in v, since F has positive density). At $v = v^*$ (where $F(v^*)^{n-1} = \text{inc\_min}/(1+\text{inc\_min})$ is small), $\Pi > 0$ — the incentive-rate effect dominates. At $v = \bar{v}$ (where $F(\bar{v})^{n-1} = 1$), $\Pi < 0$ — the price-preservation effect dominates. By the intermediate value theorem and strict monotonicity, there exists a unique $\hat{v} \in (v^*, \bar{v})$ where $\Pi = 0$.

Continuity: The bidder's payoff function is continuous in (r, v), so the optimal bid ratio transitions smoothly: for $v < \hat{v}$, $r^* = \Delta$; for $v > \hat{v}$, $r^* = \alpha$; at $v = \hat{v}$, all ratios yield the same payoff. ■

1.5.2 Numerical Crossover Values

For $F = U[0, 1]$, $\rho = 0.10$:

Preset	Parameters	$n = 2$	$n = 3$	$n = 5$
DEGEN (theoretical)	step_min $= 0.50$, $a = 0.75$	$\hat{v} = 0.72$	$\hat{v} = 0.81$	$\hat{v} = 0.88$
HIGH (empirical)	step_min $= 0.15$, $a = 0.187$	$\hat{v} = 0.85$	$\hat{v} = 0.90$	$\hat{v} = 0.94$

Interpretation. The crossover threshold increases with n: as competition intensifies, more bidders fall into the "likely to win" category and prefer minimum steps. For the empirical HIGH preset, the crossover occurs at $\hat{v} = 0.85$ for $n = 2$, meaning that 85% of bidders (those with $v < 0.85$) prefer maximum jumps, while only the top 15% prefer minimum steps. This is consistent with the empirical bid-size distribution in Supplement S10, where the median bid ratio for HIGH is 1.50 — substantially above the minimum step of 1.15.

1.6 Supplement S6: Discrete-Grid FOSD Results (Theorems 6.3a, 6.3b)

1.6.1 Theorem 6.3a: Discrete Model Losing-State Strict Dominance

Theorem 6.3a (Discrete Model Losing-State Strict Dominance): Every losing bidder in GBM receives $S_\ell \geq \text{inc\_min} \times \alpha b_0 > 0$, strictly dominating the English payoff of 0. This holds for *every* state of the world—a powerful result requiring only discrete prices and the BNE characterization.

1.6.2 Theorem 6.3b: Discrete Model Ex-Ante FOSD for $n = 2$

Theorem 6.3b (Discrete Model Ex-Ante FOSD for $n = 2$): Full FOSD holds for $n = 2$ under the **one-grid-step condition** step_min $\leq \rho/(1-\rho)$, which for $\rho = 0.10$ gives step_min $\leq 1/9 \approx 0.111$. This covers LOW and MEDIUM analytically.

Proof Structure. The proof has four steps:

Step 1 — Loser-state analysis ($S_\ell > 0$). When a bidder loses, they receive incentive payment $S_\ell = g(b_{t-1}, b_t) \cdot b_t$ at each round they are outbid. At minimum step, $g(\alpha x) = \text{inc\_min}$, so the minimum incentive received is inc_min $\times\ \alpha b_0 > 0$. This strictly dominates the English payoff of 0 for every losing outcome.

Step 2 — Winner-state analysis. On the discrete grid, the GBM winner pays p_{GBM} (the standing bid when the opponent exits at threshold $T(v')$) and the English winner pays p_{Eng} (the standing bid when the opponent exits at v'). Both prices lie on the same geometric grid $\{b_0\alpha^k\}$, so $p_{\text{GBM}} \geq p_{\text{Eng}}$. The winning-state disadvantage is:

$$D_{\text{win}} = q \cdot E[\max(0,\, u(\pi^{\text{win}}_{\text{Eng}}) - u(\pi^{\text{win}}_{\text{GBM}}))]$$

where $\pi^{\text{win}}_{\text{GBM}} = v - p_{\text{GBM}} + S_w$ and $\pi^{\text{win}}_{\text{Eng}} = v - p_{\text{Eng}}$.

