

Competing Auctions in Intermediated Markets

Bruno Mazorra

Flashbots, Barcelona, Spain

Minghao Pan

California Institute of Technology, Pasadena, CA, USA

Flashbots, Pasadena, CA, USA

Christoph Schlegel

Flashbots, Zurich, Switzerland

Abstract

We analyze competing auctions in intermediated markets, where a seller selects among parallel mechanisms for the sale of a single good, most prominently the relay-and-protocol architecture of proposer-builder separation in Ethereum.

When the intermediary can enforce single-homing on its bidders, sealed-bid second-price intermediary auctions fully unravel into the sealed first-price principal auction; open bidding-format intermediaries unravel only partially, collapsing into first-price in equilibrium under symmetric latency and sorting fast bidders to the intermediary under asymmetric latency. Any last-look advantage is removed through the availability of a credible sealed bidding channel. These results extend to multi-plexing environments (no enforcement by the intermediary).

While the unraveling result indicates that the availability of a sealed first-price bidding channel pushes the overall market to the same auction structure, the very assumption of the credibility of such a channel is problematic, as the seller may have an incentive to leak information: a first-price auction is leakage-resistant in the presence of a single “fast” bidder but not against two or more. However, if the seller can credibly commit to not leak bids, it is optimal for them to do so.

A main motivation is the forthcoming Glamsterdam update of Ethereum: our analysis suggests that the availability of an in-protocol (first-price) bidding channel severely limits the design space for out-of-protocol auctions by relays and other intermediaries.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory and mechanism design

Keywords and phrases Competing auctions, proposer-builder separation, maximal extractable value

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.19

Related Version *Full Version:* <https://arxiv.org/abs/2606.06633>

1 Introduction

In many markets, the same good can be bought through different channels, often through competing auction mechanisms. In classical online marketplaces (eBay etc.) competing auctions usually mean that there are competing sellers of the same type of good, and the sellers themselves may multi-home, offering their goods or services across different platforms at the same time (prominent examples include ride-sharing services). In other contexts, however, there may be a single seller of a single good, while competition happens between intermediaries. Competing auctions by competing sellers have been studied in a rich literature starting with [10]; competing auctions run by intermediaries for a single good have received much less attention.

This is not only a theoretical gap, but practically very relevant. The most familiar example from traditional finance is order routing: a broker with a single order to execute chooses among competing venues running different auction mechanisms (lit exchanges, dark pools, retail wholesalers, periodic auctions, RFQ systems) and is required under best-execution



© Bruno Mazorra, Minghao Pan, and Christoph Schlegel;
licensed under Creative Commons License CC-BY 4.0

8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 19; pp. 19:1–19:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

rules to route the order to whichever venue offers the best outcome. Blockchain systems, which will be our main motivating examples, are even more abundant in intermediation services, and we observe competing auctions throughout the transaction inclusion pipeline. Ethereum’s MEV ecosystem offers a particularly rich set of examples, from order flow auctions (MEV-protect, MEV-blocker, etc.) over routing and RFQ services (1inch, CowSwap etc.) to builder competition (Titan, BuilderNet etc.) for searcher bundles. In each case, the transaction or intent is the single good, and competing intermediaries run different auction mechanisms to capture it.

A forthcoming Ethereum protocol upgrade, enshrined Proposer-Builder Separation (ePBS) [6], introduces another instance of competing auction mechanisms. Currently, Proposer-Builder Separation, the market structure through which block builders bid for their blocks to be proposed, is implemented out-of-protocol with the help of intermediaries (so-called relays). Under ePBS, the protocol provides a native bid-and-commitment layer for builder bids, while off-protocol channels may still coexist with it: bids may be routed through public protocol gossip, private proposer-facing builder RPCs, or new relay-like intermediaries. Thus, we will possibly observe competition between direct in-protocol and intermediated bids.

In this paper, we study this competition between an in-protocol bidding channel and intermediated ones. Our analysis predicts that the introduction of an in-protocol sealed-bid first-price channel severely constrains the design space available to off-protocol intermediaries, and we characterize which relay formats can and cannot survive this competition.

We analyze the question of what an in-protocol bidding channel does to relay intermediation from several angles: first we study what kind of pricing rule the competing auction can maintain in the presence of the in-protocol outside option. A plausible alternative choice to a first-price auction (sealed or open in the form of an English auction) would be to run a second-price auction format. A second-price auction has the advantage that bidders have an incentive to bid their value truthfully in it.¹ In the presence of a competing first-price channel, however, bidders still bid truthfully *conditional on entering* the second-price auction (Lemma 2), but entering it at all becomes unattractive: in the single-homing case, the second-price auction unravels and bidders self-select into the first-price (in-protocol) auction (Theorem 1). In the multi-homing case, the good will almost surely be allocated through the first-price auction (Theorem 4).

Second, we study how information can be shared throughout the auction: open bidding as in the status quo is naturally undermined by the availability of a sealed-bid in-protocol channel. Bidders have an incentive to listen to an open-bid feed but to bid through a sealed channel if available. However, the seller (or an intermediary acting on the seller’s behalf) can choose to share bids received during the auction period, orchestrating a (partially) open-bid auction. Such information sharing naturally unravels in the sense that bidders have an incentive to bid as late as possible to avoid having their bids shared, effectively turning the auction into a sealed-bid auction (Propositions 7 and 14). However, in the presence of latency differences by bidders, information sharing (of the slower bidders’ bids with the faster bidders) may be feasible. We analyze this in a model where a subset of bidders is “fast” and the seller or intermediary may share bids with these bidders. We find that sharing information with a single fast bidder is not optimal for the seller (Proposition 8), whereas sharing information with two or more fast bidders (in case there are that many fast bidders) is

¹ Currently, some relays like the Ultrasound relay experiment with “second-price” features.

■ **Table 1** Summary of results. Each row corresponds to one main proposition or theorem in the paper, organized by section.

Section	Setting	Result
§ 3.1	Single-homing, sealed second-price relay	The native first-price auction is the unique symmetric equilibrium outcome (Theorem 1).
§ 3.2	Multi-plexing, second-price relay	The item is allocated through the first-price auction with probability 1 (Theorem 4).
§ 3.3.1	No last-look advantage	All-first-price is an equilibrium; the all-English profile is unstable in large markets (Proposition 7).
§ 3.3.2	One fast bidder	The first-price bidding channel is leakage-resistant (Proposition 8).
§ 3.3.2	Multiple fast bidders, $ S \geq 2$	The proposer leaks in equilibrium (Theorem 9).
§ 3.3.3	Last-look with credible proposer commitment	The proposer weakly prefers to commit to a credible sealed-bid first-price channel (Proposition 13).

a best response for the seller (Theorem 9) and fast bidders would benefit from bid disclosure (Proposition 11). However, if the seller can credibly commit ex-ante to a disclosure rule, the optimal commitment is to share no information at all (Proposition 13).

The interpretation of the “information disclosure” results in the Ethereum context is straightforward: in the absence of a credible first-price sealed-bidding channel, if there are at least two “fast” bidders, a proposer would have an incentive to leak the bids it has received (in or out-of protocol), effectively making the in-protocol bidding channel leaky. Practically, this could be orchestrated through a service where proposers forward the received bids to a relay that integrates them in their bid stream. However, there are potentially ways to establish a credible sealed bid first-price channel: proposers could use reputation (in the case of large staking pools like Lido) or hardware solutions such as Trusted Execution Environments (TEEs) to commit to not leak bids that they have received. Alternatively, they can delegate to an intermediary, i.e. a potential “sealed-bid first-price” relay. As commitment to non-disclosure is optimal if possible, such a solution can quite plausibly recover the no-disclosure equilibrium institutionally.

