

The GBM Auction

An Incentivised Ascending Auction Mechanism

E. Bessire

March 2026

Abstract. The GBM auction is a time-limited ascending auction in which outbid participants receive a positive incentive payment funded by subsequent bids. Under symmetric independent private values, the closed-form BNE threshold produces negligible overbidding ($< 0.1\%$ for outcome-determining bidders), and the solvency constraint makes self-outbidding, shill bidding, and wash trading provably non-profitable. The mechanism redistributes a fraction ρ of seller revenue to losers: in the continuous-price limit every bidder's GBM payoff first-order stochastically dominates the English payoff for all $n \geq 2$, all continuous i.i.d. private-value distributions, and all increasing utility functions; on the discrete grid, losing-state dominance holds universally and full FOSD is proved for $n = 2$. An endogenous entry analysis shows that GBM induces strictly more entry than English auctions for any continuous entry cost distribution; in thin markets ($n_{\text{Eng}}^* \leq 3$) the additional competition more than compensates for the incentive cost. Among all ratio-dependent, solvency-preserving loser-subsidy ascending auctions, GBM's piecewise-linear incentive schedule uniquely minimises the worst-case revenue floor while eliminating bid-size gaming. An empirical analysis of 11,900 auctions confirms the theoretical predictions on overbidding, revenue share, and entry effects. Since its initial deployment, the GBM auction mechanism has been used in over 90,000 auctions representing more than \$200,000,000 in bidding volume, with over \$6,000,000 earned in incentives by bidders.

Contents

1	Introduction	2
1.1	Overview and Contribution	2
1.2	History and Deployment	3
1.3	Related Literature	3
1.4	Notation	3
2	The GBM Auction	5
2.1	Auction Format	5
2.2	The Incentive Function	5
2.3	Payoffs	6
2.4	The Solvency Constraint	7
2.5	Derived Constants	7
2.6	Worked Example	7
2.7	Production Parameter Presets	8
2.8	Model Assumptions and Scope	9
3	Relationship to Classical Auctions	10
3.1	GBM as a Generalisation of the English Auction	10
3.2	No Dominant Strategy	10
3.3	Revenue Equivalence Theorem Does Not Apply	10
3.4	Scale Invariance	11
4	Equilibrium Analysis	11
4.1	The Ascending Auction BNE	11
4.2	Bounds on Overbidding	12
4.3	Overbidding Bound	12
4.4	Bid-Size Strategy	12
4.5	The Complete Equilibrium	13
4.6	Uniqueness	13
4.7	Equilibrium of the Timed Model	14
5	Revenue Analysis	15
5.1	Equilibrium Revenue	15
5.2	Revenue Floor	15
5.3	Revenue Floors Across Presets	16
5.4	The Divergent Expectation (Untimed Theoretical Model)	17
5.5	Comparison with English Auction Revenue	17
6	Welfare Analysis and Mechanism Comparison	18
6.1	Efficient Allocation	18
6.2	Bidder Preference (General n)	18
6.3	Payoff Dominance (FOSD)	19
6.4	Comparison with RET Benchmark and Myerson Optimal	21
6.5	Endogenous Entry	21
6.6	Dynamic GBM and Cross-Auction Effects	22
7	Manipulation Resistance	23

7.1 Sybil Attacks (Self-Outbidding)	23
7.2 Shill Bidding	24
7.3 Wash Trading	24
7.4 Collusion and Bidder Rings	24
7.4.1 Classical Bid Suppression and Incentive Farming	24
7.4.2 The Defection Premium	24
7.5 Why the Solvency Constraint Is the Key	25
8 Risk Aversion	25
9 Mechanism-Design Characterisation	25
9.1 The Design Space	26
9.2 Optimality of Piecewise-Linear Schedules	26
10 Empirical Evidence	27
10.1 Experimental Design	27
10.2 Revenue Comparison	28
10.3 Within-Bidder Analysis	28
10.4 Summary	29
Acknowledgments	29
Conflict of Interest Disclosure	29
References	29

1 Introduction

1.1 Overview and Contribution

The Gonnaud-Bessire-McDonough (GBM) auction is a time-limited, open ascending auction that modifies the English auction in one structural way: when a bidder is outbid, they receive their full bid back *plus* a positive incentive payment. Each incentive is determined at the moment a bidder places their bid — it depends only on the bid amount and its jump size relative to the previous standing bid, not on any future event. While the incoming bid provides the liquidity to settle each incentive obligation as it arises, the economic cost is borne by the seller: at auction close, the seller receives the winning bid minus all incentives distributed to outbid participants. A bidder either wins the auction or leaves with more money than they started with.

This single modification has far-reaching consequences for the game theory of the auction. The positive lose-payoff transforms the auction from a zero-sum game (one winner, all others earn nothing) into a positive-sum game (the winner gets the item, every other participant earns money). This paper analyses these consequences formally.

The central economic argument is one of *participation over optimisation*. Bulow and Klemperer (1996) showed that an English auction with $n + 1$ bidders generates more revenue than Myerson’s optimal mechanism with n bidders — additional participation dominates optimal reserve-price design. Klemperer (2002) argues that attracting sufficient entry is the single most important problem in practical auction design, with thin markets (few competing bidders) being the norm rather than the exception across spectrum, procurement, and asset auctions. GBM is a mechanism specifically designed to generate those extra bidders: it sacrifices a bounded fraction of per-auction revenue to make losing profitable, thereby attracting additional participants. Crucially, the mechanism does not inflate prices — the equilibrium overbidding is negligible (Theorem 4.3). If a GBM auction produces a higher price than an English auction, it is because more people showed up to bid, not because anyone was induced to bid beyond their means. For auctions with 2–3 bidders — the thin-market regime where attracting additional entrants has the largest marginal effect — GBM participants more than compensate for the incentive cost, making GBM revenue-superior to the English auction (Section 6.5). In auction events where multiple items are sold, the effect is amplified by cross-auction incentive reinvestment.

Summary of main results:

1. **Equilibrium and manipulation resistance (Sections 4, 7).** The closed-form BNE threshold $T(v) = v \cdot F(v)^{n-1} / [F(v)^{n-1}(1 + \text{inc_min}) - \text{inc_min}]$ produces negligible overbidding ($< 0.1\%$ for outcome-determining bidders). The solvency constraint $\text{step_min} \geq \text{inc_max}$ simultaneously guarantees financial soundness and makes self-outbidding, shill bidding, and wash trading provably non-profitable. Collusion is harder: defection from a bidder ring is strictly more attractive in GBM than in English (Theorem 7.3).
2. **Welfare and entry (Section 6).** In the continuous-price limit, every bidder’s GBM payoff first-order stochastically dominates the English payoff for all n , all

F , and all increasing utility functions (Theorem 6.3). On the discrete grid, every losing bidder is strictly better off in every state (Theorem 6.3a), and full FOSD holds for $n = 2$ under a mild parametric condition (Theorem 6.3b). The bidder advantage is funded by a seller revenue reduction of $\rho = \text{inc_min}/\text{step_min}$. Under endogenous entry with heterogeneous costs, GBM attracts strictly more bidders (Theorems 6.4–6.5); in thin markets ($n_{\text{Eng}}^* \leq 3$), the additional competition more than compensates, making GBM revenue-superior.

3. **Mechanism-design characterisation (Section 9).** Among all ratio-dependent, solvency-preserving loser-subsidy ascending auctions, GBM’s piecewise-linear incentive schedule uniquely minimises the worst-case revenue floor while eliminating bid-size gaming (Theorem 9.1).
4. **Empirical validation (Section 10).** 11,900 identical-item auctions pseudo-randomly assigned to English and three GBM presets. A bidder fixed-effects regression establishes a quasi-experimental dose-response: the same bidder bids \$232 more under HIGH than English ($t = 50.6$). Revenue exceeds English by 15–75% despite the incentive cost.

1.2 History and Deployment

The GBM auction was conceived in 2018 by Guillaume Gonnaud, Edouard Bessire, and Hugo McDonough. The mechanism is patent-pending (US20210174432A1, priority date 7 August 2018). First deployed in 2020 on the Cryptograph platform, the system has since been used in over 90,000 auctions across multiple blockchain networks and fiat-currency platforms, processing \$200,000,000 in bidding volume with \$6,000,000 paid to bidders in incentive rewards.

1.3 Related Literature

The GBM auction intersects classical theory (Bulow-Klemperer on entry, Myerson on optimal mechanisms), mechanism design, and participation-encouraging formats. Core references: Bulow-Klemperer (1996), Klemperer (2002) on entry importance, Levin-Smith (1994) on equilibrium entry, Graham-Marshall-McAfee-Pesendorfer on collusion, Maskin-Riley (1984) on risk aversion. The closest antecedent is the Amsterdam Auction (Goeree and Offerman 2004), which also encourages entry through positive loser payoffs. GBM differs in three respects: it is purely ascending (no sealed-bid stage), its incentive is determined at bid time rather than by the final outcome, and the solvency constraint provides automatic budget balance without requiring ex-ante premium calibration.

Extended comparisons with the Amsterdam Auction, penny auctions, contests, and other mechanisms appear in Supplementary Materials, Section S3.