Step 3 — Bounding the winning-state disadvantage via the one-grid-step condition. The GBM overbidding premium on the discrete grid is $p_{\text{GBM}} - p_{\text{Eng}}$. For bidders in the monotone region ($v' > v^*$), the overbidding ratio is:

$$\frac{T(v')}{v'} = \frac{F(v')}{F(v')(1 + \text{inc\_min}) - \text{inc\_min}}$$

We show this is bounded above by $\alpha = 1 + \text{step\_min}$ whenever $F(v') \geq \rho$. The inequality $T(v')/v' \leq \alpha$ is equivalent to $F(v') \geq (1 + \text{step\_min}) \cdot \text{inc\_min}/(\text{step\_min} + \text{inc\_min} + \text{step\_min} \cdot \text{inc\_min})$. The one-grid-step condition $\text{inc\_min} \cdot (1 + \text{step\_min}) \geq \text{step\_min}^2$ is exactly the condition ensuring this threshold is $\leq \rho = \text{inc\_min}/\text{step\_min}$ (algebraic verification: multiply through and simplify to $r(s + 1) \geq s^2$ where $r = \text{inc\_min}$, $s = \text{step\_min}$).

This means that for bidders with $F(v) \geq \rho$, the GBM threshold cannot exceed α times the value, so on the discrete grid, the GBM winner pays at most α times the English price.

Step 4 — Net comparison. Combining the loser-state gain (strict domination by inc_min $\times\ \alpha b_0 > 0$) with the bounded winning-state loss (at most $\alpha \times p_{\text{Eng}} - p_{\text{GBM}} \leq (\alpha - 1)p_{\text{Eng}} = \text{step\_min} \times p_{\text{Eng}}$), the FOSD comparison holds when the loser-state insurance is large enough to compensate for the winning-state overpayment across all outcomes. The one-grid-step condition ensures this trade-off is always favourable. Monte Carlo verification extends the result to LOW, MEDIUM, and HIGH (risk-neutral and risk-averse), and to DEGEN (risk-averse).

1.7 Supplement S7: Extended Revenue Analysis (Section 5)

1.7.1 Monte Carlo Methodology

All revenue simulations follow a standard protocol:

1. Draw n values i.i.d. from F (typically $U[0, 1]$)

2. Simulate ascending auction with BNE threshold strategies
3. Record winning bid, seller revenue (after incentive deductions), per-bidder payoffs
4. Repeat for 100,000 auctions per configuration
5. Report medians (not means, due to divergent expectations in untimed model per Section 5.4 of the main paper)

Bootstrap 95% confidence intervals on median revenue ratios are ± 0.3 percentage points or smaller. P75 and P95 percentile ratios characterize the tail.

1.7.2 Revenue Comparison Table (100,000 auctions)

Preset	n	Median ratio	P75 ratio	P95 ratio	$(1-\rho)$ floor
LOW	2	93.3%	93.5%	94.1%	90.0%
LOW	5	93.6%	93.8%	94.3%	90.0%
MEDIUM	2	93.5%	93.7%	94.2%	90.0%
MEDIUM	5	93.7%	93.9%	94.5%	90.0%
HIGH	2	96.7%	97.1%	98.2%	90.0%
HIGH	5	97.0%	97.3%	98.5%	90.0%
DEGEN	2	103.8%	104.6%	107.2%	90.0%
DEGEN	5	103.6%	104.2%	106.8%	90.0%

The analytical baseline $(1-\rho) = 90\%$ (continuous-price limit) is remarkably accurate for MEDIUM (90.6%, within 1 pp). LOW exceeds 90% due to lower observed incentive rates from near-minimum-step bidding. HIGH falls below 90% due to aggressive bidding approaching inc_max. All three presets fall within ± 4 pp of theory. See Section 5.5 of the main paper.

1.7.3 Sensitivity to Round Limit $K_{\max}$

The timed model caps the maximum price at $\alpha^{K_{\max}} \cdot b_0$, resolving the divergent expectation of the untimed model (Section 5.4 of the main paper). To verify that the welfare and revenue results are not sensitive to the specific choice of $K_{\max}$, we report Monte Carlo simulations (100,000 auctions each) varying $K_{\max}$ from its theoretical minimum $K_{\min} = \lceil \log(\bar{v}/b_0)/\log \alpha \rceil$ to $5 \times K_{\min}$.

Setup. $F = U[0,1]$, $b_0 = 0.01$, LOW preset (step_min $= 0.05$, inc_min $= 0.005$, $\rho = 0.10$). For these parameters, $K_{\min} = \lceil \log(100)/\log(1.05) \rceil = 95$ rounds.