We proceed as follows: in the next two subsections, we provide background on the literature on competing auctions, and institutional background on Ethereum’s blockbuilding stack. In Section 2 we discuss our modeling choices for the subsequent analysis. The discussion is more detailed than the usual scope of such a section, primarily because some modeling choices are very much informed by institutional details of the block building ecosystem in Ethereum which we want to discuss in this context. In Section 3 we derive our results, which are summarized in Table 1. In Section 4, we conclude with practical implications of our analysis and open questions.

1.1 Competing Auctions & Related Literature

The literature on *competing auctions* [10, 15, 3, 7] studies the constellation where multiple sellers each offer their own good through a chosen mechanism, with buyers single-homing across sellers. The standard result is that competition pushes mechanisms toward second-price with a reserve in the large-market limit and disciplines reserve prices. Our setting differs structurally: a single seller has a single good and faces multiple parallel mechanisms run by intermediaries, with the winning mechanism determined ex-post by a max-revenue

rule. The strategic forces are therefore quite different: bidders are not choosing between substitute goods but between channels for the same allocation, and competition operates on intermediaries rather than on sellers.

The literature on *platform competition* [4, 1, 16] introduces the distinction between single-homing and multi-homing that also plays a role in our analysis. The general finding – that buyer-side multi-homing intensifies competition between platforms, while single-homing produces winner-take-all dynamics – is consistent with the broad shape of our results, though our setting is narrower and our focus is on equilibrium auction outcomes rather than platform pricing.

Our information leakage results connect to the literature on *seller-controlled information disclosure in auctions*. The classical result here is the linkage principle of Milgrom and Weber [11], which establishes that under affiliated and interdependent values, the seller benefits from committing to reveal information that links the price more strongly to bidders’ signals. This logic does not apply in our setting: at the builder level, values are private (although possibly drawn from a correlated distribution) and known with certainty, so additional information about other bidders’ valuations has no effect on a bidder’s own valuation. The relevant disclosure question becomes strategic rather than informational: the seller’s choice is not about helping bidders refine their valuations, but about changing the strategic environment by giving some bidders the ability to react to others’ bids. The closer reference is therefore the literature on information design in auctions under private values [2]. The commitment-vs-no-commitment distinction in our results has a natural interpretation in terms of whether disclosure is delegated to a trusted intermediary.

The closest non-blockchain analog of our setting is in *market microstructure*, where a substantial literature studies how a single trade is routed across competing exchanges, dark pools, and wholesalers under best-execution obligations [8, 12]. A particularly close institutional parallel to our latency-advantaged-bidder setting is the case of last-look provision in FX markets [13], where dealers selectively respond to client quote requests based on observed market information.

1.2 From MEV-Boost to ePBS

Ethereum’s current block-building market is organized through MEV-Boost, an off-chain implementation of proposer-builder separation. Validators outsource execution-payload construction to specialized builders. Builders compete by submitting blocks and associated bids to relays, and the proposer signs the highest-paying blinded block header returned by MEV-Boost. The relay is therefore the central intermediary: it runs the auction, hides the block contents until the proposer has committed, and then releases the full payload. In the standard, non-optimistic flow, the relay also verifies block validity and proposer payment before forwarding the bid. The builder sets itself as the block’s `feeRecipient` and includes a final transfer to the proposer’s registered `feeRecipient` address; the relay checks that this transfer matches the advertised bid.

Relay designs have evolved to reduce latency on the critical path. In optimistic relay designs, builders post collateral, allowing the relay to accept bids before completing full validation. The relay still intermediates the auction and forwards the blinded header to the proposer, but payment and validity checks are performed ex-post. If the winning block is invalid or fails to pay the promised amount, the builder’s collateral is used to compensate the proposer. Optimistic relaying therefore makes the MEV-Boost auction more latency-efficient, since builders can submit competitive bids later in the slot, but it does not remove the relay as the trusted off-chain intermediary.

Thus, MEV-Boost implements PBS without changing Ethereum consensus, but it does so by relying on off-chain infrastructure. Relays solve the fair-exchange problem between proposers and builders: builders need protection against payload theft or payload leakage before commitment, while proposers need assurance that the promised block is valid and pays the advertised amount. The resulting market is naturally modeled as an open, continuous, relay-mediated auction. Builders can update bids during the slot, relays intermediate the information flow, and latency determines how late bids can arrive and still be competitive.

Enshrined proposer-builder separation, or ePBS, changes this architecture by moving the core commitment-and-payment logic into the protocol. In the current EIP-7732 design [6], the beacon block no longer contains the full execution payload on the attestation-critical path. Instead, it contains a signed builder commitment, the `SignedExecutionPayloadBid`, which commits to a later execution payload and specifies the value to be paid to the proposer. Builders become protocol-recognized entities with beacon-chain balances, so the committed payment can be deducted from the builder's balance and later credited to the proposer. Timely reveal is checked by a Payload Timeliness Committee (PTC), while full execution validation is deferred to a later point in the slot. The current EIP-7732 design has some potential flaws if payload-delivery failure is not sanctioned and the PTC deadlines are positioned too far in the slot [9]. As there are plausible mitigation mechanisms for this problem, we assume throughout the analysis that these “free-option” considerations do not play a role for the questions we seek to answer in this paper.

These protocol changes also alter the market structure. Under MEV-Boost, the relay observes and intermediates the bidding process, and the auction is effectively open, continuous, and latency-sensitive. Under ePBS, the protocol provides a native bid-and-commitment layer, while off-protocol channels may still coexist with it. The post-ePBS market may therefore contain competing auction formats: a native in-protocol mechanism and external auction channels that can differ in observability, timing, trust assumptions, and payment guarantees. Our goal is to understand when these mechanisms coexist, when one unravels into the other, and how the change from trusted relay-mediated exchange to protocol-mediated exchange affects builder and proposer incentives.

Trustless vs Trusted Payments. In MEV-Boost, payments are mediated by trusted off-chain infrastructure: the relay verifies that the builder's block contains the promised payment and releases the payload only after the proposer has signed the blinded block. In ePBS, by contrast, payment can be enforced in protocol. A builder's signed bid commits to a payment amount, which can be deducted from the builder's beacon-chain balance and later credited to the proposer.

The main cost of trustless payments is that builders must pre-fund their bids. This can increase entry barriers, since a builder needs sufficient in-protocol balance to participate. For ordinary blocks, this constraint may be mild: most bids are small, and if deposit and withdrawal frictions are low, the required working capital need not be large. For high-value blocks, however, the balance constraint can become binding. A builder with high value but insufficient balance may be unable to express its full willingness to pay through the trustless mechanism. In such cases, bids may still be routed through trusted relationships, relays, or other intermediaries that can provide credit, guarantees, or faster settlement. Subsequently, we will not model this particular friction explicitly, but the reader should keep in mind that inter-mediated bids may still have a role and certain advantages over in-protocol bidding, even if trustless bidding would be optimal based on “pure auction considerations” that are the main focus of our analysis.

2 Modeling Consideration

2.1 Private Values

We will work throughout the paper in a private value setting. Some justification is needed here, as this often seems to be a point of possible confusion. Most block value is generated from financial transactions, which inherently have a common-value component (with at least some uncertainty and heterogeneous signal on this value). However, this does not mean that block builders have a common value for proposing the next block. An uncertain common value is effectively realized up-stream at the searcher level. Block builders receive bundles and compose them into a block whose value for the builder is deterministic and known with certainty. This would be different if builders did searching themselves, a market dynamic called “searcher-builder integration” that has been discussed extensively in the past [14]. However, this dynamic has never materialized in reality, and at the current point in time the dominant Ethereum builders (Titan, BuilderNet, Quasar) are neutral and do not search. In particular, “price discovery” in the PBS auction, sometimes cited as an advantage of the current open auction format, is – if at all – relevant for searchers to update their bids to the builders. Thus, our analysis focuses on the practically relevant case of private values. As different builders tend to receive similar bundles (or even the same bundles from the same searcher), these private values are naturally positively correlated. However, we refrain from modeling the correlation structure explicitly, for the sake of analytical simplicity. We note, however, that positive correlation generally pushes in the same direction as the logic of most of our results, so that they likely generalize to models of positively correlated private valuations.² Thus, throughout the paper we assume that valuations are drawn i.i.d. from a distribution F that is absolutely continuous and strictly increasing on a compact interval $[v, \bar{v}]$. We denote its density by f .