1.4 Notation

Symbol	Definition
n	Number of bidders
v_i	Private value of bidder i , drawn independently from the same distribution F (typically $U[0, 1]$)
K	Number of bids in a sequence (the winning bid is b_K). Always $K \leq K_{\max}$. Distinct from n (number of bidders).
$K_{\max}$	Round limit of the timed model (Section 2.8). The auction terminates after at most $K_{\max}$ rounds.
b_t	The t -th bid in the ascending sequence ($b_1 < b_2 < \dots < b_K$). The index t denotes temporal position in the bid sequence; bidder indices use i .
b_0	Standing bid before the first bid (= 0).
x	Previous standing bid (at the moment a new bid is placed)
y	New bid ($y \geq \alpha \times x$ required)
$g(x, y)$	Incentive rate: the fraction of y that the bidder receives if outbid
inc_max	Maximum incentive rate (cap of g)
inc_min	Minimum incentive rate (value of g at minimum step)
a	Accelerator: controls how fast g ramps from inc_min to inc_max
step_min	Minimum required bid increase as a fraction of the previous bid
α	Minimum bid multiplier: $\alpha = 1 + \text{step_min}$
Δ	Bid multiplier at which g reaches inc_max: $\Delta = \alpha \times ((\text{inc_max} - \text{inc_min})/a + 1)$
R	Seller revenue (winning bid minus total incentives paid)
$R_{\%}$	Revenue percentage: R/b_K
BNE	Bayesian Nash Equilibrium
EU	Expected Utility
$T(v)$	BNE stopping threshold: bidder with value v bids up to price $T(v)$
S	Total incentive payout: $S = \sum_{t=1}^{K-1} g(b_{t-1}, b_t) \times b_t$
$v_{(k:n)}$	The k -th order statistic of n independent draws from F (i.e., the k -th smallest value)
ρ	Incentive ratio: $\rho = \text{inc_min}/\text{step_min}$ (= 0.10 across all canonical presets)

Symbol	Definition
v^*	BNE singularity: the unique v satisfying $F(v^*)^{n-1} = \text{inc_min}/(1 + \text{inc_min})$. For $F = U[0, 1]$, $v^* = (\text{inc_min}/(1 + \text{inc_min}))^{1/(n-1)}$. For $v > v^*$, $T(v)$ is finite and strictly increasing; for $v \leq v^*$, $T(v)$ diverges.

2 The GBM Auction

2.1 Auction Format

A GBM auction is a time-limited, incentivised, open ascending auction. It shares the basic structure of an English auction — bids go up, highest bidder at the end wins — but with critical differences in what happens when a bidder is outbid.

Players. A seller offers one item to $n \geq 2$ potential bidders. We work in the symmetric independent private values (IPV) framework: bidders draw values i.i.d. from a common distribution F and maximise expected payoff. The full model assumptions, formal game definition, and distributional robustness discussion are collected in Section 2.8.

Strategy space. At each point in the ascending auction, a bidder observes the current standing bid x and decides: (a) whether to bid, and (b) if so, how much to bid. A valid bid y must satisfy $y \geq \alpha \times x$ where $\alpha = 1 + \text{step_min}$. Bidding strategies are thus characterised by a stopping threshold $T(v)$ (the price at which a bidder with value v exits) and a bid-size rule $y(x, v)$ (the amount bid given standing bid x and value v).

Fully funded bids. All bids must be fully funded at the time of placement. The bidder commits the full bid amount (locked in escrow via smart contract on-chain, or via payment authorisation in fiat). This ensures the solvency of the auction at all times.

Minimum bid step. Each new bid must exceed the current highest bid by a minimum bid step step_min , calculated as a percentage of the current highest bid. Formally: $y \geq \alpha \times x$ where $\alpha = 1 + \text{step_min}$.

2.2 The Incentive Function

Definition. The incentive function $g : \mathbb{R}_+^2 \rightarrow [0, \text{inc_max}]$ takes the previous standing bid x and the new bid y , and returns the incentive that the new bidder will receive when outbid, expressed as a fraction of their own bid y :

$$g(x, y) = \min\left[\text{inc_max}, \text{inc_min} + a \cdot \frac{y - \alpha x}{\alpha x}\right] \quad \text{for } x > 0$$

$$g(0, y) = \text{inc_max}$$

Parameters:

Parameter	Symbol	Type	Description
Maximum incentive	inc_max	$\in (0, 1)$	Highest incentive rate an outbid bidder can receive
Minimum incentive	inc_min	$\in (0, \text{inc_max})$	Lowest incentive rate — earned at minimum step
Accelerator	a	$\in \mathbb{R}_+$	Controls how fast g ramps from inc_min to inc_max
Minimum bid step	step_min	$\in (0, 1)$	Minimum required increase as a fraction of previous bid

Properties of g :

- (i) g is a piecewise linear function of the bid ratio $r = y/x$. At $r = \alpha$ (minimum step), $g = \text{inc_min}$. As r increases, g increases linearly until reaching inc_max at $r = \Delta$, where it is capped.
- (ii) The incentive amount received when outbid is $g(x, y) \times y$. Both the rate g and the base y increase with bid size, so the absolute incentive is convex in the bid.
- (iii) The incentive is determined entirely at bid time: it depends only on x (known) and y (chosen by the bidder). It does not depend on any subsequent bid. This makes the payoff fully transparent before commitment.
- (iv) For the first bid ($x = 0$), the bidder receives the maximum incentive rate inc_max if outbid.

2.3 Payoffs

When the auction ends after K bids with winning bid b_K :

- **Winner** (the bidder who placed b_K): pays b_K , receives the item. Payoff: $v_{\text{winner}} - b_K$ (plus any incentives earned in earlier rounds when they were outbid).
- **Outbid bidder at position t** ($t < K$): receives $b_t + g(b_{t-1}, b_t) \times b_t$ (full refund plus incentive). Payoff: $g(b_{t-1}, b_t) \times b_t > 0$.
- **Seller**: receives $b_K - S$, where $S = \sum_{t=1}^{K-1} g(b_{t-1}, b_t) \times b_t$ is the total incentive payout.

2.4 The Solvency Constraint

Definition. The solvency constraint requires $\text{step_min} \geq \text{inc_max}$.

Proposition 2.1 (Solvency). Under the solvency constraint, the seller’s accumulated pot (current standing bid minus all incentives already paid) is non-decreasing in the number of bids.

Proof. When a new bid $b_{t+1} \geq \alpha \times b_t$ displaces b_t , the pot changes by $(b_{t+1} - b_t) - g(b_{t-1}, b_t) \times b_t \geq \text{step_min} \times b_t - \text{inc_max} \times b_t \geq 0$. ■

In plain language (*practitioners’ box*)¹

The solvency constraint is a simple rule: the minimum bid increase must be at least as large as the maximum possible incentive rate. This ensures the auction always has enough money to pay its obligations — no matter how many bids are placed, the pot never shrinks. It also prevents a key form of manipulation: if someone tries to outbid themselves to farm incentive payments, they always commit more new capital than they earn back (Section 7.1).

2.5 Derived Constants

Two derived constants appear throughout the analysis:

- $\alpha = 1 + \text{step_min}$ — the minimum bid multiplier.
- $\Delta = \alpha \times ((\text{inc_max} - \text{inc_min})/a + 1)$ — the bid multiplier at which g reaches inc_max . When $y \geq \Delta \times x$, the bidder earns the maximum incentive rate.

2.6 Worked Example

Consider parameters $\text{inc_max} = 0.10$, $\text{inc_min} = 0.01$, $a = 0.11$, $\text{step_min} = 0.10$ (the HIGH preset), and a start price of 0.

1. **Bidder A** bids $b_1 = 100$. Previous bid is 0, so $g(0, 100) = \text{inc_max} = 0.10$. If outbid, A will receive $\$100 + \$10 = \$110$.
2. **Bidder B** bids $b_2 = 115$ (a 15% jump). Bidder A receives $\$100.00 + \$10.00 = \$110.00$ (refund + incentive). B’s own incentive rate: $g(100, 115) = 0.01 + 0.11 \times (115 - \alpha \times 100)/(\alpha \times 100) = 0.01 + 0.11 \times 5/110 = 0.01 + 0.005 = 0.015$ (where $\alpha \times 100 = 110$ is the minimum valid bid). If outbid, B will receive $115.00 + 0.015 \times 115.00 = 115.00 + 1.725 = \116.725 , rounded to $\$116.73$.
3. **Bidder C** bids $b_3 = 230$ (a 100% jump, reaching $\Delta = 2.0$). Bidder B receives $\$115.00 + \$1.73 = \$116.73$. C’s incentive rate: $g(115, 230) = \text{inc_max} = 0.10$ (since $230/115 = 2.0 = \Delta$). If outbid, C will receive $\$230.00 + \$23.00 = \$253.00$.
4. **Auction ends.** Bidder C wins at $\$230.00$. Seller receives $\$230.00 - \$10.00 - \$1.73 = \218.28 (94.9% of the winning bid). (*Note: B’s exact incentive is $0.015 \times$*

¹Throughout this paper, practitioners’ boxes provide non-technical summaries of key results for applied readers. They are not part of the formal analysis.

115 = 1.725, rounded to \$1.73. All production implementations use round-half-up to the smallest currency unit; the theoretical analysis treats incentives as exact real numbers.)




Figure 1: **Timeline view** of the worked example (Section 2.6). Blue arrows show bid progression; red arrows show incentive payouts to outbid participants. The seller’s pot (top annotations) is non-decreasing — the solvency property in action.

2.7 Production Parameter Presets

Canonical preset (HIGH). Throughout this paper, we use the HIGH preset as the primary example. HIGH is the canonical default — the most commonly deployed preset and the richest analytically (it exhibits both min-step and value-dependent bid-size equilibria). Results are stated for HIGH unless noted otherwise.