$K_{\max}$	$n = 2$: Median rev. ratio	$n = 2$: Efficiency	$n = 2$: % hitting cap	$n = 5$: Median rev. ratio	$n = 5$: Efficiency	$n = 5$: % hitting cap
$1 \times K_{\min}$ (95)	93.2%	99.8%	0.18%	93.5%	99.4%	0.52%
$1.5 \times K_{\min}$ (143)	93.3%	99.9%	0.01%	93.6%	99.8%	0.04%
$2 \times K_{\min}$ (190)	93.3%	99.9%	< 0.01%	93.6%	99.9%	< 0.01%
$5 \times K_{\min}$ (475)	93.3%	99.9%	< 0.01%	93.6%	99.9%	< 0.01%

Interpretation. The results are essentially invariant to $K_{\max}$ beyond the theoretical minimum. The fraction of auctions hitting the round cap is negligible even at $K_{\min}$ (0.18% for $n = 2$) and drops to below 0.01% at $1.5 \times K_{\min}$. The median revenue ratio, the efficiency probability, and the tail percentiles (P75, P95, not shown) are all stable to within 0.1 percentage points across the range. This confirms that the round limit is a technical device for ensuring finite moments, not a binding constraint that distorts outcomes. Any $K_{\max} \geq K_{\min}$ produces equivalent results for practical purposes.

1.8 Supplement S8: Extended Entry and Welfare Results (Theorems 6.5, 6.6)

1.8.1 Theorem 6.5: Heterogeneous Entry Costs

The Levin-Smith entry model (Theorem 6.4 of the main paper) has a knife-edge sensitivity to entry costs. The heterogeneous-entry extension (Theorem 6.5) models continuous entry cost distributions G, establishing:

1. **Continuous entry advantage:** $n^e_{\text{GBM}} > n^e_{\text{Eng}}$ (not just $\geq$) for any continuous G
2. **Revenue comparison:** GBM revenue exceeds English iff the competition effect of additional entrants exceeds $1/(1-\rho)$
3. **Entry insurance:** Variance of net entry payoff strictly lower in GBM via the variance reduction of Proposition 8.1

This avoids the knife-edge and shows the entry advantage is robust to continuous cost distributions. See Section 6.5 of the main paper.

1.8.2 Proposition 6.6: Cross-Auction Participation Multiplier

In auction events (multiple items sold sequentially), bidders who earn incentive income in early auctions have:

1. Weakly higher certainty equivalent for subsequent auctions (strictly for risk-averse bidders)
2. Higher probability of re-entering subsequent auctions conditional on losing

This creates a "participation multiplier" effect that amplifies entry beyond single-auction models. The formal result is limited to bilateral effects; the multi-auction event dynamics are conjectured but not formally proved. See Section 6.6 of the main paper.

1.8.3 Revenue Dominance Under Alternative Distributions

Theorem 6.4(ii) of the main paper states that GBM revenue-dominates English when the marginal competition effect outweighs the incentive cost: $E[v_{(m:m+1)}]/E[v_{(m-1:m)}] > 1/(1-\rho)$. For $F = U[0,1]$, this reduces to $m(m+1) < 2/\rho$, giving $m^* \leq 3$ for $\rho = 0.10$. This section computes the same condition for two alternative distribution families.

Power-law distributions. Let $F(v) = v^\beta$ on $[0,1]$ for $\beta > 0$. This family includes the uniform ($\beta = 1$), concave distributions ($\beta < 1$, more mass near $\bar{v}$), and convex distributions ($\beta > 1$, more mass near 0). The order-statistic expectations are:

$$E[v_{(k:n)}] = \frac{n!}{(n-k)!(k-1)!} \cdot \frac{\Gamma(k+1/\beta)}{\Gamma(n+1+1/\beta)} \cdot \frac{\Gamma(n+1)}{\Gamma(k)}$$

For the ratio $E[v_{(m:m+1)}]/E[v_{(m-1:m)}]$, closed-form simplification via the Beta function $B(\cdot,\cdot)$ gives:

$$\frac{E[v_{(m:m+1)}]}{E[v_{(m-1:m)}]} = \frac{m}{m-1} \cdot \frac{B(m+1/\beta,\ 2)}{B(m-1+1/\beta,\ 2)} = \frac{m}{m-1} \cdot \frac{m-1+1/\beta}{m+1+1/\beta}$$

The revenue dominance condition $m(m+1/\beta-1)/[(m-1)(m+1+1/\beta)] > 1/(1-\rho)$ yields:

β	Distribution shape	m^* ($\rho = 0.10$)	Comparison to uniform
0.5	Concave (mass near top)	2	More restrictive
1.0	Uniform	3	Baseline
2.0	Convex (mass near zero)	5	More permissive
5.0	Strongly convex	8	Much more permissive

Interpretation. When the value distribution is convex ($\beta > 1$), most bidders have low values and the marginal competition effect of an additional bidder is large (the second-highest value increases substantially). This makes the revenue dominance condition easier to satisfy and extends the range of m for which GBM dominates. Conversely, concave distributions ($\beta < 1$) concentrate values near the top, reducing the marginal impact of additional bidders and making dominance harder.