2.2 Single vs Multi-plexing of Bids

The in-protocol auction in ePBS is permissionless. The protocol can neither restrict who enters the auction nor restrict whether they submit multiple bids through different channels (in-protocol and directly to the proposer or through a relay) or just a single bid. Thus, the ePBS auction is by default “multi-plexing” friendly. The case of relays is more subtle. Relays already today restrict certain functionalities (optimistic relaying) to registered and staked participants while being generally open to all participants for their standard service. In a future with ePBS, it is still to be seen whether relays become “more permissioned” than they are today. One potential reason they could develop in this direction is that they want to offer additional services to remain attractive even with an in-protocol mechanism in place; for example, block-merging by the relay is easier to coordinate with a permissioned set of builders. More importantly for our context, a relay might want to sustain an auction format that is otherwise not sustainable in the multi-plexing case. The current open-bidding PBS

² A trivial form of introducing correlation for which our result would hold without modification of the proofs is through a block-level public state: Fix a block and condition on its realized public state; let F denote the value distribution induced by that state. Assume there is a family $\mathcal{F}$ of distributions supported on the same compact interval. For the realized distribution $F \in \mathcal{F}$, bidders’ values satisfy

$$v_i \mid F \sim_{\text{i.i.d.}} F, \quad i = 1, \dots, n.$$

Thus values can be correlated before conditioning: the public state can shift the distribution faced by all bidders in the block. Economically, F captures the public market conditions of the current block. Whenever a proposition refers to F , it can be read as conditional on an arbitrary realized $F \in \mathcal{F}$.

auction, for example, is hard to maintain in an ePBS world: if a relay offers an open-bidding mechanism, then bidders have a natural incentive to listen to the bidding feed of the relay but to place their own bid in a sealed way directly to the proposer, with the relay auction unraveling as a consequence. If a relay wants to fight this logic, it might restrict bidding to a permissioned set of builders and impose that, in case a builder bids with them in a particular auction round, they need to commit to not bidding through an external channel in the same round. This assumes, of course, that deviations are detectable with sufficient probability by the relay. If a relay manages to implement single-homing within its service, it makes the entire market single-homing: bidders need to select into one of the two mechanisms before they bid.

We find both scenarios – single-homing enforced at the relay level and multi-homing – plausible at first view. Thus, we think it is valuable to analyze both cases.

2.3 Latency and Last-Look Advantage

The performance and equilibrium analysis of auctions in time-sensitive environments may be influenced by bidder latency considerations. The current MEV-Boost auction, for example, is naturally modeled, at an idealized level, as an English or candle-style auction. Relays receive continuously updated bids during the slot, builders can react to the current competitive level, and the process ends only when the proposer must commit to a block. But in the current MEV-Boost world, bid streams and timing frictions make latency economically meaningful: a builder who observes a competing bid and can still respond before the deadline while others cannot obtain a form of *last look* [14] advantage.

For our analysis of open bidding and bid disclosure we, therefore, consider two types of environments: We first consider the clean benchmark in which no bidder has a last-look advantage. In that benchmark, the relay auction may reveal information internally among relay participants, but bidders can submit a last bid that is not observed by competitors. Moreover, the relay winner does not observe the native first-price bids before choosing the proposer-facing bid and vice versa.

The no-last-look benchmark is an idealization. We then study the more interesting case of last look explicitly. We treat technological latency as exogenous, but the exploitation of latency as endogenous: low-latency builders can use their advantage only if some information is made available before the deadline. That information may be public, relay-mediated, or selectively disclosed by the seller.

This leads to two versions of the last-look problem. In the first, the seller cannot credibly commit to a disclosure strategy. After receiving early bids, the seller may decide history by history which private bids to leak and to whom (of the latency advantaged bidders).³ In the second, the seller can commit ex-ante to a specific disclosure rule. Practically, this disclosure can be implemented through a trusted intermediary, such as a relay in the PBS case. The distinction matters because a sealed-bid procedure can be ex-ante attractive, while after the early bids have arrived the seller may have an incentive to leak them to latency-advantaged builders.

³ For the PBS application, practically, the proposer will still decide ex-ante, before the auction, how to configure their node (or more likely which third-party side-car software to install). The point is that while the proposer chooses their strategy ex-ante in this application, they cannot credibly commit to a disclosure strategy that is not an equilibrium strategy.

3 Results

Proofs omitted for space are provided in the full version.

3.1 Single-Homing: First-Price vs Second-Price Auction

We first analyze the single-homing case. There are $n \geq 2$ bidders who bid for the right to decide on the content of the next block. The market structure is as follows: there are two competing auctions. The first auction is a first-price sealed-bid auction (the native ePBS auction) and the second a competing relay auction. The bidders can submit a bid to one of the two auctions or neither auction. The two auctions are run on the submitted bid. Whichever auction yields higher revenue allocates the item and charges the winner.

Bidders choose ad interim, i.e. after they know their valuation but before they send their first bid, through which channel they bid in the auction. The natural interpretation is that of an intermediary (relay) that can detect deviations (with high probability) where bidders bid through both channels and punish this “multi-plexing” behavior.

We first study whether a competing auction can enforce a second-price pricing rule. We restrict attention to symmetric equilibria.⁴

► **Theorem 1** (Single-homing unraveling). *If bidders must choose to participate in exactly one of the first-price auction and the sealed second-price auction, then, up to F -null sets, there is a unique symmetric Bayes–Nash equilibrium. In this equilibrium, every type chooses the first-price auction and bids according to the standard symmetric first-price equilibrium.*

We separate the proof into several lemmas.

► **Lemma 2.** *In equilibrium, if an agent bids in the second-price auction, then she bids truthfully almost surely.*

Proof. By symmetry, we may assume that bidders who participate in the second-price auction bid truthfully almost surely. The reason is simple. Truthful bidding weakly dominates every other bid in the second-price auction: lowering the bid below one’s value can only remove wins at profitable prices, while raising the bid above one’s value can only add wins at loss-making prices. Moreover, in a symmetric equilibrium any positive-measure set of non-truthful second-price bids would make this domination strict on a positive-probability event. If a positive measure of types overbid, then by symmetry there is positive probability that all second-price participants have values below some cutoff but bids above it; the winner then pays more than her value, while truthful bidding would avoid this loss. If a positive measure of types underbid, then with positive probability all second-price participants have values above some cutoff but bids below it; at least one bidder then loses even though a truthful bid would win at a price below her value. Thus any non-truthful second-price bid used on a positive-measure set would admit a strictly profitable deviation to the truthful bid, contradicting equilibrium. Hence, conditional on participating in the second-price auction, bidders bid truthfully almost surely. ◀

⁴ A natural question is whether the unraveling results in this and the following section follow more directly from revenue equivalence. Revenue equivalence determines the expected payment in any efficient symmetric mechanism with zero surplus at the lowest type, but it presupposes a fixed allocation rule. In our setting, the allocation rule is itself an equilibrium object of interest. In principle, the good can be allocated inefficiently – for example, if a high-value bidder selects themselves into the second-price auction but is outbid by a lower-value bidder in the first-price auction. We therefore cannot invoke revenue equivalence to short-cut the analysis.