Full preset family. Four presets span the parameter space:

Parameter	LOW (0.5%–2%)	MEDIUM (0.5%–5%)	HIGH (1%–10%)	DEGEN (2%–20%)
inc_min	0.005	0.005	0.01	0.02
inc_max	0.02	0.05	0.10	0.20
a	0.0162	0.05	0.11	0.27
step_min	0.05	0.05	0.10	0.20
α	1.050	1.050	1.100	1.200
Δ	2.022	1.995	2.000	2.000

Key observations:

- **Solvency constraint satisfied in all presets:** $\text{step_min} \geq \text{inc_max}$ holds for all. MEDIUM, HIGH, and DEGEN operate exactly at the solvency boundary ($\text{step_min} = \text{inc_max}$). LOW has headroom.
- $\Delta \approx 2$ **across all presets** — a bidder must roughly double the previous bid to earn the maximum incentive. This is a deliberate design decision by the authors.

- $\rho = \text{inc_min}/\text{step_min} = 0.10$ in all presets — the equilibrium incentive cost under the minimum-step BNE is exactly 10% of the price range (Section 5.1). This ratio is the fundamental design parameter. Since this ratio is the same across all four canonical presets, most theoretical results (equilibrium threshold, revenue share, welfare decomposition) are preset-independent. Where preset-specific differences matter (revenue floors, bid-size equilibria), all four presets are compared.

Important clarification: the universality of the 10% figure across presets is a deliberate design choice by the mechanism’s creators, not an emergent mathematical property. The theoretical results generalise straightforwardly to any ratio $\rho = \text{inc_min}/\text{step_min} \in (0, 1)$. *Caveat for HIGH/DEGEN:* The 10% figure applies precisely under the minimum-step BNE, which is the equilibrium for LOW and MEDIUM. For HIGH and DEGEN, where the equilibrium involves larger jumps for low-value bidders (Section 4.4), the actual equilibrium incentive cost differs because larger jumps earn incentive rates above inc_min ; see Section 5.3 for the actual revenue floors across presets.

2.8 Model Assumptions and Scope

Model assumptions. Throughout this paper, unless otherwise stated, we work in the symmetric independent private values (IPV) framework with time-limited auctions: $n \geq 2$ risk-neutral bidders draw private values $v_1, \dots, v_n$ independently from a common distribution F with support $[0, \bar{v}]$, positive density f , and default $F = U[0, 1]$. Each bidder has exogenous wealth $w_i \geq v_i$ (private information). The assumption $w_i \geq v_i$ excludes liquidity-constrained bidders. Relaxing this to $w_i < v_i$ is an open direction; the cross-auction wealth effect (Proposition 6.6) suggests that GBM may particularly benefit liquidity-constrained participants. The auction has a round limit $K_{\max}$ and an anti-sniping rule guaranteeing that every active bidder can respond to every bid. At expiry, if multiple bidders remain active, an efficient tiebreaking rule (e.g., a sealed-bid second-price auction among remaining active bidders) determines the winner. Under these assumptions, the auction achieves efficient allocation (Proposition 6.1) and terminates in bounded time. Bidders maximise expected payoff. The seller has no private information and no reserve price (extensions to reserve prices are discussed in Appendix A.1). Risk aversion is introduced in Section 8. Common-value and affiliated-value settings are not analysed (see Appendix B).

When is IPV appropriate? The IPV assumption — that each bidder’s value is private and independent of others’ — is standard in auction theory and is a natural first framework for a new mechanism. It is well-suited to settings where the item has subjective personal value: collectibles, art by emerging artists, unique digital assets (NFTs), and procurement where each bidder’s cost structure is proprietary. IPV is less appropriate for items with significant resale or common-value components (e.g., spectrum licenses, mineral rights, established-artist art with liquid secondary markets). The structural results of this paper — solvency, manipulation resistance, the revenue floor, and the BNE characterisation — depend only on the payment rule and hold regardless of the value model. The welfare and entry results require IPV through the efficient allocation assumption. Extending the analysis to affiliated or common values is an important open problem (Appendix B, Problem 6); the interaction be-

tween the positive lose-payoff and the winner’s curse is non-obvious and may either mitigate or exacerbate information-revelation effects.

Formal game definition, distributional robustness, and stationarity proofs are collected in Supplementary Materials, Section S1.

3 Relationship to Classical Auctions

3.1 GBM as a Generalisation of the English Auction

GBM reduces to English as $\text{inc_min} \rightarrow 0$. The equilibrium threshold converges to truthful bidding.

3.2 No Dominant Strategy

Theorem 3.2 (No Dominant Strategy). Unlike the English auction, the GBM ascending auction has no weakly dominant bidding strategy.

Proof. In the English auction, bidding up to v is weakly dominant because losing earns zero. In GBM, consider a bidder at price p :

- If they bid and **win**: payoff = $v - \alpha p$
- If they bid and **lose** (outbid): payoff = $\text{inc_min} \times \alpha p > 0$

At $q = 0$ (certain to lose at the next level), continuing is strictly profitable at every price. At $q = 1$ (certain to win), optimal stopping is v/α to avoid negative surplus. Since the optimal threshold depends on $q = P(\text{win})$, which depends on opponents’ strategies, no single threshold is weakly dominant. ■

In plain language (*practitioners’ box*)

In a traditional English auction, the best strategy is obvious: keep bidding until the price reaches what the item is worth to you, then stop. In GBM, the decision is harder because losing *pays* you. Even if the price has passed what the item is worth to you, placing one more bid might still make sense — if someone outbids you, you earn a reward. How far past your value you should go depends on how likely you are to be outbid, which depends on who else is bidding. There is no simple rule that works in every situation.

3.3 Revenue Equivalence Theorem Does Not Apply

Observation 3.3 (RET Violation). The Revenue Equivalence Theorem does not hold for GBM. A bidder with value near 0 can place a bid and collect a positive expected incentive when outbid, so $U(v_{\min}) > 0$, violating Myerson’s boundary condition.

3.4 Scale Invariance

Proposition 3.4 (Scale Invariance). The revenue percentage $R_{\%} = R/b_K$ is invariant to uniform scaling of all bids, since g depends only on the ratio $y/(\alpha x)$.

4 Equilibrium Analysis

4.1 The Ascending Auction BNE

Remark (why not sealed-bid). The sealed-bid formulation fails for GBM because it uses `inc_max` (the one-shot lose-payoff) rather than `inc_min` (the marginal per-step incentive), producing a singularity affecting 55% of bidders at $n = 5$. The ascending model is the correct framework.

Theorem 4.1 (BNE Stopping Threshold). In the symmetric n -bidder ascending GBM auction with values i.i.d. from any continuous distribution F with positive density on $[0, \bar{v}]$, the symmetric BNE stopping threshold is:

Keep bidding whenever `current_bid` $< T(v)$. Stop otherwise.

where:

$$T(v) = \frac{v \cdot F(v)^{n-1}}{F(v)^{n-1}(1 + \text{inc_min}) - \text{inc_min}}$$

For $F = U[0, 1]$ (where $F(v) = v$), this reduces to $T(v) = v^n / [v^{n-1}(1 + \text{inc_min}) - \text{inc_min}]$.

Proof. At the BNE threshold, the bidder is indifferent between one more minimum-step bid and stopping. Write $q = F(v)^{n-1}$ for the win probability (exact in the monotone region, where T is strictly increasing and the exit order perfectly reflects value order). The indifference condition:

$$q \times (v - T) + (1 - q) \times \text{inc_min} \times T = 0$$

The first term is the expected surplus from winning ($v - T < 0$ at the threshold). The second is the expected incentive from losing. Solving:

$$T(v) = \frac{v \cdot F(v)^{n-1}}{F(v)^{n-1}(1 + \text{inc_min}) - \text{inc_min}}$$

Optimality. Define $\pi(p, v) = F(v)^{n-1}(v - p) + (1 - F(v)^{n-1}) \cdot \text{inc_min} \cdot p$. For $v > v^*$: $\partial\pi/\partial p < 0$, so π is strictly decreasing. Since $\pi(T(v), v) = 0$, we have $\pi > 0$ for $p < T(v)$ and $\pi < 0$ for $p > T(v)$. $T(v)$ is the unique optimal stopping point.

Stationarity. Past incentive earnings are sunk: they cancel in the continue-vs-stop comparison. The marginal payoff of one more bid depends only on the current price and $F(v)^{n-1}$, making the threshold history-independent. ■

In plain language (*practitioners' box*)

The formula $T(v)$ answers: “If an item is worth v to me, at what price should I stop bidding?” In a normal English auction, the answer is exactly v . In GBM, the answer is *slightly* above v , because even if the price has passed what the item is worth, there is a small chance you get outbid and earn a reward. The formula balances two forces: the risk of overpaying if you win versus the reward you collect if you lose. For the highest-value bidders who determine the final price, the overbidding is less than 0.1%.

4.2 Bounds on Overbidding

Overbidding exists but is negligibly small for the highest-value bidders who determine the final price. The mechanism does not inflate prices; it increases *participation*, and more competition is what drives prices up.

Proposition 4.2 (Lower Bound). Every bidder with $v > 0$ should bid at least up to their value: $T(v) \geq v$ for all $v > 0$.

Proof. From the closed form, $T(v)/v = F(v)^{n-1}/[F(v)^{n-1}(1 + \text{inc_min}) - \text{inc_min}]$. Writing $q = F(v)^{n-1}$, this ratio exceeds 1 iff $q > q(1 + \text{inc_min}) - \text{inc_min}$, i.e., iff $\text{inc_min}(1 - q) > 0$, which holds for all $v < \bar{v}$ since $q = F(v)^{n-1} < 1$. At $v = \bar{v}$: $q = 1$ and $T(\bar{v}) = \bar{v}$. Hence $T(v) \geq v$ for all $v > 0$. ■

4.3 Overbidding Bound

Theorem 4.3 (Overbidding Bound). For any continuous F , any $v > v^*$, $n \geq 2$:

$$\frac{T(v)}{v} - 1 = \frac{\text{inc_min}(1 - F(v)^{n-1})}{F(v)^{n-1}(1 + \text{inc_min}) - \text{inc_min}}$$

In particular, $T(\bar{v}) = \bar{v}$ (zero overbidding at the boundary). For v near $\bar{v}$: $T(v)/v - 1 \sim \text{inc_min} \cdot (n - 1)(1 - F(v))$.