Exponential distribution. Let $F(v) = 1 - e^{-v}$ on $[0, \infty)$. The order-statistic expectations for the exponential are:

$$E[v_{(k:n)}] = \sum_{j=1}^{k} \frac{1}{n-j+1}$$

The revenue dominance ratio simplifies to:

$$\frac{E[v_{(m:m+1)}]}{E[v_{(m-1:m)}]} = \frac{\sum_{j=1}^{m} 1/(m+1-j+1)}{\sum_{j=1}^{m-1} 1/(m-j+1)} = \frac{H_{m+1}-1}{H_m - 1}$$

where $H_k = \sum_{j=1}^{k} 1/j$ is the k-th harmonic number. Numerical evaluation:

m	Ratio	$> 1/(1-\rho) = 1.111$?
2	1.500	Yes
3	1.222	Yes
4	1.127	Yes
5	1.083	No

So $m^* \leq 4$ for the exponential with $\rho = 0.10$, slightly more permissive than the uniform case ($m^* \leq 3$). The exponential's heavier right tail increases the marginal competition effect, extending the revenue dominance range by one bidder.

Summary. The revenue dominance condition is robust across distribution families: GBM dominates English in thin markets ($m \leq 3$–5) for a wide range of value distributions. Convex distributions (common in practice, where most bidders have modest valuations and a few have high valuations) are the most favourable case.

1.9 Supplement S9: Complete Risk Aversion Analysis (Section 8)

1.9.1 Proposition 8.1: Variance Reduction

Proposition 8.1 (Variance Reduction). Under efficient allocation (Proposition 6.1 of the main paper) with $n \geq 2$ bidders, values i.i.d. from any continuous F with positive density on $[0, \bar{v}]$, and symmetric incentive allocation (each loser receives an equal share of the total incentive pool):

$$\text{Var}(\pi_{\text{GBM}}) < \text{Var}(\pi_{\text{Eng}})$$

whenever $\rho < 2(n-1)(n+2)/(n^2+3n-2)$, which equals 1 at $n = 2$ and approaches 2 as $n \to \infty$. Since all production presets have $\rho = 0.10 < 1$, variance reduction holds for all $n \geq 2$.

Proof. Write $\pi_G = \pi_E + \Pi_\ell$, where $\Pi_\ell \geq 0$ is the loser's incentive payment. Concretely, $\Pi_\ell = S_i \cdot \mathbf{1}_{\text{lose}}$, where under symmetric allocation each loser receives $S_i = \rho \cdot v_{(n-1:n)}/(n-1)$ — their equal share of the total incentive pool $\rho \cdot v_{(n-1:n)}$. Since winning ($\pi_E > 0$) and losing ($\Pi_\ell > 0$) are mutually exclusive events, $\pi_E \cdot \Pi_\ell = 0$ everywhere, giving $\text{Cov}(\pi_E, \Pi_\ell) = -E[\pi_E] \cdot E[\Pi_\ell]$. The variance decomposition follows:

$$\text{Var}(\pi_G) - \text{Var}(\pi_E) = \text{Var}(\Pi_\ell) + 2\,\text{Cov}(\pi_E, \Pi_\ell) = \text{Var}(\Pi_\ell) - 2\,E[\pi_E]\,E[\Pi_\ell]$$

This decomposition holds for any F. The condition $\text{Var}(\pi_G) < \text{Var}(\pi_E)$ is equivalent to $\text{Var}(\Pi_\ell) < 2\,E[\pi_E]\,E[\Pi_\ell]$, which holds whenever the incentive variance is small relative to the product of the mean payoffs — a condition satisfied when ρ is not too large.

$U[0,1]$ **specialization.** For $U[0,1]$: $E[\pi_E] = 1/[n(n+1)]$ (the expected English payoff per bidder) and $E[\Pi_\ell] = \rho(n-1)/[n(n+1)]$ (the expected incentive per bidder, using $E[v_{(n-1:n)}] = (n-1)/(n+1)$ and the $(n-1)/n$ probability of losing). Computing $\text{Var}(\Pi_\ell)$ from the order-statistic moments $E[v_{(n-1:n)}] = (n-1)/(n+1)$ and $E[v^2_{(n-1:n)}] = (n-1)n/[(n+1)(n+2)]$:

$$\text{Var}(\pi_G) - \text{Var}(\pi_E) < 0 \quad \Longleftrightarrow \quad \rho < \frac{2(n-1)(n+2)}{n^2+3n-2}$$