► **Lemma 3.** *If two bidders with types $v, w \in (\underline{v}, \bar{v}]$ participate in the first-price auction with positive probability, then $v > w$ implies $b(v) > b(w)$.*

Proof of the theorem. For notational simplicity, normalize the support to $[0, 1]$. Let μ be the distribution associated with F . Suppose by contradiction that there exists a symmetric equilibrium with positive participation in the second-price auction. Let $E \subseteq [0, 1]$ denote the set of types who participate in the second-price auction.

We first show that $\alpha := \text{ess inf } E > 0$. A type with valuation v , by Lemmas 2 and 3, obtains from the second-price auction payoff

$$U_2(v) := \int_{[0, v] \cap E} (v - t) \sum_{k=1}^{n-1} \binom{n-1}{k} k \cdot \mu(b^{-1}([0, t]) \cap E^c)^{n-1-k} \mu([0, t] \cap E)^{k-1} d\mu(t).$$

Here $\binom{n-1}{k}$ samples k other participants of the second-price auction, and the following factor k selects one participant to have value t .

Suppose that the function

$$t \mapsto (v - t) \sum_{k=1}^{n-1} \binom{n-1}{k} k \cdot \mu(b^{-1}([0, t]) \cap E^c)^{n-1-k} \mu([0, t] \cap E)^{k-1}$$

is maximized at some $t = t^*$. Then

$$U_2(v) \leq \mu([0, v] \cap E) (v - t^*) \sum_{k=1}^{n-1} \binom{n-1}{k} k \mu([0, t^*] \cap E)^{k-1} \mu(b^{-1}([0, t^*]) \cap E^c)^{n-1-k}. \quad (1)$$

If the type chooses to bid t^* in the first-price auction instead, her utility payoff is at least

$$(v - t^*) \mu(b^{-1}([0, t^*]) \cap E^c)^{n-1} + (v - t^*) \sum_{k=1}^{n-1} \binom{n-1}{k} \mu(E) \mu([0, t^*] \cap E)^{k-1} \mu(b^{-1}([0, t^*]) \cap E^c)^{n-1-k}. \quad (2)$$

Here $\binom{n-1}{k} \mu(E) \mu([0, t^*] \cap E)^{k-1}$ comes from first sampling k participants in the second-price auction, letting one participant bid arbitrarily, and having the remaining $k - 1$ participants bid less than t^* .

By comparing (2) with (1) term by term, as long as

$$(n - 1) \mu([0, v] \cap E) < \mu(E),$$

we have (2) > (1). This is true when v is small enough. In other words, when v is small enough, v prefers the first-price auction to the second-price auction, and so $\alpha > 0$.

Let $\varepsilon > 0$ be small such that $\alpha + \varepsilon \in E$. When participating in the second-price auction, bidder $\alpha + \varepsilon$ has a chance of winning bounded above by

$$(n - 1) \mu([\alpha, \alpha + \varepsilon]).$$

So her utility payoff is at most

$$\varepsilon (n - 1) \mu([\alpha, \alpha + \varepsilon]).$$

However, if she bids α in the first-price auction, she wins whenever all other participants have value smaller than α , so her utility payoff is at least

$$\varepsilon \mu([0, \alpha])^{n-1}.$$

Thus, taking ε sufficiently small,

$$\mu([0, \alpha])^{n-1} > (n-1)\mu([\alpha, \alpha + \varepsilon]),$$

and so, bidder $\alpha + \varepsilon$ prefers the first-price auction, a contradiction. $\blacktriangleleft$

Theorem 1 implies that, under single-homing, the sealed second-price relay does not attract participation: in the unique symmetric equilibrium, all bidders bid in the first-price auction, and the item is allocated through that channel. In the Ethereum interpretation, the block is therefore allocated through the native direct-communication channel rather than through the relay. This conclusion, however, relies on the relay's ability to enforce single-homing. In a permissionless environment, a relay may restrict the identities that participate in its auction, but it cannot generally verify that an identity corresponds to a unique economic bidder. A bidder who wants to bid in both channels can use Sybil identities, submitting one bid through the relay and another through the native auction while appearing as distinct participants. Thus, single-homing has bite only to the extent that such deviations can be detected. When Sybil participation is feasible, bidders can circumvent the restriction by splitting bids across identities, so the effective environment is one of multi-homing. We analyze that case next.

3.2 Multi-Homing: First-Price vs Second-Price Auction

In this section, we prove that for the multi-plexing case with competing first-price and second-price auctions, there is an essentially unique symmetric pure-strategy Nash equilibrium in which the second-price auction unravels. The precise statement is the following.

► **Theorem 4** (Essential uniqueness). *Let b^{FP} be the unique symmetric pure-strategy equilibrium bidding function of the first-price auction. Every symmetric pure-strategy Nash equilibrium $v \mapsto (b(v), w(v))$ such that b is a.e. strictly increasing whenever the agent bids in the first-price auction is of the form $v \mapsto (b^{\text{FP}}(v), w(v))$ with $w(v) \leq b^{\text{FP}}(v)$.*

Under any such equilibrium, the item is always allocated through the first-price auction. Hence the second coordinate w is payoff-irrelevant and may, without loss of generality, be taken to be identically zero.

Throughout this subsection, fix a symmetric pure-strategy equilibrium $v \mapsto (b(v), w(v))$ of the two-auction game, and define

$$P := \{v : w(v) > b(v)\}.$$

We identify strategies that differ only on an F -null set. We also use the standard observation that first-price overbidding is weakly dominated, so we may take $b(v) < v$ for every $v > \underline{v}$.

► **Lemma 5.** *Fix a type v and a first-price bid $x \leq v$. Then, for every second-price bid $s > x$ and $s \neq v$, the action (x, v) weakly dominates (x, s) and $w(v) = v$ for almost every $v \in P$.*

► **Lemma 6.** *For almost all types v , the agent participates in the first-price auction.*

Proof. Observe that not participating in the first-price auction is payoff equivalent to bidding $\underline{v}$, and so wlog, we consider

$$P_0 := \{v \in P : b(v) = \underline{v}\}$$

as the set of agents that do not participate in the first-price auction. We first show that $\Pr[V \in P_0] = 0$. By Lemma 5, we may take $w(v) = v$ for almost every $v \in P$.

Suppose, towards a contradiction, that $\Pr[V \in P_0] > 0$, and let $\alpha := \text{ess inf } P$. Define

$$\varphi(x) := \Pr[b(V) \leq x, w(V) \leq x].$$

Thus $\varphi(x)^{n-1}$ is the probability that all opponents submit both bids below x .