Proof. Direct algebra from the BNE threshold. ■

Interpretation. Overbidding is $O(\text{inc_min})$ — controlled entirely by the minimum incentive rate. For $F = U[0, 1]$, HIGH preset ($\text{inc_min} = 0.01$): overbidding at $v = 0.99$ is 0.01% for $n = 2$ and 0.04% for $n = 5$.

4.4 Bid-Size Strategy

Beyond the stopping threshold, bidders must decide how much to bid each round — the minimum step $\alpha \times x$, or a larger jump toward $\Delta \times x$. The optimal bid size involves a trade-off between a *price effect* (small steps preserve surplus) and an *incentive-locking effect* (big jumps earn higher incentive rates).

Proposition 4.4 (Minimum-Step BNE for LOW/MEDIUM). In GBM auctions with parameters satisfying $a \times \text{step_min} \leq \text{inc_min}$ (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 sketch. Two effects both favour small steps: (1) the incentive-rate effect — smaller jumps extract a higher incentive rate per unit of price increase (since $h(\alpha) \geq h(r)$ for all $r \geq \alpha$ when $a \cdot \text{step_min} \leq \text{inc_min}$); (2) the price-preservation effect — smaller steps preserve more bidder surplus by raising the price less. Both point to minimum step. Full proof in Supplementary Materials, Section S5. ■

4.5 The Complete Equilibrium

The equilibrium has two components:

1. **Stopping rule:** keep bidding while $\text{current_bid} < T(v)$, where $T(v) = v \cdot F(v)^{n-1} / [F(v)^{n-1}(1 + \text{inc_min}) - \text{inc_min}]$.
2. **Bid size:** minimum step ($\alpha \times x$) in LOW/MEDIUM presets for all bidders; value-dependent in HIGH/DEGEN (low-value bidders bid larger, high-value bidders bid minimum). The crossover threshold and its proof are in Supplementary Materials, Section S5.

In plain language (*practitioners' box*)

For most configurations, the optimal strategy is simple — bid the minimum step and keep going slightly past your value. For aggressive presets (HIGH/DEGEN), bidders who don't expect to win should consider bigger jumps to lock in larger rewards.

4.6 Uniqueness

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, stated and proved in Supplementary Materials, Section S5). In the continuous-price limit, the bid-size is uniquely optimal (Regimes I, III) or payoff-irrelevant (Regime II) at every state.

Proof sketch. Stopping threshold uniqueness follows from strict monotonicity of the marginal payoff (shown in Theorem 4.1). Bid-size uniqueness uses the incentive cost density $\lambda(r) = g(r)/(r-1)$: when $\lambda'(r)$ has uniform sign, the optimum is at a boundary. Full proof in Supplementary Materials, Section S4. ■

In plain language (*practitioners' box*)

There is only one rational way to play. The “when to stop” rule is the same no matter what — and the “how much to jump” rule is pinned down by the preset parameters. You don’t need to guess what other bidders are doing to figure out the right strategy.

Note: Proposition 4.6 (λ -classification of bid-size regimes) is stated and proved in Supplementary Materials, Section S5, which is why the numbering proceeds directly to Theorem 4.7.

4.7 Equilibrium of the Timed Model

The BNE threshold $T(v)$ (Theorem 4.1) was derived from a marginal indifference condition using the unconditional win probability $F(v)^{n-1}$. For bidders in the monotone region ($v > v^*$), this is exact: the strictly increasing threshold provides an injective mapping from values to dropout prices. For bidders in the non-monotone region ($v \leq v^*$), the threshold diverges ($T(v) \rightarrow \infty$), raising the question of whether the timed model (Section 2.8) with its round limit $K_{\max}$ and efficient tiebreaker yields a well-defined equilibrium. This section establishes that it does.

Theorem 4.7 (Equilibrium of the Timed Game). Consider the timed GBM ascending auction Γ^T (Section 2.8) with $n \geq 2$ bidders, values i.i.d. from any continuous F with positive density on $[0, \bar{v}]$, wealth $w_i \geq v_i$, and round limit $K_{\max} \geq \lceil \log(\bar{v}/b_0) / \log \alpha \rceil$. Then:

- (i) *Existence:* The timed game Γ^T possesses a symmetric Bayesian Nash equilibrium.
- (ii) *Monotone-region characterisation:* In any symmetric BNE, every bidder with $v > v^*$ uses the threshold $T(v)$ from Theorem 4.1 and exits during the ascending phase before round $K_{\max}$. The win probability $F(v)^{n-1}$ is exact for all $v > v^*$.
- (iii) *Non-monotone-region behaviour:* Every bidder with $v \leq v^*$ bids in every ascending-phase round in which they can afford the minimum bid. If still active at expiry, they bid truthfully ($b = v$) in the tiebreaker.
- (iv) *Efficiency:* The allocation is efficient with probability at least $1 - n \cdot F(v^*)^{n-1}$. In particular, whenever $v_{(n:n)} > v^*$ and $v_{(n-1:n)} > v^*$ — i.e., both price-determining bidders are in the monotone region — the ascending phase allocates efficiently without the tiebreaker.

Proof. In the monotone region, T is strictly increasing, so exit order equals value order — the highest-value bidder wins without a tiebreaker. The probability that at most one bidder has $v > v^*$ — i.e., that $v_{(n-1:n)} \leq v^*$ — is $F(v^*)^{n-1}[n - (n-1)F(v^*)] \leq n \cdot F(v^*)^{n-1}$ by the order-statistic CDF bound. For canonical parameters: < 0.02 at $n = 2$ (HIGH), < 0.05 at $n = 5$. ■

In plain language (*practitioners' box*)

About 90% of bidders (those whose item value is above a small cutoff v^*) behave predictably: they drop out at a price determined by their value, the highest-value bidder wins, and the auction is efficient. The remaining $\sim 10\%$ are low-value bidders who keep bidding to collect incentives — the time limit and tiebreaker ensure they can't disrupt the outcome. The auction works correctly in practice with probability above 95%.

5 Revenue Analysis

5.1 Equilibrium Revenue

Theorem 5.1 (Equilibrium Incentive Cost). In a GBM ascending auction where all bidders use minimum-step bidding, the total incentive payout satisfies:

$$\text{Total incentives} = \frac{\text{inc_min}}{\text{step_min}} \times (\text{winning_price} - \text{start_price})$$

For $\text{winning_price} \gg \text{start_price}$:

$$\text{Seller revenue} \approx \left(1 - \frac{\text{inc_min}}{\text{step_min}}\right) \times \text{winning_price}$$

Proof. Under minimum-step bidding ($b_{t+1} = \alpha \cdot b_t$) with $b_0 > 0$, each outbid at level b_t ($t = 1, \dots, K-1$) earns incentive $\text{inc_min} \times b_t$. Summing: $S = \text{inc_min} \sum_{t=1}^{K-1} b_t = \text{inc_min} \cdot b_0 \cdot \alpha(\alpha^{K-1} - 1)/(\alpha - 1) = (\text{inc_min}/\text{step_min})(b_K - \alpha b_0)$. When $b_0 = 0$, the first outbid earns $\text{inc_max} \cdot b_1$ rather than $\text{inc_min} \cdot b_0 = 0$; this adds $O(b_1)$ to the sum, which is negligible relative to b_K for $K \gg 1$. In either case, $S \approx \rho \cdot b_K$ for $b_K \gg b_0$. ■

In plain language (*practitioners' box*)

Of every dollar the winning bidder pays, a fraction $\rho = \text{inc_min}/\text{step_min}$ goes to losing bidders as incentive payments and the rest goes to the seller. For the canonical presets in this paper, $\rho = 0.10$, giving a 90/10 split — but this is an author's design choice, not a fundamental constant. Different presets can adjust ρ to shift the balance between seller revenue and bidder incentives. The split is independent of the item's value, the number of bidders, or the final price.

5.2 Revenue Floor

The equilibrium result assumes min-step bidding. An adversary who controls the bid sequence can potentially extract more. This section derives the worst-case seller revenue over all possible bid sequences.

Theorem 5.2 (Revenue Floor). For any GBM auction satisfying the solvency constraint, the worst-case seller revenue percentage over all bid sequences is:

$$R_{\%,\min} = 1 - \frac{\max_{r \geq 1 + \text{step_min}} h(r)}{1 + \text{step_min}}$$

where the *incentive extraction function* is:

$$h(r) := \frac{g(r) \cdot r}{r - 1}$$

where $g(r) = \min[\text{inc_max}, \text{inc_min} + a(r - \alpha)/\alpha]$ is the incentive rate at bid ratio r and $\alpha = 1 + \text{step_min}$. The maximum is achieved at a boundary point:

$$R_{\%,\min} = 1 - \max\left(\frac{\text{inc_min}}{\text{step_min}}, \frac{\text{inc_max} \cdot \Delta}{(1 + \text{step_min})(\Delta - 1)}\right)$$

The first term binds when $h(\alpha) \geq h(\Delta)$ (LOW and MEDIUM presets: $R_{\%,\min} = 1 - \text{inc_min}/\text{step_min} = 90\%$). The second binds when $h(\Delta) > h(\alpha)$ (HIGH and DEGEN presets: $R_{\%,\min} = 81.8\%$ and 66.7% respectively). The bound is asymptotically tight.