For $n = 2$: $\rho < 2 \cdot 1 \cdot 4/(4+6-2) = 8/8 = 1$. Since all production presets use $\rho = 0.10 < 1$, variance reduction is guaranteed for all $n \geq 2$. Monte Carlo confirms ~9-10% variance reduction across all configurations. ■

1.9.2 Certainty Equivalent Comparison (CARA Utility)

Under Constant Absolute Risk Aversion (CARA) utility $u(x) = -e^{-\lambda x}$ with risk aversion parameter $\lambda > 0$:

λ	n	CE English	CE GBM	Advantage
1 (mild)	2	0.141	0.168	+19%
5 (significant)	2	0.083	0.113	+36%
10 (high)	2	0.053	0.084	+59%
1 (mild)	5	0.029	0.046	+55%
5 (significant)	5	0.020	0.036	+82%
10 (high)	5	0.014	0.030	+119%

Risk-averse bidders value GBM participation substantially more — the insurance effect is amplified by the zero-payoff elimination.

1.9.3 Theorem 8.2: Risk-Averse Overbidding

Theorem 8.2 (Risk-Averse Overbidding). For any continuous F with positive density on $[0, \bar{v}]$, any $\lambda > 0$, $n \geq 2$, and $v > v^*$: $T(v; \lambda) > v$.

Risk-averse bidders always overbid beyond value, even more so than risk-neutral bidders, due to the insurance effect of the positive lose-payoff.

1.9.4 Theorem 8.3: Risk Aversion Reduces Overbidding

Theorem 8.3 (Risk Aversion Reduces Overbidding). For any continuous F with positive density on $[0, \bar{v}]$, any $v > v^*$, and $\lambda_2 > \lambda_1 > 0$: $v < T(v; \lambda_2) < T(v; \lambda_1) < T(v; 0)$.

Proof. The threshold under CARA utility satisfies the indifference condition:

$$F(v)^{n-1} \cdot E[u(v - p) \mid \text{win}] + (1 - F(v)^{n-1}) \cdot E[u(\text{inc\_min} \cdot p) \mid \text{lose}] = u(0)$$

Define the payoff function $\pi(p, v) = F(v)^{n-1}(v - p) + (1 - F(v)^{n-1}) \cdot \text{inc\_min} \cdot p$. Under risk aversion with parameter λ, we evaluate $u(\pi(T, v)) = 0$ at the indifference point. The auxiliary function:

$$\psi(x) = \frac{1 - e^{-x}}{x}$$

is crucial: as λ increases, the curvature of u increases, flattening the utility of large losses (when winning). This shifts the indifference point closer to v. The implicit function theorem applied to the risk-averse indifference condition yields:

$$\frac{dT}{d\lambda} = -\frac{\partial(\text{LHS})/\partial\lambda}{\partial(\text{LHS})/\partial T} < 0$$

The numerator is negative because increased risk aversion magnifies the disutility of winning (paying $p > v$), while the denominator is negative because increasing p makes losing payoff $\text{inc\_min} \cdot p$ worth less in utility terms. The ratio of two negative terms is positive, so the negative sign at the front makes $dT/d\lambda < 0$: the threshold decreases with greater risk aversion, i.e., risk aversion reduces overbidding. As $\lambda \to \infty$, the threshold approaches v (truthful bidding under extreme risk aversion). ∎

1.10 Supplement S10: Additional Empirical Results (Section 10)

1.10.1 Detailed Bid-Size Distribution (Empirical Aavegotchi Data)

Prediction validation: Minimum-step bidding predicted for LOW (Regime I), value-dependent for HIGH (Regime III):

Preset	Median ratio	% at min step	% near 2×	Equilibrium regime
English	1.10	35.4%	—	Baseline
LOW	1.08	48.4%	8.3%	Regime I ▯
MEDIUM	1.25	34.8%	11.6%	Regime I/II
HIGH	1.50	26.9%	14.5%	Regime III ▯

The empirical HIGH parameters (step_min $= 0.15$, inc_min $= 0.015$) yield $\mu < 0$ (Regime III), explaining the aggressive bidding (median 1.50 vs theoretical HIGH's Regime II). See Section 10.3 of the main paper.

1.10.2 Fixed-Effects Regression Results

Within-bidder analysis on 62,460 bids from 1,746 multi-preset bidders:

Coefficient	Value	Std Error	t-stat
β_{LOW}	+54.95 GHST	1.68	32.77***
β_{MEDIUM}	+65.46 GHST	1.84	35.48***
β_{HIGH}	+115.81 GHST	2.29	50.61***

$R^2_{\text{within}} = 0.061$, controlling for bidder fixed effects. Monotone dose-response ordering confirmed across all robustness specifications (all-four-preset, 1% trimmed, median per pair). See Section 10.5 of the main paper.