Fix $v \in P$. Since $w(v) = v$, the payoff from the equilibrium action is bounded above by

$$U(v) \leq (v - b(v))\varphi(b(v))^{n-1} + \int_{[b(v), v] \cap P} (v - t)(n - 1)\varphi(t)^{n-2} dF(t).$$

Let $t^* \in [b(v), v]$ maximize $(v - t)\varphi(t)^{n-2}$. Then

$$U(v) \leq (v - t^*)\varphi(t^*)^{n-1} + (v - t^*)(n - 1)\varphi(t^*)^{n-2} \Pr[V \in P \cap [b(v), v]].$$

Now let type v deviate by bidding t^* in the first-price auction and by making her second-price bid irrelevant. This deviation gives payoff at least

$$(v - t^*)\varphi(t^*)^{n-1} + (v - t^*)(n - 1)\varphi(t^*)^{n-2} \Pr[V \in P_0 \cap [t^*, \infty)].$$

Since $t^* \leq v$, we have

$$\Pr[V \in P_0 \cap [t^*, \infty)] \geq \Pr[V \in P_0 \cap [v, \infty)].$$

For $v \in P$ sufficiently close to α , the term $\Pr[V \in P \cap [b(v), v]]$ is arbitrarily small, whereas $\Pr[V \in P_0 \cap [v, \infty)]$ is bounded away from zero. Hence the deviation is strictly profitable, a contradiction. Therefore

$$\Pr[V \in P_0] = 0.$$

It remains to rule out types with $b(v) = \underline{v}$ and $w(v) \leq b(v)$. Such types obtain zero payoff. But any type $v > \underline{v}$ can bid $(x, 0)$ for some $x \in (\underline{v}, v)$ and obtain strictly positive payoff on the event that all opponents' values are below x . Hence this cannot be a best response. Therefore $b(v) > \underline{v}$ for almost every v . $\blacktriangleleft$

Proof of Theorem 4. By the previous lemma and by assumption b is globally strictly increasing. We will show that $P = \emptyset$. Suppose, for contradiction, that there exists $v \in P$. By Lemma 5, $(b(v), v)$ is a best-reply. Now, consider the tuple-action (x, v) for an agent i with valuation v . His interim payoff is

$$\Pi(v, x) = (v - x) \Pr[B \leq x, W \leq x] + \mathbb{E}[(v - W) \mathbf{1}\{W > \max\{B, x\}, W \leq v\}],$$

where

$$W = \max_{j \neq i} \{w(v_j)\} \quad \text{and} \quad B = \max_{j \neq i} \{b(v_j)\}.$$

Now, we assume that Π is differentiable at $(v, b(v))$ ⁵. The derivative of Π with respect to x is

$$\frac{\partial \Pi(v, x)}{\partial x} = -\Pr[B \leq x, W \leq x] + (v - x)f_B(x) \Pr[W \leq x \mid B = x].$$

⁵ This assumption can be dropped by taking the left and right limits of Π and computing them directly. For simplicity, we do not include this argument.

Conditional on $B = b(v)$, since b is strictly increasing, we must have that there exists another $j \neq i$ such that $v_j = v$, so $W > b(v)$ and $\Pr[W \leq b(v) \mid B = b(v)] = 0$. Therefore,

$$\frac{\partial \Pi(v, b(v))}{\partial x} = -\Pr[B \leq b(v), W \leq b(v)].$$

However, $\Pr[B \leq b(v), W \leq b(v)] \geq \Pr[\max_{j \neq i} \{v_j\} \leq b(v)] = F(b(v))^{n-1}$ and since b is strictly increasing, $F(b(v))^{n-1} > 0$. Therefore, $\frac{\partial \Pi(v, b(v))}{\partial x} < 0$. This contradicts the fact that bidding $b(v)$ in the FP is a best-reply by first-order conditions, so there is no $v \in P$. ◀

When agents are allowed to use Sybil identities, the equilibrium analysis becomes more nuanced. The first-price only equilibrium is still an equilibrium even with Sybil-strategies. However, uniqueness can generally fail: For instance, an agent can submit the same bid through two identities in a second-price auction. If those two identities are the highest bidders, the agent pays her own bid, so the outcome effectively replicates a first-price bid.

3.3 Sealed and Open Bidding

Next, we look into the question whether some form of open bidding could emerge in an ePBS world. There are two fundamental and related questions to be answered:

1. In the presence of a credible sealed bid first-price bidding channel, can an alternative auction format with open bidding survive?
2. If in the first-price bidding channel (generally we think here of the in-protocol channel), the seller could choose to leak received bids, would he do so?

The first question is the English auction version of the question we studied in the last section. The analysis in this section applies to both single-homing and multi-homing environments; the participation distinction is orthogonal to the disclosure question. The answer in the symmetric case with equally fast bidders is that the English auction will unravel (or all bidders will bid at the “last” possible moment, essentially turning the auction into a sealed-bid first-price auction). In the asymmetric case, we will observe partial un-raveling: the slow bidder will self-select into the sealed-bidding channel, to avoid leaking information to bidders with a last look. The single- vs multi-homing distinction does not matter for these kinds of results. The conclusion from these results is that in the presence of a credible sealed-bid first-price bidding channel, the overall market is pushed in one way or another to a sealed-bid first-price auction architecture.

The answer to the second question will add important qualifications to the conclusion drawn from the answer to the first question: whether or not the seller has an incentive to keep the first-price bidding channel sealed depends on how many bidders have a last look advantage. In the symmetric case, as well as in the case of one fast bidder with last look, the seller has no incentive to leak bids. In the case of at least two fast bidders, it generally has an incentive to leak bids. Practically, we could imagine different ways of seller bid leakage in ePBS, e.g. the proposer could forward bids to a relay that integrates them into its bid stream.

However, if the seller can credibly commit to not leak information (e.g. through reputation, TEE commitments or delegation to an intermediary such as a possible “first-price sealed-bid relay”) it is optimal for him to do so.

We formalize the unraveling of open-bidding relays in Propositions 7 and 14, and the seller’s disclosure incentives (without and with commitment) in Propositions 8 and 13 and Theorem 9.

3.3.1 No Last Look

We begin with a deliberately relay-friendly idealization. The relay is allowed to run a fully converged English auction among its participants, but the openness of the auction is purely internal: after the relay selects its representative, the relay representative and the native-channel builders submit proposer-facing bids without observing each other's final bids. Thus the model removes any cross-channel last-look advantage. The purpose of this benchmark is to ask whether an English relay can be stable even in this favorable case. Proposition 7 shows that it cannot: in large markets, all builders joining the English relay is not an equilibrium.

Fix a first-stage participation profile. Let S be the set of builders who choose the relay, and let $A = [n] \setminus S$ be the set of builders who choose the native first-price channel. Write $s = |S|$ and $m = |A|$.

If $s \geq 2$, the relay runs an internal English knockout auction. Let Y_S be the highest value among relay participants and let Z_S be the second-highest value. The internal auction selects the builder with value Y_S and sets the internal English-auction price equal to Z_S . The relay winner then chooses a proposer-facing bid b_S . Since bidding below Z_S would not be consistent with the internal English outcome, while bidding above Y_S is weakly dominated, we restrict attention to bids satisfying

$$Z_S \leq b_S \leq Y_S.$$

If $s = 1$, the relay has no internal competition and therefore generates no English-auction information; the sole relay participant simply chooses a proposer-facing bid subject to the usual no-overbidding constraint.

Each builder $j \in A$ submits a native first-price bid b_j to the proposer. The proposer compares the relay bid b_S , when present, with the native bids $(b_j)_{j \in A}$ and accepts the highest bid. The key modeling restriction is that the relay representative chooses b_S without observing the native bids, and native builders bid without observing the internal relay state. Hence the English auction reveals information only within the relay and creates no last-look opportunity across channels.


To test stability of the all-English profile, let $U_E^n(F)$ be a builder's ex-ante payoff when all n builders choose the relay, and let $U_{\text{out}}^n(F)$ be the ex-ante payoff of a single builder who deviates to the native first-price channel while the other $n - 1$ builders remain in the relay. The all-English profile can be a first-stage equilibrium only if

$$U_E^n(F) \geq U_{\text{out}}^n(F).$$

The all-first-price profile is a Bayes–Nash equilibrium of the two-stage game: a unilateral deviation to the relay creates no internal relay competition and hence no informational advantage, so the deviation is payoff-equivalent to remaining in the native first-price channel. The all-English profile is less robust. For compactly supported distributions with non-flat upper tails, full participation in the English relay is unstable once the number of builders is sufficiently large.