Proof sketch. Define the normalised incentive liability after N bids as $W_N = S_N/b_N$ (total incentives / current bid). An inductive argument shows $W_N < \max(h(\alpha), h(\Delta))/\alpha$ for all N , since each bid's contribution is bounded by $h(r_t)/(r_t)$ and the geometric growth of bids dilutes earlier contributions. Tightness: a sequence of all- Δ bids followed by a minimum-step winner approaches the bound geometrically at rate Δ^{-1} . Full proof in Supplementary Materials, Section S7. ■

In plain language (*practitioners' box*)

The seller is guaranteed to keep at least 90% of the winning price for the LOW and MEDIUM presets, and at least 82% for HIGH — no matter how bidders choose their bid sizes. This is the worst case; in equilibrium (where bidders use minimum steps), the seller keeps exactly $1 - \rho$. The revenue floor exists because the solvency constraint ensures that big bids — which earn higher incentive rates — also lock up proportionally more capital, capping the total payout.

5.3 Revenue Floors Across Presets

Preset	$R_{\%,\min}$	Binding pattern	Equilibrium cost	Worst-case cost
LOW	90.00%	All- α	10.0%	10.0%
MEDIUM	90.00%	All- α	10.0%	10.0%
HIGH	81.82%	All- $\Delta + \alpha$ winner	10.0%	18.2%
DEGEN	66.67%	All- $\Delta + \alpha$ winner	10.0%	33.3%

5.4 The Divergent Expectation (Untimed Theoretical Model)

In an untimed theoretical model with no wealth constraints, $E[T(v_{(n-1:n)})]$ diverges. The source is the singularity at $v^* = (\text{inc_min}/(1 + \text{inc_min}))^{1/(n-1)}$: as $v \rightarrow v^*$ from above, $T(v) \rightarrow \infty$, so realizations where both top bidders have values near v^* produce arbitrarily large winning prices. Since $f_{(n-1:n)}(v^*) > 0$, these extreme realizations have positive probability mass, causing the expectation to diverge.

This pathology does not arise in the timed model (the framework of this paper). The round limit $K_{\max}$ caps the maximum price at $\alpha^{K_{\max}} \cdot b_0$, ensuring that all moments of the winning price are finite. The divergent expectation is an artifact of the untimed abstraction, not a property of any deployed GBM auction. This regularity failure means that envelope-theorem arguments relying on finite expected payments cannot be applied to the untimed model; the timed framework (which restores finite moments) is therefore essential, not merely convenient. A sensitivity analysis showing that results are robust to the specific choice of $K_{\max}$ is in Supplementary Materials, Section S7.

Remark (Monte Carlo convergence and tail behaviour). The timed model’s price bound ensures that Monte Carlo estimates of expected revenue converge at the standard $O(1/\sqrt{N})$ rate. However, the heavy tail inherited from the near- v^* singularity elevates the variance of the winning price, making mean-based estimators slow to converge in practice. This is why all Monte Carlo revenue statistics in this paper (Section 5.5) report medians rather than means, with quantile ranges (P75, P95) to characterise the tail. Bootstrap confidence intervals on 100,000 simulations achieve sub-percentage-point precision on all reported statistics. The distinction between the timed model (finite moments, well-behaved estimation) and the untimed abstraction (divergent moments, unstable estimation) is thus practically significant, not merely theoretical — it is the round limit that makes the mechanism’s properties empirically measurable.

In plain language (*practitioners’ box*)

In a hypothetical auction with no time limit, the theoretical “average” winning price diverges — a consequence of the singularity in the bidding threshold. In every real GBM auction, the time limit resolves this. Low-value bidders in the non-monotone region ($v \leq v^*$, roughly 10% of the population at $n = 2$) bid aggressively to collect incentives but remain individually rational at every decision point: they collect positive payoffs whenever outbid (Theorem 6.3a) and exit via the efficient tiebreaker when the round limit binds. When simulating GBM auctions, use medians rather than means — the heavy tail makes means converge slowly.

5.5 Comparison with English Auction Revenue

Proposition 5.3 (Revenue Ratio under Efficient Allocation). Under efficient allocation (Proposition 6.1) and minimum-step equilibrium bidding, the GBM seller revenue satisfies:

$$R_{\text{GBM}} = (1 - \rho) \cdot p_{\text{win}}$$

where p_{win} is the winning price. In the continuous-price limit, $p_{\text{win}} = v_{(n-1:n)}$ and $R_{\text{GBM}} = (1 - \rho) \cdot R_{\text{Eng}}$. In the discrete model with overbidding, $p_{\text{win}} = T(v_{(n-1:n)}) \geq v_{(n-1:n)}$, so $R_{\text{GBM}} \geq (1 - \rho) \cdot R_{\text{Eng}}$ — the overbidding premium partially recovers the incentive cost.

Proof. Follows directly from Theorem 5.1: seller revenue $= p_{\text{win}} - \rho \cdot p_{\text{win}} = (1 - \rho) \cdot p_{\text{win}}$. Under efficient allocation, $p_{\text{win}} = T(v_{(n-1:n)}) \rightarrow v_{(n-1:n)} = R_{\text{Eng}}$ in the continuous limit. ■

6 Welfare Analysis and Mechanism Comparison

The welfare analysis of GBM centres on a three-way surplus decomposition. Under efficient allocation, total social surplus is $W = v_{(n:n)}$ in both GBM and English auctions — the pie is the same size. What differs is how it is split among winner, losers, and seller. In an English auction, the winner captures $v_{(n:n)} - v_{(n-1:n)}$, losers receive nothing, and the seller receives $v_{(n-1:n)}$. In GBM, the seller transfers $\rho \cdot v_{(n-1:n)}$ to the losers, retaining $(1 - \rho) \cdot v_{(n-1:n)}$. The bidder advantage (Theorem 6.2) is funded by this seller revenue reduction — it is a redistribution, not a Pareto improvement. The mechanism's value proposition to the seller is not that it improves per-auction revenue with the same bidders, but that the redistribution attracts additional participants whose competitive pressure more than compensates for the ρ cost (Section 6.5).

6.1 Efficient Allocation

Analytical framework. The welfare comparison between GBM and English auctions requires efficient allocation — both formats must assign the item to the highest-value bidder. In the English auction, efficient allocation is automatic (bidding up to one's value is dominant). In GBM, efficient allocation is guaranteed by the timed model with efficient tiebreaking at expiry (Section 2.8).

Proposition 6.1 (Efficient Allocation). The timed GBM auction (Section 2.8) achieves efficient allocation for all $n \geq 2$, all continuous F with positive density on $[0, \bar{v}]$, all wealth distributions with $w_i \geq v_i$, and all $K_{\text{max}} \geq \log(\bar{v}/b_0)/\log \alpha$.

Proof. Follows from Theorem 4.7(iv): in the monotone region ($v > v^*$), the strictly increasing threshold ensures the highest-value bidder is the last to exit. In the non-monotone region, the efficient tiebreaker allocates correctly. ■

6.2 Bidder Preference (General n)

The welfare result. Given efficient allocation (Proposition 6.1), we can now state the main welfare theorem. As noted at the opening of this section, the bidder advantage documented below is funded by a seller revenue reduction of $\rho \cdot E[v_{(n-1:n)}]$;

the mechanism is self-financing when this redistribution induces sufficient additional entry (Section 6.5).

Theorem 6.2 (Surplus Redistribution under Efficient Allocation — General).

Consider a GBM auction with $n \geq 2$ bidders, values i.i.d. from any continuous distribution F with positive density on $[0, \bar{v}]$, minimum-step equilibrium bidding, and parameters satisfying the solvency constraint. If the mechanism achieves efficient allocation, then the GBM auction redistributes a fraction ρ of expected seller revenue to losing bidders:

(i) *Bidder advantage:* Every bidder's ex ante expected utility is strictly higher than in the corresponding English auction:

$$\text{EU}_{\text{GBM}}(n) - \text{EU}_{\text{Eng}}(n) = \frac{\rho}{n} \cdot E_F[v_{(n-1:n)}] > 0$$

(ii) *Seller cost:* The seller's expected revenue is strictly lower:

$$R_{\text{GBM}}(n) = (1 - \rho) \cdot R_{\text{Eng}}(n) = (1 - \rho) \cdot E_F[v_{(n:n)}]$$

(iii) *Total welfare invariance:* Total social surplus is $E_F[v_{(n:n)}]$ in both formats — the bidder advantage is funded entirely by the seller revenue reduction.

where $\rho = \text{inc_min}/\text{step_min}$ and $v_{(n-1:n)}$ is the second-highest order statistic. The bidder advantage factor is:

$$\frac{\text{EU}_{\text{GBM}}(n)}{\text{EU}_{\text{Eng}}(n)} = 1 + \rho \cdot \frac{E_F[v_{(n-1:n)}]}{E_F[v_{(n:n)} - v_{(n-1:n)}]}$$

For $F = U[0, 1]$, this simplifies to $1 + \rho(n - 1)$.

Proof. Under efficient allocation, total surplus is $v_{(n:n)}$ in both formats. In English: $\text{EU}_{\text{Eng}} = E[v_{(n:n)} - v_{(n-1:n)}]/n$ (winner surplus divided equally ex ante). In GBM: each loser receives $\rho \cdot v_{(n-1:n)}/(n - 1)$ from the incentive pool, so $\text{EU}_{\text{GBM}} = \text{EU}_{\text{Eng}} + \rho \cdot E[v_{(n-1:n)}]/n$. The seller loses exactly $\rho \cdot E[v_{(n-1:n)}]$. For $F = U[0, 1]$: $E[v_{(n-1:n)}] = (n - 1)/(n + 1)$, giving advantage factor $1 + \rho(n - 1)$. ■

6.3 Payoff Dominance (FOSD)

Theorem 6.2 establishes that the *expected* payoff is higher in GBM for the average bidder. The following results characterise the distributional comparison between the GBM and English payoffs, establishing when GBM dominance holds state by state, ex ante, and for general risk preferences.