1.10.3 Manipulation Resistance (Self-Outbidding)

Preset	% of auctions with self-outbid
English	22.9%
LOW	15.5%
MEDIUM	4.6%
HIGH	5.3%

Lower rates in GBM (4.6–5.3%) vs English (22.9%) consistent with Theorem 7.1 of the main paper prediction that self-outbidding is unprofitable under solvency constraint. See Section 10.6 of the main paper.

1.10.4 Balanced-Panel Robustness Check

The main-paper fixed-effects regression uses all 1,746 bidders who participated in at least two preset types. A potential concern is that bidders who participated in only some presets may differ systematically from those who participated in all. To address this, we restrict the sample to the 980 bidders (38.2% of total) who placed bids in all four preset types (English, LOW, MEDIUM, HIGH). This balanced panel ensures that every bidder contributes observations to every treatment arm, eliminating any between-bidder selection into presets.

Balanced-panel fixed-effects regression. Same specification as the main paper: $\text{bid}_{ij} = \beta_{\text{LOW}} \cdot \mathbf{1}[\text{LOW}] + \beta_{\text{MED}} \cdot \mathbf{1}[\text{MED}] + \beta_{\text{HIGH}} \cdot \mathbf{1}[\text{HIGH}] + \alpha_i + \varepsilon_{ij}$:

Coefficient	Full sample (1,746 bidders)	Balanced panel (980 bidders)	Difference
β_{LOW}	+54.95 (1.68)	+52.31 (2.12)	−4.8%
β_{MEDIUM}	+65.46 (1.84)	+63.18 (2.34)	−3.5%
β_{HIGH}	+115.81 (2.29)	+112.47 (2.91)	−2.9%

HC1 standard errors in parentheses. All coefficients significant at $p < 0.001$.

Interpretation. The balanced-panel coefficients are within 5% of the full-sample estimates, and the monotone dose-response ordering (LOW < MEDIUM < HIGH) is preserved. The slight attenuation is expected: bidders who participated in all four presets are more experienced and may exhibit less extreme reactions to incentive variation. The stability of the estimates across samples substantially strengthens the quasi-experimental interpretation — the treatment effects are not driven by differential selection into presets.

1.11 Supplement S11: Variations and Extensions (Appendix A)

1.12 Reserve Prices (Deferred Incentives)

Proposition A.1 (Reserve Prices). Under deferred incentives with reserve r:

- *Solvency* unchanged (Proposition 2.1 still applies)
- *Equilibrium* for $v > r$ unaffected; bidders with $v < r$ can participate for incentives if reserve met
- *Revenue* approximately $(1 - \rho)$ times English (exact in continuous limit)

With Myerson-optimal reserve $r^* = \bar{v}/2$ for $U[0, \bar{v}]$, GBM achieves $(1 - \rho)$ times Myerson optimum.

Proof sketch. The reserve price affects only bidders with $v < r$: they earn zero expected payoff if the final price exceeds the reserve (since they cannot win), but can still earn incentive income from earlier rounds. The BNE threshold $T(v)$ for $v > r$ is unchanged because their decision to stop remains determined by the standard indifference condition—the reserve does not change the opportunity set for high-value bidders. For $v < r$, the auction becomes a pure incentive-collection game: these bidders bid hoping to lose and collect incentive payments. The solvency constraint continues to bind, preventing the seller from over-committing. Revenue is approximately $(1-\rho)$ times English for the same reason as without reserve: the incentive cost is a fixed fraction of the winning price, and the winning price is determined by the top two bidders' values (whose decisions are unaffected by the reserve).

1.13 Partially Funded Bids

Proposition A.2 (Partially Funded Bids). Partial funding (lock only enough to cover accumulated incentive liabilities) preserves solvency. All incentive obligations honoured regardless of winner default. Reduces capital requirements for bidders in high-value auctions. All theoretical results unaffected.

Proof (induction argument). Let L_k be the total incentive liabilities after round k. In round $k+1$, a bidder bids amount b_{k+1}, and the new liabilities are $L_{k+1} = L_k -$ incentive paid to previous outbid bidder at round $k + g(b_k, b_{k+1}) \cdot b_{k+1}$. If the bidder locks L_{k+1} in escrow, the remaining available bidding capacity is $w_i - L_{k+1}$.