► **Proposition 7** (Asymptotic instability of all-English). *There exists $N(F) < \infty$ such that, for every $n \geq N(F)$, the all-English profile is not a first-stage equilibrium.*

Therefore, even in the idealized no-last-look benchmark, the relay cannot rely on single-homing alone to make an English format stable against a native sealed first-price alternative. The next question is what happens when the relay or proposer can create a genuine last-look opportunity by disclosing early bids.



■ **Figure 1** Last look vs First-price auction game representation.

3.3.2 Last Look

We now move to the more realistic case in which a subset of builders is latency-advantaged. These builders can submit a final response after observing information released close to the auction deadline; in this sense, they have the technological ability to obtain *last look*. Last look is not automatic, however. It matters only if the seller, the relay, or the auction format reveals relevant bid information before the deadline.

We first consider the case in which the seller, interpreted here as the block proposer, cannot credibly commit ex-ante to a disclosure strategy. As highlighted in Footnote 3, this does not exclude the possibility that the seller chooses their strategy ex-ante, but it forces them to choose a (Nash) equilibrium strategy. Let $L \subseteq [n]$ denote the set of builders who are technologically able to respond to late information before the auction closes. A builder in L does not automatically have last look; it obtains last look only if relevant information is revealed before the deadline. In the model below, we write $S \subseteq L$ for the builders who may receive such disclosures.

The game has three stages. First, the builders in $[n] \setminus S$ submit bids to the seller. Second, after observing these bids, the seller chooses a verifiable message σ_i to send to each bidder $i \in S$, where σ_i consists of a collection of bid pairs $\{(j, b_j)\}$ corresponding to an arbitrary subset of $[n] \setminus S$. Third, the bidders in S bid simultaneously. The block is allocated to the highest bidder under a pay-as-bid rule, with ties broken lexicographically in favor of bidders in S .⁶

Figure 1 summarizes the three-stage game.

We say that an auction is *k-leakage-resistant* if, for every set S with $|S| = k$, in every sequential Bayes–Nash equilibrium the seller reveals no information in the second stage. We say that it is *k-leaking* if for every set S with $|S| = k$, and every sequential Bayes–Nash equilibrium, the seller leaks information with probability one.

► **Proposition 8.** *The first-price auction is 1-leakage-resistant.*

Proof. Let $S = \{i\}$ and let r denote the highest bid submitted in stage 1. If the seller reveals r , then bidder i can condition on it and, whenever $v_i \geq r$, win by bidding exactly r , since ties are broken in favor of S . Thus leakage yields revenue exactly r . If instead the

⁶ We could formulate a more general version of the game by adding a fourth stage in which the seller chooses which bid to accept. However, in any sequential equilibrium, the seller would choose the highest bid. Hence, without loss of generality, we restrict attention to the three-stage game.

seller reveals nothing, bidder i must choose his bid without observing the realized value of r . Hence, with positive probability, his equilibrium bid is strictly above r , and on that event the seller receives strictly more than r . Therefore the seller's expected revenue is strictly higher under no leakage, so in any sequential equilibrium the seller does not leak. ◀

With a unique latency-advantaged builder, the seller's incentive is compatible with no disclosure: the first-price allocation is preserved, builders are ex-ante symmetric, and the mechanism weakens incentives to compete purely on latency rather than on order-flow acquisition or block construction. However, with multiple latency-advantaged builders this is no longer true.

► **Theorem 9.** *The first-price auction is k -leaking for all $2 \leq k \leq n - 1$.*

Proof of the theorem. Let us fix a sequential Bayes–Nash equilibrium and write $s = |S|$. Let R be the random variable

$$R := \max\{b_j : j \in [n] \setminus S\},$$

that is, the highest bid among the bidders in the first stage. Let M denote the seller's message in the second stage, with $M = \emptyset$ denoting the event that the seller reveals no information at that stage. Define

$$\bar{r} = \text{ess sup}(R \mid M = \emptyset).$$

Observe that first-stage bids are bounded away from $\bar{v}$. Fix $a \in (\underline{v}, \bar{v})$, let $q = F(a)^{n-1}$, and choose $\eta > 0$ such that $(\bar{v} - \eta - a)q > \eta$. A bidder with value $v \geq \bar{v} - \eta$ can bid a and win whenever all other values are below a , obtaining a payoff at least $(\bar{v} - \eta - a)q > \eta$. On the other hand, any bid $x > \bar{v} - \eta$ gives $v - x < \eta$. Hence no high type bids above $\bar{v} - \eta$, while lower types do not do so by no-overbidding. Therefore $R \leq \bar{v} - \eta$ almost surely, and so $\bar{r} < \bar{v}$.

► **Lemma 10.** *Fix the continuation game among the bidders in S after stage 2. Consider the two cases in which the seller either reveals the realized highest first-stage bid r to all bidders in S , or reveals nothing. In both cases, every equilibrium outcome is outcome-equivalent to a symmetric Bayes–Nash equilibrium. Moreover, every symmetric mixed Bayes–Nash equilibrium is outcome-equivalent to a pure symmetric Bayes–Nash equilibrium.*

Proof. Conditional on either public message, the bidders in S are ex-ante identical, their values are still i.i.d. according to F , and the continuation game is a continuous first-price auction game with a stochastic or static reserve price. Hence the symmetrization argument in [5] applies, so we may restrict attention to symmetric equilibria. The same result also implies that any symmetric mixed equilibrium is outcome-equivalent to a pure symmetric equilibrium. Since the seller's continuation payoff depends only on the induced distribution of bids and payments, this replacement is without loss of generality. ◀

By Lemma 10, we may assume without loss of generality that the continuation bidding strategies are symmetric and pure.

We compare the continuation game after no disclosure with the continuation game in which the seller reveals the realized highest bid to all bidders in S . Let the corresponding symmetric pure bidding strategies be $b_\emptyset$ and b_r .

When the seller reveals $M = r$ to all bidders in S , the continuation game is the standard first-price auction with reserve r . Its symmetric equilibrium is unique, and each player has bidding strategy

$$b_r(v) = \begin{cases} v - \frac{\int_r^v F(t)^{s-1} dt}{F(v)^{s-1}}, & v \geq r, \\ 0, & v < r. \end{cases}$$

19:16 Competing Auctions in Intermediated Markets

Now consider the continuation game after no disclosure. Let

$$H(x) = \Pr[R \leq x \mid M = \emptyset]$$

be the cdf of R conditional on $M = \emptyset$ and $b_\emptyset$ the bid function of the symmetric Bayes-Nash equilibrium with $M = \emptyset$. Define

$$v^* = \inf\{v : b_\emptyset(v) \geq \bar{r}\}.$$

A first observation is that $v^* > \bar{r}$. Indeed, in a first-price auction bidding above one's value is weakly dominated, so $b_\emptyset(v) \leq v$ for every v , which already implies $v^* \geq \bar{r}$. If $v^* = \bar{r}$, then by continuity we would have $b_\emptyset(\bar{r}) = \bar{r}$. But then a bidder of type $\bar{r}$ obtains payoff 0 from bidding $\bar{r}$, while by shading slightly below $\bar{r}$ she still beats R with positive probability under $M = \emptyset$ and wins against the other bidders in S with positive probability, yielding strictly positive payoff. This contradicts optimality, so $v^* > \bar{r}$.