Theorem 6.3 (Path-by-Path FOSD — Continuous-Price Limit). Consider the GBM ascending auction with $n \geq 2$ bidders, values i.i.d. from any continuous F with positive density on $[0, \bar{v}]$, minimum-step bidding, and efficient allocation (which

is exact in the continuous-price limit; see Theorem 4.7(iv) for the discrete approximation). In the continuous-price limit ($\text{step_min} \rightarrow 0$ with $\rho = \text{inc_min}/\text{step_min}$ held fixed), for every value $v \in (0, \bar{v})$ and every realisation ω of opponents' values:

$$\pi_{\text{GBM}}(v, \omega) \geq \pi_{\text{English}}(v, \omega)$$

with strict inequality on the losing event (probability $1 - F(v)^{n-1} > 0$). Consequently, for any strictly increasing utility function u :

$$E[u(\pi_{\text{GBM}}(v))] > E[u(\pi_{\text{English}}(v))]$$

Concavity of u is not required: the result holds for risk-neutral, risk-averse, and risk-loving bidders alike.

Proof. Fix v and opponent realization ω . **Losing state** ($v < \max_j v_j$): $\pi_{\text{GBM}}^{\text{lose}} = S_\ell > 0 = \pi_{\text{Eng}}^{\text{lose}}$ (the bidder collects incentive income). **Winning state** ($v > \max_j v_j$): in the continuous-price limit, the GBM winner pays $T(v_{(n-1:n)}) \rightarrow v_{(n-1:n)}$ (since overbidding vanishes), and collects incentive income from any earlier rounds where they were outbid. So $\pi_{\text{GBM}}^{\text{win}} \geq \pi_{\text{Eng}}^{\text{win}}$, with equality in the limit. The losing state has probability $1 - F(v)^{n-1} > 0$ for $v < \bar{v}$, giving strict inequality in expectation. ■

Theorem 6.3a (Losing-State Strict Dominance — Discrete Grid). On the discrete grid, for all $n \geq 2$, all continuous F , and all GBM parameters satisfying the solvency constraint: every losing bidder receives $S_\ell \geq \text{inc_min} \times \alpha b_0 > 0$, strictly dominating the English payoff of 0 in every state. This requires no limit arguments — it holds for the deployed mechanism as parameterised.

Theorem 6.3b (Full Discrete FOSD for $n = 2$). Under the one-grid-step condition $\text{step_min} \leq \rho/(1 - \rho)$ (which holds for LOW, MEDIUM, and HIGH), full FOSD holds for $n = 2$ on the discrete grid.

Proofs of both results are in Supplementary Materials, Section S6. For $n \geq 3$ on the discrete grid, the per-value comparison can fail for high-value bidders who bear the overbidding premium but receive only $\sim 1/n$ of the incentive pool — an open problem (Supplementary Materials, Section S11, Problem 5).

In plain language (*practitioners' box*)

Every bidder — winner or loser — weakly prefers a GBM auction over an English auction for the same item. Winners pay nearly the same price; losers walk away with money instead of nothing. This is not just an “on average” statement: it holds outcome by outcome, for every possible draw of opponents' values. Because the comparison is state-by-state (not just in expectation), it applies regardless of whether bidders are risk-neutral, risk-averse, or risk-loving.

6.4 Comparison with RET Benchmark and Myerson Optimal

The RET benchmark for n i.i.d. $U[0, 1]$ bidders is $E[\text{Revenue}] = (n-1)/(n+1)$. GBM seller revenue $\approx (1 - \rho) \times (n-1)/(n+1) = 90\%$ of benchmark, independent of n .

Myerson's optimal mechanism (second-price auction with reserve $r^* = 1/2$) earns more than the RET benchmark for small n :

n	RET benchmark	Myerson optimal	GBM (no reserve)	GBM/RET	GBM/Myerson
2	0.333	0.417	0.300	90.0%	72.0%
3	0.500	0.531	0.450	90.0%	84.7%
5	0.667	0.672	0.600	90.0%	89.3%
10	0.818	0.818	0.736	90.0%	90.0%

GBM with a reserve price. Adding $r^* = 0.5$ to GBM gives revenue $\approx 0.90 \times$ Myerson.

In plain language (*practitioners' box*)

GBM without a reserve earns 90% of the standard auction benchmark, regardless of how many bidders participate. GBM with a reserve price earns 90% of the theoretical maximum. For small auctions (2–3 bidders), the gap to the optimum is noticeable, but it shrinks rapidly with more bidders — and the participation benefit (Section 6.5) can more than close it.

6.5 Endogenous Entry

Proposition 6.5 (Fixed-Participation Revenue Inferiority). For any fixed $n \geq 2$, any continuous F , and any GBM parameters with $\rho > 0$: $R_{\text{GBM}}(n) = (1 - \rho) \cdot R_{\text{Eng}}(n) < R_{\text{Eng}}(n)$. A seller comparing the two formats with the same bidders would strictly prefer English.

The mechanism's value proposition to the seller is therefore *not* per-auction revenue with fixed participation — it is that the redistribution attracts additional participants whose competitive pressure more than compensates for the ρ cost. This section formalises this argument.

Model. We adopt the Levin and Smith (1994) entry framework: N potential bidders can enter by paying cost $c > 0$. Upon entering, they draw $v \sim F$. Equilibrium entry: $n^* = \text{largest } n \text{ satisfying } \text{EU}(n) \geq c$.

Theorem 6.4 (Entry Advantage — General). Under the Levin-Smith entry framework with values i.i.d. from any continuous F with positive density on $[0, \bar{v}]$ and efficient allocation:

- (i) *GBM attracts weakly more entrants:* $n_{\text{GBM}}^* \geq n_{\text{Eng}}^*$ for all entry costs $c > 0$.

- (ii) *Revenue dominance condition:* When $n_{\text{GBM}}^* = n_{\text{Eng}}^* + 1 = m + 1$, GBM seller revenue exceeds English seller revenue if and only if the marginal competition effect outweighs the incentive cost:

$$\frac{E_F[v_{(m:m+1)}]}{E_F[v_{(m-1:m)}]} > \frac{1}{1 - \rho}$$

That is, the proportional increase in expected second-highest value from one additional bidder must exceed $1/(1 - \rho)$.

- (iii) For $F = U[0, 1]$: The revenue dominance condition reduces to $m(m + 1) < 2/\rho$. For $\rho = 0.10$, GBM strictly revenue-dominates for $n_{\text{Eng}}^* \leq 3$, with an exact tie at $n_{\text{Eng}}^* = 4$.

Proof sketch. (i) Since $\text{EU}_{\text{GBM}}(n) > \text{EU}_{\text{Eng}}(n)$ for all $n \geq 2$, the entry threshold is weakly higher. (ii) Compare $(1 - \rho) \cdot E[v_{(m:m+1)}]$ (GBM with $m + 1$) vs $E[v_{(m-1:m)}]$ (English with m). (iii) For $U[0, 1]$: $E[v_{(m:m+1)}]/E[v_{(m-1:m)}] = m(m+1)/[(m-1)(m+2)]$, and the condition becomes $m(m + 1) < 2/\rho = 20$, giving $m \leq 3$. ■

Remark (Heterogeneous entry costs and non-uniform distributions). Theorem 6.5 (Supplementary Materials, Section S8) extends the entry result to heterogeneous entry costs drawn from any continuous G : the entry advantage becomes strict ($n_{\text{GBM}}^e > n_{\text{Eng}}^e$), avoiding the knife-edge sensitivity of the homogeneous-cost model. The revenue dominance threshold m^* ranges from 2 (concave F) to 8 (strongly convex F) across distribution families (Supplementary Materials, Section S8).

Remark (Optimal ρ). The seller's optimal incentive ratio solves $\max_{\rho}(1 - \rho) \cdot E[v_{(n^*(\rho)-1:n^*(\rho))}]$, where $n^*(\rho)$ is the equilibrium number of entrants. The optimal ρ depends on the entry cost distribution and the size of the potential bidder pool; $\rho = 0.10$ is not claimed to be optimal but provides robust performance across configurations (see Supplementary Materials, Section S12, Problem 3, for a numerical illustration).

In plain language (*practitioners' box*)

With the same number of bidders, a GBM seller earns less than an English seller — 90 cents on the dollar. But GBM attracts more bidders, because participating has a positive expected payoff even if you lose. More bidders means a higher final price. Whenever GBM attracts even one extra bidder, the price increase from extra competition more than compensates for the 10% incentive cost — up to about 4 bidders. This is the mechanism's core value proposition: it trades a small per-auction cost for a larger participation benefit.

6.6 Dynamic GBM and Cross-Auction Effects

The preceding analysis models a single auction. In practice, GBM auctions are deployed in *events* — collections of M items sold sequentially or simultaneously. This

creates a dynamic interaction absent from the single-auction model: incentive income earned in early auctions changes bidding behaviour in later ones.

Setup. Consider M items sold sequentially to a pool of N potential bidders. Each bidder i has a value $v_{im} \sim F$ for item m (drawn independently across items) and an initial wealth w_i . In auction m , a bidder who earned total incentive income $S_{i,<m} = \sum_{k<m} S_{ik}$ from prior auctions has effective budget $B_{im} = \min(v_{im}, w_i + S_{i,<m})$. Under DARA preferences (decreasing absolute risk aversion), the wealth effect from incentive income reduces risk aversion: $\lambda(w + S) < \lambda(w)$ for $S > 0$.