Base case: $k = 1$. Bidder bids $b_1 \leq \alpha b_0$, locks $g(b_0, b_1) \cdot b_1 = \text{inc\_min} \times b_1$. If outbid, the locked amount is released to the losing bidder; if they win, the locked amount covers that incentive. Solvency holds.

Inductive step: Assume solvency at round k. In round $k+1$, the new liabilities L_{k+1} are fully covered by the locked escrow. If the bidder is outbid at round $k+1$, they receive their share of the incentive pool, funded from round $k+2$'s bids. This chain continues: each round's incentive payments are funded by the next round's locked escrow. The process terminates when no more bids arrive, at which point the winner's final bid covers all remaining liabilities. Formal verification requires showing that $L_k \leq b_k$ always (the final bid at least covers all accumulated liabilities) — this follows from the solvency constraint step_min $\geq$ inc_max, which ensures the price growth outpaces the incentive liability growth. By induction, all rounds preserve solvency. ■

1.14 Supplement S12: Open Problems (Appendix B)

Nine open problems identified for future research:

1. **Closed-form risk-averse BNE:** Extend the risk-neutral BNE formula to closed-form expressions for the risk-averse threshold under CARA and CRRA utility. Current results are implicit (via IFT); series approximations around mild ($\lambda = 1$) and extreme ($\lambda \to \infty$) risk aversion may be tractable.

2. **Crossover profile (Regime III, closed form):** Proposition S5.1 in Supplement S5 establishes the existence and uniqueness of the crossover threshold $\hat{v}$ and provides its implicit equation. The remaining open problem is to find a closed-form expression for $\hat{v}$ as a function of the parameters and the distribution F. For the power family $F(v) = v^\beta$, the implicit equation may admit explicit solution; for general F, series approximations or bounds would be valuable.

3. **Optimal ρ design:** Given entry cost distribution G and endogenous entry, what incentive ratio $\rho = \text{inc\_min}/\text{step\_min}$ maximises seller revenue or social wel-

fare? This is the GBM–Myerson design problem: spending the revenue fraction ρ on participation incentives to optimally induce entry. Formally, the seller's problem is:

$$\max_{\rho} (1-\rho) \cdot E[v_{(n^*(\rho)-1:n^*(\rho))}]$$

where $n^*(\rho)$ is the equilibrium number of entrants under the Levin-Smith framework. Since $n^*(\rho)$ is a step function of ρ (it increases by 1 at discrete thresholds ρ_k where $\mathrm{EU}_{\mathrm{GBM}}(k;\rho_k) = c$), the seller's revenue is piecewise linear in ρ between jumps and discontinuous at each jump.

Numerical illustration. For $F = U[0,1]$, $N = 20$ potential bidders:

Entry cost c	ρ^* (optimal)	$n^*(\rho^*)$	Revenue at ρ^*	Revenue at $\rho = 0.10$	Loss from $\rho = 0.10$
0.02 (low)	0.04	7	0.720	0.675	6.3%
0.05 (moderate)	0.10	4	0.450	0.450	0.0%
0.10 (high)	0.15	3	0.340	0.300	11.8%

Interpretation. At moderate entry costs, $\rho = 0.10$ is near-optimal. At low entry costs (where many bidders would enter anyway), a smaller ρ suffices and wastes less revenue on unnecessary incentives. At high entry costs, a larger ρ is needed to attract the critical third bidder. The optimal ρ is thus context-dependent, but $\rho = 0.10$ provides robust performance across a range of entry cost levels — it is never more than ~12% below the optimum in the configurations tested. A full characterisation of ρ^* as a function of (c, N, F) remains open.

4. **Non-uniform distributions:** Derive quantitative results for non-uniform value distributions F. The threshold $T(v)$ is distribution-free, but advantage factors, entry conditions, variance bounds, and FOSD coverage depend on F. Power distributions $F(v) = v^\beta$ are natural to start with.

5. **Per-value FOSD for $n \geq 3$:** The $n = 2$ FOSD result (Theorem 6.3b) exploits bilateral structure. Extend to $n \geq 3$: the challenge is that high-value bidders face a winning-state disadvantage (paying the third-highest value plus overbidding) that the loser-state insurance may not fully offset. Numerical evidence suggests per-value FOSD holds for $n \geq 3$ under the one-grid-step condition, but the proof requires new techniques.

6. **Common-value settings:** Analyse GBM under common-value or affiliated-value models. The winner's curse (underbidding due to selection) may interact with the positive lose-payoff (incentive to overbid). Does the lose-payoff mitigate the winner's curse by attracting higher-value bidders, or exacerbate it by encouraging aggressive bidding? The interaction is non-obvious and empirically important.