Case 1: $v^* < +\infty$. For all $v \geq v^*$, equilibrium bids satisfy $b_\emptyset(v) \geq \bar{r}$, so such bids beat the first-stage maximum whenever $M = \emptyset$. Hence, if a bidder of type v deviates to the bid of a type $z \geq v^*$, her payoff is

$$(v - b_\emptyset(z))F(z)^{s-1}.$$

Optimality at $z = v$ gives

$$\left. \frac{d}{dz} \left((v - b_\emptyset(z))F(z)^{s-1} \right) \right|_{z=v} = 0,$$

that is,

$$-b'_\emptyset(v)F(v)^{s-1} + (v - b_\emptyset(v))(s-1)F(v)^{s-2}f(v) = 0.$$

Equivalently,

$$\frac{d}{dv} \left((v - b_\emptyset(v))F(v)^{s-1} \right) = F(v)^{s-1}.$$

Using the boundary condition $b_\emptyset(v^*) = \bar{r}$, we obtain

$$b_\emptyset(v) = v - \frac{(v^* - \bar{r})F(v^*)^{s-1} + \int_{v^*}^v F(t)^{s-1} dt}{F(v)^{s-1}}.$$

We now compare the two bid functions. For every $v \geq v^*$,

$$b_{\bar{r}}(v) - b_\emptyset(v) = \frac{(v^* - \bar{r})F(v^*)^{s-1} - \int_{\bar{r}}^{v^*} F(t)^{s-1} dt}{F(v)^{s-1}}.$$

Since F is strictly increasing, we have $F(t)^{s-1} < F(v^*)^{s-1}$ for all $t \in [\bar{r}, v^*)$, and therefore

$$\int_{\bar{r}}^{v^*} F(t)^{s-1} dt < (v^* - \bar{r})F(v^*)^{s-1}.$$

Hence

$$b_{\bar{r}}(v) > b_\emptyset(v) \quad \text{for all } v \geq v^*.$$

Also, if $v \in (\bar{r}, v^*)$, then by definition $b_\emptyset(v) < \bar{r}$, whereas

$$b_{\bar{r}}(v) = v - \frac{\int_{\bar{r}}^v F(t)^{s-1} dt}{F(v)^{s-1}} > \bar{r}.$$

Thus

$$b_{\bar{r}}(v) > b_\emptyset(v) \quad \text{for all } v > \bar{r}.$$

Case 2: $v^* = +\infty$. In this case, for all $v \in [\underline{v}, \bar{v}]$, $b_\emptyset(v) < \bar{r}$. However, for $v \in [\bar{r}, \bar{v}]$, $b_{\bar{r}}(v) \geq \bar{r}$, and therefore $b_{\bar{r}}(v) > b_\emptyset(v)$.

Let v_i denote the valuation of bidder $i \in S$. The seller's continuation revenue after revealing r is

$$\Pi_{\text{rev}}(r) = \mathbb{E} \left[\max \{ r, \max_{i \in S} b_r(v_i) \} \right],$$

whereas after no disclosure it is

$$\Pi_\emptyset(r) = \mathbb{E} \left[\max \{ r, \max_{i \in S} b_\emptyset(v_i) \} \right].$$

From the pointwise comparison above, for every valuation profile $(v_i)_{i \in S}$,

$$\max \{ \bar{r}, \max_{i \in S} b_{\bar{r}}(v_i) \} \geq \max \{ \bar{r}, \max_{i \in S} b_\emptyset(v_i) \},$$

and the inequality is strict whenever $\max_{i \in S} v_i > \bar{r}$. Since $\bar{r} < \bar{v}$,

$$\Pr \left[\max_{i \in S} v_i > \bar{r} \right] = 1 - F(\bar{r})^s > 0.$$

Therefore

$$\Pi_{\text{rev}}(\bar{r}) > \Pi_\emptyset(\bar{r}).$$

The seller's continuation revenue is continuous in r in both cases. Hence there exists $\delta > 0$ such that

$$\Pi_{\text{rev}}(r) > \Pi_\emptyset(r) \quad \text{for all } r \in [\bar{r} - \delta, \bar{r}].$$

By definition of $\bar{r}$ as the essential supremum of R conditional on $M = \emptyset$,

$$\Pr[R \in [\bar{r} - \delta, \bar{r}] \mid M = \emptyset] > 0.$$

Hence, on a positive-probability set of histories at which the equilibrium prescribes no disclosure, the seller would strictly prefer to reveal the realized highest bid. This contradicts sequential optimality. Therefore the seller cannot choose $M = \emptyset$ with positive probability. ◀

The previous theorem shows that the allocation induced by the three-stage game is not equivalent to the allocation in the standard first-price auction. Since the first-price auction is welfare-optimal with symmetric i.i.d. bidders, the three-stage game is not welfare-optimal. Moreover, when value distributions are regular, the first-price auction is revenue-optimal within the class of auctions without a reserve price (see Proposition 13). Thus allowing the seller to leak information cannot improve on the no-disclosure benchmark; the theorem implies that revenue is strictly lower. Also, information disclosure can have an adverse effect on competition in the block-building market by increasing market concentration. As we show in the following proposition, disclosure raises the payoff of agents with an informational advantage, thereby strengthening their position relative to less-informed agents.

► **Proposition 11** (Last-look rents). *Suppose the seller reveals the highest first-stage bid to every bidder in S . Then, in any symmetric equilibrium, every bidder in S obtains a strictly larger ex-ante utility and probability of winning than every bidder in $[n] \setminus S$.*

Proof. As in Lemma 10, the symmetrization argument in [5] applies to the bidders in $A := [n] \setminus S$. Let $m := |A|$ and $s := |S|$, so that $m + s = n$. Therefore, we can reduce the analysis to the case in which bidders in A use a monotone bidding strategy b_A and bidders in S use a monotone bidding strategy b_S . Let $X_A(v)$ and $X_S(v)$ denote the interim winning probabilities of a bidder with value v in groups A and S , respectively.

Almost surely, every player bids strictly below her valuation. In particular, $b_A(v) < v$. A bidder in A with value v wins only if the other bidders in A have values below v and the bidders in S have values below $b_A(v)$. Hence

$$X_A(v) = F(v)^{m-1} F(b_A(v))^s \leq F(v)^{m-1} F(v)^s = F(v)^{n-1},$$

with strict inequality on a positive-probability set.

By contrast, a bidder in S with value v wins whenever the other bidders in S have values below v and the highest first-stage bid is at most v . Therefore

$$X_S(v) = F(v)^{s-1} \Pr \left[\max_{j \in A} b_A(v_j) \leq v \right].$$

Since $b_A(t) \leq t$, the event that all bidders in A have values below v implies $\max_{j \in A} b_A(v_j) \leq v$. Hence

$$X_S(v) \geq F(v)^{s-1} F(v)^m = F(v)^{n-1}.$$

Thus $X_S(v) > X_A(v)$ on a positive-measure set. By the envelope theorem,

$$u_S(v) = \int_{\underline{v}}^v X_S(t) dt \geq \int_{\underline{v}}^v X_A(t) dt = u_A(v),$$

with strict inequality on a positive-measure set. Taking expectations over $v \sim F$ gives the result. $\blacktriangleleft$

3.3.3 Last-Look with Seller Commitment

The previous section studied the case in which the seller could not commit to a disclosure policy. In that environment, disclosure is chosen history by history after early bids have arrived. We showed that this creates an ex-post incentive to leak information when there are at least two latency-advantaged bidders: after observing the early bids, the seller can raise continuation revenue by revealing information to the bidders who are still able to respond.

We now ask a different question. Suppose instead that the seller can commit ex-ante to a particular auction rule, or equivalently can choose among relays that commit to different rules (and make it credible that they will ignore or not-leak direct bids sent to them). One relay may run a sealed first-price auction, another may run an open English-style auction, and another may commit to giving last look to a specified set of builders. In this sense, the seller's decision is over which auction format they prefer to pick ex-ante and commit to.

This distinction matters. Ex-post, after early bids have been observed, leaking information can be profitable. Ex-ante, however, the proposer must take into account that last look changes bidders' incentives and shifts surplus toward the latency-advantaged builders. The result below shows that, when values are i.i.d. from a regular distribution, this ex-ante effect dominates: among committed last-look rules, the proposer weakly prefers the sealed first-price auction.