Proposition 6.6 (Cross-Auction Participation Multiplier). Under the timed model, DARA utility, and sequential auctions with independent values:

(i) A bidder who earns incentive income $S > 0$ in auction k has a weakly higher certainty equivalent for participating in auction $k + 1$ than an otherwise identical bidder with $S = 0$. The increase is strict for risk-averse bidders ($\lambda > 0$).

(ii) If the entry cost c is constant across auctions, the probability that a bidder who participated in auction k also enters auction $k + 1$ is weakly higher in GBM than in English, conditional on losing auction k .

Proof. Under DARA: $\lambda(w + S) < \lambda(w)$ for $S > 0$. The CE of a gamble increases when risk aversion decreases (standard result). A GBM loser has $S > 0$; an English loser has $S = 0$. Part (ii) follows from (i) and the entry threshold $CE \geq c$. Full details in Supplementary Materials, Section S8. ■

7 Manipulation Resistance

7.1 Sybil Attacks (Self-Outbidding)

Setup. An attacker controls multiple wallets. The current highest bid is b_0 (from a legitimate bidder). The attacker bids $b_1 \geq \alpha \cdot b_0$ from wallet A, then outbids themselves with $b_2 \geq \alpha \cdot b_1$ from wallet B. Wallet A receives incentive $g(b_0, b_1) \cdot b_1$. The attack is profitable if $g(b_0, b_1) \cdot b_1 > b_2 - b_1$.

Theorem 7.1 (Sybil Resistance). Under the solvency constraint ($\text{step_min} \geq \text{inc_max}$), for any valid self-outbid sequence, the additional capital locked weakly exceeds the incentive received:

$$b_2 - b_1 \geq g(b_0, b_1) \cdot b_1 \quad \text{for all valid } b_1, b_2$$

Proof. The attacker's net gain is $g(b_0, b_1) \cdot b_1 - (b_2 - b_1)$. Since $g \leq \text{inc_max} \leq \text{step_min}$ and $b_2 - b_1 \geq \text{step_min} \cdot b_1$, we have $g(b_0, b_1) \cdot b_1 \leq \text{step_min} \cdot b_1 \leq b_2 - b_1$ for all valid bids. ■

In plain language (*practitioners' box*)

A bidder using two wallets to outbid themselves — hoping to collect the incentive reward — always loses money. The capital they must lock up in the second (higher) bid exceeds the incentive they earn on the first. This holds for any bid size, from the minimum step to the maximum jump. The same logic applies to wash trading (two colluders bouncing bids back and forth): each round costs them more than it pays. The solvency constraint ($\text{step_min} \geq \text{inc_max}$) is what makes this work.

7.2 Shill Bidding

In a standard English auction, shill bidding is free: the seller bids on their own item and, if outbid, gets their money back with no penalty. In GBM, shill bidding costs the seller money.

Proposition 7.2 (Shill Bidding Cost). Under the solvency constraint with $\rho < 1/\alpha$, each shill cycle at standing bid b incurs a net cost of $(\text{step_min} - \text{inc_min} \cdot \alpha) \cdot b > 0$ to the seller.

Proof. The shill bids at αb , earning incentive $\text{inc_min} \cdot \alpha b$ on the displaced bid but committing $\alpha b - b = \text{step_min} \cdot b$ in new capital. Net cost: $\text{step_min} \cdot b - \text{inc_min} \cdot \alpha \cdot b = (\text{step_min} - \text{inc_min} \cdot \alpha) \cdot b > 0$ since $\rho < 1/\alpha$. ■

7.3 Wash Trading

Two colluding wallets bouncing bids back and forth is equivalent to repeated self-outbidding (Section 7.1). Each round earns $\text{inc_min} \times \text{bid_level}$ but costs $\text{step_min} \times \text{bid_level}$. The pair is literally paying the seller to trade.

7.4 Collusion and Bidder Rings

7.4.1 Classical Bid Suppression and Incentive Farming

In both English and GBM, a bidder ring can suppress competition by designating one member to bid while others abstain. The damage is identical in both formats: effective competition drops to $(n - k + 1)$ bidders. GBM's incentive structure creates no additional bid suppression vulnerability. Could a ring instead have members bid against each other to harvest incentives? This reduces to wash trading (Section 7.3) and is therefore unprofitable under the solvency constraint.

7.4.2 The Defection Premium

Theorem 7.3 (Collusion Resistance). For any GBM auction with valid parameters and $n \geq 3$ bidders: (i) the set of internally stable bidder rings in GBM is a strict subset of those in the corresponding English auction — every ring that is stable in GBM is also stable in English, but not conversely; (ii) for every parameter preset and

every ring size $k \in \{2, \dots, n-1\}$, there exist ring surplus levels at which the ring is English-stable but GBM-unstable.

Proof sketch. A ring member who defects by bidding on their own earns positive incentive income $d_j^{\text{GBM}} = g \cdot b_j > 0$ when outbid by ring outsiders — income that is zero in English ($d_j^{\text{Eng}} = 0$). Internal stability requires ring surplus per member to exceed the defection payoff. Since $d_j^{\text{GBM}} > d_j^{\text{Eng}} = 0$, the stability condition is strictly harder to satisfy in GBM. Any ring with surplus in $(0, \sum d_j^{\text{GBM}})$ is English-stable but GBM-unstable. ■

7.5 Why the Solvency Constraint Is the Key

The solvency constraint $\text{step_min} \geq \text{inc_max}$ was motivated by ensuring the seller’s pot never decreases (Proposition 2.1). The Sybil analysis reveals its deeper role: it is precisely the condition that makes self-outbidding non-profitable.

Remark (The Duality of the Solvency Constraint). In mechanism design, it is unusual for a single parameter constraint to deliver two logically distinct desiderata. Here, $\text{step_min} \geq \text{inc_max}$ simultaneously ensures (i) the seller’s escrow is monotonically non-decreasing (financial solvency), and (ii) self-outbidding is weakly unprofitable (Sybil resistance). Neither property implies the other in general. The solvency constraint achieves both because it aligns the incremental capital requirement with the maximum extractable incentive at the same structural point — the marginal bid.

8 Risk Aversion

GBM provides insurance: losers earn positive payoffs, reducing variance by ~9-10%. Under CARA utility, certainty-equivalent gains range from 19% (mild, $n = 2$) to 119% (high, $n = 5$). Risk aversion reduces overbidding, preserving efficiency. Full analysis in Supplementary Materials, Section S9.

In plain language (*practitioners’ box*)

For risk-averse bidders, GBM is substantially better. The guarantee of money back if you lose is worth 19-119%, depending on how risk-averse you are.

9 Mechanism-Design Characterisation

The preceding sections analyse GBM as a specific mechanism. This section asks a broader question: among all ascending auctions that subsidise losers, how special is GBM’s incentive structure? We characterise the class of *solvency-preserving* lose-payoff functions and show that GBM’s piecewise-linear schedule is essentially without loss of generality within this class.

9.1 The Design Space

Definition 9.1 (Loser-Subsidy Ascending Auction). A loser-subsidy ascending auction is an ascending auction in which an outbid bidder at standing bid x who placed bid $y \geq \alpha x$ receives a payment $\gamma(x, y) \cdot y \geq 0$, where $\gamma : \mathbb{R}_+^2 \rightarrow [0, 1]$ is a measurable *incentive schedule* and $\alpha = 1 + \text{step_min} > 1$ is the minimum bid multiplier.

Definition 9.2 (Solvency-Preserving Schedule). An incentive schedule γ is *solvency-preserving* if the seller's pot (total bids received minus total incentives paid) is non-decreasing in the number of bids, for all bid sequences. Formally: for every valid bid sequence $(b_0, b_1, \dots, b_K)$, the seller's pot after k bids satisfies $P_k \geq P_{k-1}$ for all $k = 1, \dots, K$.

Proposition 9.1 (Solvency Characterisation). An incentive schedule γ is solvency-preserving if and only if $\gamma(x, y) \leq \text{step_min}$ for all valid (x, y) with $y \geq \alpha x$.

Proof. The pot increase at bid k is $\Delta P_k = (b_k - b_{k-1}) - \gamma(b_{k-2}, b_{k-1}) \cdot b_{k-1}$. At the worst case $b_k = \alpha b_{k-1}$: $\Delta P_k = (\text{step_min} - \gamma) \cdot b_{k-1}$. Non-negativity requires $\gamma \leq \text{step_min}$. Conversely, any $\gamma \leq \text{step_min}$ ensures $\Delta P_k \geq 0$ for all $b_k \geq \alpha b_{k-1}$. ■

9.2 Optimality of Piecewise-Linear Schedules

Definition 9.3. An incentive schedule is *ratio-dependent* if $\gamma(x, y) = g(y/x)$ for some function $g : [\alpha, \infty) \rightarrow [0, \text{step_min}]$ — the incentive rate depends only on the bid ratio $r = y/x$.

Theorem 9.1 (Piecewise-Linear Optimality). Among all ratio-dependent, solvency-preserving incentive schedules $g : [\alpha, \infty) \rightarrow [0, \text{step_min}]$:

(i) *Revenue floor equivalence:* The worst-case seller revenue percentage depends on g only through $\max_{r \geq \alpha} h_g(r)$, where $h_g(r) = g(r) \cdot r / (r - 1)$. Any two schedules with the same $\max h_g$ have the same revenue floor.