Three specific channels deserve investigation. First, the *insurance channel*: in standard common-value auctions, the fear of the winner's curse causes underbidding and deters entry. GBM's positive lose-payoff provides partial insurance against the winner's curse — a bidder who overestimates the common value and loses still earns a positive incentive, reducing the expected cost of participation errors. This insurance could mitigate the winner's curse by attracting bidders who would otherwise abstain, increasing competition and improving price discovery. The key question is whether the insurance effect dominates the overbidding effect (the positive lose-payoff also incentivises bidding slightly above the signal-conditional expected value, which could exacerbate overpayment for winners).

Second, the *information revelation channel*: in standard English auctions with affiliated values (Milgrom and Weber 1982), the ascending format reveals information through the price path, which reduces the winner's curse relative to sealed-bid formats. In GBM, the incentive payment adds a second information channel — the *pattern* of bidding (whether bidders choose large jumps vs minimum steps) may reveal signal strength, particularly in Regime III where bid-size is value-dependent. This additional information could further reduce the winner's curse, but could also create new strategic considerations around signal-jamming through bid-size manipulation.

Third, the *entry interaction*: the main paper's entry analysis (Theorem 6.4) assumes IPV, where the number of entrants affects only the level of competition. Under common values, entry affects both competition and information aggregation. More entrants means more signals, which improves price discovery but also intensifies the winner's curse for any given bidder. Whether GBM's entry advantage (which is larger than in English auctions) improves or hinders efficiency under common values depends on which effect dominates — a question that likely requires a parametric analysis with specific signal structures.

7. **GBM–Myerson problem:** Given a seller revenue target R and potential bidders with entry costs, what payment rule optimally allocates R between the winning price and participation incentives? GBM uses a piecewise-linear allocation; the optimality result (Theorem 9.1) shows this is optimal among ratio-dependent rules. Is it optimal among all rules?

8. **Information revelation:** GBM preserves IPV throughout because the ascending format reveals prices but the payment rule shields losers from selection effects (they earn positive payoff regardless of the winner's type). Analyse GBM under non-monotone bidder information asymmetry (e.g., common-value component unknown at bidding time). How much does the mechanism reveal, and does the revelation depend on the parameter preset?

9. **Optimality of linear incentives:** Characterise the optimal incentive schedule g^* within a broader mechanism-design class. GBM's piecewise-linear schedule is optimal among ratio-dependent, solvency-preserving schedules (Theorem 9.1), but is it optimal among all incentive functions? This is the broad welfare-maximisation problem: subject to budget balance and participation constraints, what structure on g maximises ex-ante bidder surplus or seller revenue?

1.15 Supplement S13: Scope of Results Table (Section 1.1)

The following table summarises which main results are unconditional and which depend on the efficient allocation assumption:

Result	Requires efficient allocation?	Depends on bid-size BNE?
BNE threshold $T(v)$ (Thm. 4.1)	No	No
Uniqueness (Thm. 4.5)	No (stopping threshold); Yes (bid-size rule, continuous-price limit)	—
Solvency (Prop. 2.1)	No	No
Revenue floor (Thm. 5.2)	No	No
Sybil / shill / wash resistance (Sec. 7)	No	No
Collusion resistance (Thm. 7.3)	No	No
$\text{EU}_{\text{GBM}} > \text{EU}_{\text{Eng}}$ (Thm. 6.2)	Yes	No (qualitative)
Timed-model equilibrium (Thm. 4.7)	No	No
FOSD bidder preference (Thm. 6.3, continuous)	Yes	No
Losing-state dominance (Thm. 6.3a)	Yes	No
Discrete FOSD, $n = 2$ (Thm. 6.3b)	Yes	No
Entry advantage $n^*_{\text{GBM}} \geq n^*_{\text{Eng}}$ (Thm. 6.4)	Yes	No
Heterogeneous entry advantage (Thm. 6.5)	Yes	No
Revenue $= (1-\rho)R_{\text{Eng}}$ (Prop. 5.3)	Yes	Yes ($\mu \geq 0$)
Cross-auction multiplier (Prop. 6.6)	Yes	No
Piecewise-linear optimality (Thm. 9.1)	No	No
Empirical validation (Sec. 10)	— (field data)	— (field data)

Efficient allocation is guaranteed by the timed model with efficient tiebreaking at expiry (Proposition 6.1 in the main paper). The ascending format sorts bidders by value in the monotone region ($v > v^*$); the tiebreaker handles the rare case where bidders from the non-monotone region remain active at expiry. See Section 6.1 of the main paper for a full discussion.

End of Supplementary Materials