We assume that the seller can choose among the following auction formats. For a set $S \subseteq [n]$, define the auction that reveals bids from $[n] \setminus S$ to bidders in S as follows:

1. The set of builders $A := [n] \setminus S$ submits bids simultaneously.
2. The bids submitted by A are revealed to the builders in S .
3. The builders in S then submit bids simultaneously.
4. The right to decide the content of the next block is allocated to the highest bidder under a pay-as-you-bid payment rule.

Subsequently, we use the short-hand notation $\text{Show}(S)$ to denote the auction that reveals bids from $[n] \setminus S$ to bidders in S . Observe that the cases $S = \emptyset$ and $S = [n]$ are equivalent to the first-price auction. For the set of latency-advantaged parties $L \subseteq [n]$, one can view the case $S = L$ as the current status quo, i.e. an open first-price auction with last-look advantage for a set of parties.

First, we assume that the block proposer can *commit* to a specific auction format; that is, once they announce to whom they reveal bids, they cannot deviate from it during its execution.

► **Definition 12.** *A distribution F is regular if its virtual valuation function*

$$\phi(v) = v - \frac{1 - F(v)}{f(v)}$$

is non-decreasing in v , where f is the density of F .

► **Proposition 13.** *If F is regular, then for any proper set $S \subseteq [n]$ and any Bayes–Nash equilibrium of $\text{Show}(S)$,*

$$\text{Rev}(\text{FP}) \geq \text{Rev}(\text{Show}(S)).$$

In other words, under commitment the block proposer weakly prefers the sealed first-price auction.

Suppose, finally, that the proposer commits not to leak bids submitted through the sealed first-price channel, but still accepts bids coming from an open-bid relay. In this case, a high-latency builder has no reason to send her bid through the open relay. If she sends a bid b through the relay, low-latency builders may observe it and still have time to react. If she sends the same bid b through the sealed channel, the proposer compares it in the same way, but the bid is not revealed before the deadline. Hence, whenever the builder has a chance of winning, submitting through the relay can only reduce that chance of winning.

► **Proposition 14** (Open relays under no leakage). *Suppose the proposer commits not to leak bids submitted through the sealed first-price channel, while still accepting bids from an open-bid relay. Then high-latency builders bid through the sealed channel. Low-latency builders may still use or observe the open relay, but the relay no longer gives them a last-look advantage over sealed bids. The outcome is equivalent to the standard sealed first-price auction.*

4 Conclusion

We have analyzed competing auctions in intermediated markets where a single seller can route a single sale through multiple parallel mechanisms, with proposer-builder separation as the leading application. We found that second-price intermediary auctions unravel fully against a sealed first-price principal auction (Theorem 1 under single-homing, Theorem 4 under multi-plexing), open-bidding intermediaries unravel partially – collapsing into first-price

under symmetric latency, sorting fast bidders to the intermediary under asymmetric latency (Propositions 7 and 14) – and Sybil bidding can restore the second-price intermediary by allowing bidders to convert it into a first-price-equivalent mechanism. On the disclosure side, a first-price auction is leakage-resistant against a single privileged bidder (Proposition 8) but leaks in equilibrium against two or more (Theorem 9), and the optimal commitment under regularity is to share no information at all (Proposition 13).

The institutional reading of these results is that ePBS will likely reshape the relay market: An in-protocol sealed-bid first-price channel limits the design space for relays going forward. Relays running second-price or open-bidding mechanisms cannot sustain participation against the sealed first-price alternative. However, this prediction relies on the in-protocol channel being credibly sealed; the assumption of a non-leaking in-protocol bidding channel is itself fragile: the disclosure analysis identifies that a proposer receiving direct bids has an incentive to leak bids in several scenarios. Thus, proposers forwarding bids to a relay’s open-bid feed is a possible outcome. On the other hand, there is value for proposers from a credible commitment not to leak bids. This could be created through reputation mechanisms or solutions such as TEEs, in the case of stakers large enough to provide individual commitments (Lido and similar institutional stakers), or through intermediation in the form of “sealed-bid” relays. Thus, there are several factors pushing relays to effectively move to sealed-bid first-price auction formats post ePBS. Still, there is a potential role for them as credible intermediaries and for services such as fast bid and block propagation as well as non-collateralized bidding. There is also, in principle, the narrower path to maintain an open-bid relaying service to which proposers forward bids. However, the viability of this latter path depends crucially on the nature of competition in the builder market.

Several limitations of our analysis are worth flagging. We do not model collusion among builders, which is a standard concern for sealed-bid auctions in classical settings but plausibly less binding in blockchain environments where on-chain transparency and programmable coordination weaken the open-vs-sealed distinction as a collusion-suppression tool. We do not model the searcher-builder pipeline upstream of the auction; our values are private at the builder level, which is the relevant abstraction for the auction itself but abstracts from common-value dynamics among searchers. We treat each slot in isolation rather than as part of a repeated game; reputational mechanisms that sustain relay commitment are mentioned but not formally modeled. The analysis also takes the set of available mechanisms and intermediaries as given; the entry and design problem for relays – what mechanism a new relay should propose, given the equilibrium predictions above – is a natural follow-up question.

References

- 1 Mark Armstrong. Competition in two-sided markets. *The RAND journal of economics*, 37(3):668–691, 2006.
- 2 Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. *Journal of economic theory*, 137(1):580–609, 2007. doi:10.1016/J.JET.2007.02.001.
- 3 Roberto Burguet and József Sákovics. Imperfect competition in auction designs. *International Economic Review*, 40(1):231–247, 1999.
- 4 Bernard Caillaud and Bruno Jullien. Chicken & egg: Competition among intermediation service providers. *RAND journal of Economics*, pages 309–328, 2003.
- 5 Shuchi Chawla and Jason D Hartline. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 181–196, 2013. doi:10.1145/2492002.2483188.

- 6 Francesco D’Amato, Nico Flaig, Barnabé Monnot, Michael Neuder, Potuz, Justin Traglia, and Terence Tsao. EIP-7732: Enshrined Proposer-Builder Separation. <https://eips.ethereum.org/EIPS/eip-7732>, 2024. Draft EIP, accessed April 2026.
- 7 Glenn Ellison, Drew Fudenberg, and Markus Möbius. Competing auctions. *Journal of the European Economic Association*, 2(1):30–66, 2004.
- 8 Thierry Foucault, Marco Pagano, and Ailsa Röell. *Market liquidity: theory, evidence, and policy*. Oxford University Press, 2023.
- 9 Bruno Mazonra, Burak Öz, Christoph Schlegel, and Fei Wu. The free option problem of epbs. *arXiv preprint arXiv:2509.24849*, 2025. doi:10.48550/arXiv.2509.24849.
- 10 R Preston McAfee. Mechanism design by competing sellers. *Econometrica: Journal of the econometric society*, pages 1281–1312, 1993.
- 11 Paul Milgrom and Robert J Weber. The value of information in a sealed-bid auction. *Journal of Mathematical Economics*, 10(1):105–114, 1982.
- 12 Maureen O’Hara and Mao Ye. Is market fragmentation harming market quality? *Journal of Financial Economics*, 100(3):459–474, 2011.
- 13 Roel Oomen. Last look. *Quantitative Finance*, 17(7):1057–1070, 2017.
- 14 Mallesh Pai and Max Resnick. Structural advantages for integrated builders in mev-boost. In *International Conference on Financial Cryptography and Data Security*, pages 128–132. Springer, 2024. doi:10.1007/978-3-031-69231-4_9.
- 15 Michael Peters and Sergei Severinov. Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory*, 75(1):141–179, 1997.
- 16 Jean-Charles Rochet and Jean Tirole. Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667, 2006.