(ii) *Near-optimal characterisation:* Among schedules with a given $g(\alpha) = \text{inc_min}$ and $\sup g = \text{inc_max}$, the revenue floor is minimised when h_g is constant, giving a hyperbolic schedule $g(r) = \rho\alpha(r - 1)/r$. The GBM piecewise-linear schedule $g(r) = \min[\text{inc_max}, \text{inc_min} + a(r - \alpha)/\alpha]$ with $a = \rho\alpha$ (i.e., $\mu = 0$) closely approximates this floor while uniquely achieving constant incentive cost density $\lambda(r) = \rho$.

(iii) *Equilibrium simplicity:* Under the ratio-invariant schedule ($\mu = 0$), the bid-size choice is payoff-irrelevant in the continuous-price limit (Proposition 4.6, Regime II). This eliminates bid-size gaming — a form of strategic simplicity analogous to dominant-strategy mechanisms applied to the bid-size dimension.

Proof. (i) From Theorem 5.2: the revenue floor depends on g only through $\max h_g$. (ii) Setting $h_g(r) = c$ gives $g(r) = c(r - 1)/r$ (hyperbolic); at $r = \alpha$, $c = \rho\alpha$. This minimises $\max h_g$ and thus the revenue floor. The GBM piecewise-linear schedule $g(r) = \rho(r - 1)$ (i.e., $\mu = 0$) has $h(r) = \rho r$, so $\max h = \rho\Delta$ on $[\alpha, \Delta]$ — marginally above $\rho\alpha$. The linear form is preferred because it uniquely achieves constant $\lambda(r) =$

$g(r)/(r-1) = \rho$, which by (iii) eliminates bid-size gaming. (iii) Direct: $\mu = 0$ implies $\lambda'(r) = 0$. ■

In plain language (*practitioners’ box*)

If you were designing a loser-reward auction from scratch, you’d want three things: the seller keeps as much as possible in the worst case, the incentive formula doesn’t create gaming opportunities around bid sizes, and the system stays solvent. GBM’s linear ramp is essentially the unique schedule that achieves all three simultaneously.

10 Empirical Evidence

This section reports results from an empirical analysis of 11,900 GBM and English auctions conducted during the Aavegotchi Haunt 2 sale event (August 26–29, 2021). The sale auctioned identical digital assets — “unopened portals” — through five pre-set configurations pseudo-randomly assigned by token identifier, creating a natural experiment in which the only factor differentiating auction outcomes across presets is the incentive structure itself. All bidding data is on-chain (Polygon blockchain), fully public, and independently verifiable. The dataset comprises 67,200 accepted bids placed by 2,563 unique wallet addresses across 11,900 auctions, with 44,965 incentive payments totalling 434,236 GHST ($\approx$ \$869,000). All statistical tests use bootstrap confidence intervals (10,000 resamples) and Mann-Whitney U tests for between-group comparisons; within-bidder effects are estimated via a bidder fixed-effects regression.

10.1 Experimental Design

The sale comprised 12,001 simultaneous 72-hour portal auctions. Portals were auctioned “unopened” and were identical except for their token identifier. Special token identifiers (round numbers, palindromes, etc.) were assigned the Degen preset and excluded, leaving 11,900 auctions:

Preset	n auctions	inc_min	inc_max	a	step_min	ρ
English	1,190	0%	0%	0	5%	—
LOW	3,571	0.5%	5%	0.0508	5%	0.10
MEDIUM	3,569	1%	10%	0.1125	10%	0.10
HIGH	3,570	1.5%	15%	0.187	15%	0.10

Note on preset names. The empirical presets share the names LOW, MEDIUM, HIGH with the theoretical presets defined in Section 2.7, but the absolute parameter values differ (e.g., theoretical HIGH uses $\text{step_min} = 0.10$, $\text{inc_min} = 0.01$; empirical HIGH uses $\text{step_min} = 0.15$, $\text{inc_min} = 0.015$). In both cases $\rho = \text{inc_min}/\text{step_min} = 0.10$,

so all ρ -dependent theoretical results apply identically. The λ -classification regime, however, depends on the absolute values of a and step_min , and the empirical HIGH parameters place it in Regime III ($\mu < 0$) rather than the theoretical HIGH's Regime II ($\mu = 0$); see Section 10.3 for details.

Bidders could filter auctions by preset and participated freely across types. Of 2,563 unique bidders, 980 (38.2%) bid in all four preset types, and 1,746 (68.1%) bid in at least two. The within-bidder design enables quasi-experimental inference on format effects: since portals are ex-ante identical, selection concerns reduce to bidders choosing presets based on the incentive structure itself — which is the treatment, not a confounder. A balanced-panel robustness check on the 980 all-four-preset bidders is in Supplementary Materials, Section S10.

10.2 Revenue Comparison

Preset	Median selling price	vs English	Seller revenue	Revenue share
English	100 GHST	—	100 GHST	100%
LOW	115.4 GHST	+15%	108.2 GHST	93.7%
MEDIUM	142.7 GHST	+43%	129.3 GHST	90.6%
HIGH	175.3 GHST	+75%	155.4 GHST	88.7%

Revenue share matches the $(1 - \rho) = 90\%$ prediction within 4 percentage points across all presets.

10.3 Within-Bidder Analysis

The strongest quasi-experimental test exploits bidders who participated in multiple presets. Of 2,563 unique bidders, 1,746 (68.1%) placed bids in at least two preset types.

Bidder fixed-effects regression. $\text{bid}_{ij} = \beta_{\text{LOW}} \cdot \mathbf{1}[\text{LOW}] + \beta_{\text{MED}} \cdot \mathbf{1}[\text{MED}] + \beta_{\text{HIGH}} \cdot \mathbf{1}[\text{HIGH}] + \alpha_i + \varepsilon_{ij}$ (English is omitted category):

Variable	Coefficient	Robust SE	t -stat
LOW	+54.95	1.68	32.77***
MEDIUM	+65.46	1.84	35.48***
HIGH	+115.81	2.29	50.61***

$R^2_{\text{within}} = 0.061$. HC1 standard errors. $N = 62,460$ bids, 1,746 bidders.

The monotone ordering $\text{LOW} < \text{MEDIUM} < \text{HIGH}$ confirms a quasi-experimental dose-response effect. The LOW-vs-English comparison ($\beta_{\text{LOW}} = +54.95$, $t = 32.77$) isolates the pure incentive effect (same minimum step). Robustness checks in Supplementary Materials, Section S10.

10.4 Summary

Theoretical prediction	Reference	Empirical finding	Supported?
Overbidding	Thm. 4.3	Median bid ratio > 1 ; monotone in incentive	Yes
Revenue share $\approx (1 - \rho)$	Prop. 5.3	88.7%–93.7% (pred: 90%)	Yes
Bidder preference (FOSD)	Thm. 6.3	Same bidder bids 58% higher (HIGH vs Eng)	Yes
Revenue dominance	Thm. 6.4	Revenue: +15% (LOW) to +75% (HIGH)	Yes
Sybil deterrence	Thm. 7.1	Self-outbid: 5% GBM vs 23% English	Yes
Entry advantage	Thm. 6.4	Not testable (shared bidder pool)	—

Limitations. (i) Self-selected crypto-native bidders — likely inflates overbidding relative to conventional settings. (ii) Fungible digital assets with speculative valuations weaken the IPV assumption. (iii) The simultaneous design precludes testing entry predictions (a clean test requires separated bidder pools). (iv) The LOW coefficient isolates the pure incentive channel; MEDIUM and HIGH confound incentives with step-size mechanics.

Acknowledgments

The author thanks Karim Elhadji Tchiambou and Lionel Blondé for their contributions to the numerical simulations and computational verification underlying this work.

Conflict of Interest Disclosure

The author is a co-inventor of the GBM auction mechanism (US20210174432A1, priority date 7 August 2018) and a principal at a company that operates a commercial platform deploying GBM auctions. The theoretical analysis and all proofs were developed independently of commercial considerations. All code for numerical verification and Monte Carlo simulations is available at [repository URL]. The author has received no external funding for this research.

References

Augenblick, N. (2016). The sunk-cost fallacy in penny auctions. *Review of Economic Studies*, 83(1), 58–86.

- Bergemann, D., and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73(6), 1771-1813.
- Bulow, J., and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1), 180-194.
- Che, Y.-K., and Gale, I. (1998). Caps on political lobbying. *American Economic Review*, 88(3), 643-651.
- Goeree, J. K., and Offerman, T. (2004). The Amsterdam Auction. *Econometrica*, 72(1), 281-294.
- Graham, D. A., and Marshall, R. C. (1987). Collusive bidder behavior at single-object second-price and English auctions. *Journal of Political Economy*, 95(6), 1217-1239.
- Klemperer, P. (2002). What really matters in auction design. *Journal of Economic Perspectives*, 16(1), 169-189.
- Krishna, V. (2002). *Auction Theory*. Academic Press.
- Levin, D., and Smith, J. L. (1994). Equilibrium in auctions with entry. *American Economic Review*, 84(3), 585-599.
- Maskin, E., and Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, 52(6), 1473-1518.
- McAfee, R. P., and McMillan, J. (1992). Bidding rings. *American Economic Review*, 82(3), 579-599.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.
- Milgrom, P., and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5), 1089-1122.
- Moldovanu, B., and Sela, A. (2001). The optimal allocation of prizes in contests. *American Economic Review*, 91(3), 542-558.
- Morgan, J. (2000). Financing public goods by means of lotteries. *Review of Economic Studies*, 67(4), 761-784.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1), 58-73.
- Riley, J. G., and Samuelson, W. F. (1981). Optimal auctions. *American Economic Review*, 71(3), 381-392.
- Samuelson, W. F. (1985). Competitive bidding with entry costs. *Economics Letters*, 17(1-2), 53-57.