CSEF

Centre for Studies in Economics and Finance

WORKING PAPER NO. 720

Adaptive Reserve Prices in Repeated Auctions

Federica Carannante, Marco Pagnozzi and Elia Sartori

June 2024

University of Naples Federico II **University of Salerno** **Bocconi University, Milan**

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS AND STATISTICS – UNIVERSITY OF NAPLES FEDERICO II
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 – e-mail: csef@unina.it
ISSN: 2240-9696

CSEF
Centre for Studies in Economics and Finance

WORKING PAPER NO. 720

Adaptive Reserve Prices in Repeated Auctions

Federica Carannante[*], **Marco Pagnozzi**[†], **and Elia Sartori**[‡]

Abstract

We analyze how the seller adjusts the reserve price in infinitely repeated auctions using the information conveyed by past bids. Bidders are myopic and have constant valuations; losers are replaced by new bidders, and winners leave with an exogenous probability. Our model is a stylized description of the market for online display advertisements, where publishers sell impressions through real-time first- or second-price auctions. The optimal reserve price is either equal to the value of the last winner, or lower than it when the winner's value is sufficiently high. In this second case, the reserve price decreases in the winner's value in a first-price auction, while it is independent of it in a second-price auction. Because past winners who are outbid substitute for the reserve price in a second-price auction, the seller often sets a lower reserve price and obtains a higher revenue than in a first-price auction. Long-run trade may be non-monotonic in the probability that winners leave.

Acknowledgments: We would like to thank Paul Klemperer, Johannes Schneider, Jacopo Perego and audiences at Universitat Pompeu Fabra, University of Bolzano, the 4th UniBg Economics Winter Symposium, the 5th UniBg IO Winter Symposium, EARIE 2023 in Rome, ESEM 2023 in Barcelona, Economic Theories of the Information Society workshop in Prague (CERGE-EI) and 12th EIEF-IGIER-UNIBO Workshop on Industrial Organization for numerous comments and suggestions.

[*] Princeton University. Email: fcarannante@princeton.edu

[†] Università di Napoli Federico II and CSEF. Email: pagnozzi@unina.it

[‡] CSEF and Università di Napoli Federico II. Email: elia.sartori@unina.it

1 Introduction

In repeated auctions for similar objects, bids provide information about valuations, which can be used by sellers to adjust the reserve price. For example, in the market for online display advertising, advertisers repeatedly bid in real-time auctions to acquire multiple impressions to the same user. In this market, reserve price adjustments are extensively used,[1] and recent evidence shows that they substantially increase the seller's revenue (Ostrovsky and Schwarz 2023; Choi and Mela 2023).

This paper analyzes a stylized model of the market of online display advertising. We consider a seller who runs an infinite sequence of auctions, each for a single identical object and with a fixed number of bidders. The seller observes all bids in past auctions, and adjusts the reserve price in each period exploiting the information conveyed by past bids. Bidders are drawn from a pool of ex-ante identical potential buyers, have a constant private valuation for the objects on sale and a stochastic capacity, representing the maximum number of objects that they are willing to acquire.[2]

We characterize the optimal reserve price and its implications for the auctions dynamics, long run trade and seller's revenue. We compare the two most commonly-used auction formats in the display advertising market: the first-price auction (FPA) and the second-price auction (SPA). Consistent with the actual functioning of the market, we assume that the auction format is fixed in all auctions, and that the only instrument available to the seller is the reserve price, that the seller chooses in every auctions.

In order to focus on the seller's dynamic pricing strategy, we make the simplifying assumption that bidders are myopic: in each auction all bidders play the equilibrium bid of a static auction with a fixed number of symmetric bidders, given the reserve price chosen by the seller. As we are going to discuss in Section 2.1, this assumption implies that bidders neglect both that their bids reveal information that the seller may use in future auctions and that the reserve price chosen by the seller may reflect information about a competitor's value.[3] We believe that the assumption of myopic bidding is a reasonable approximation of the actual behavior of many bidders in auctions for display advertising, because of the use of automated bidding. The algorithms that select optimal bids are often trained in simulated static auctions where variations in reserve prices predominantly reflect sellers' heterogeneity — e.g., different sellers' outside options — or information about a common valuation of the object on sale.[4]

The Market Display advertising consists of visual ads, such as banners or videos, that appear on websites, mobile apps and social media platforms. The US market for online display advertising is estimated to be worth \$143 billion in 2022, with more than 13 billion online advertisements sold every day.[5] The online advertising system allows advertisers to deliver advertisements targeted at the individual level, through impressions. The process begins when a user visits a webpage, triggering an auction to allocate the ad space. An impression is recorded once the ad space is filled and the advertisement becomes visible to that specific user.

[1] Sellers use intermediaries that typically have programs designed to increase their revenue by setting appropriate reserve prices — see, e.g., https://support.google.com/admanager/answer/11385824?hl=en.

[2] This can be interpreted either as a true capacity constraint (e.g., arising from a fixed budget) or as the event that the advertiser achieves his final objective of selling an item to the user. What matters in our model is that the winner of an auction may not participate in the subsequent ones, and that this decision is unknown to the seller.

[3] Pagnozzi and Sartori (2023) examine the effects of bidder sophistication in a framework analogous to the one presented in this paper: a sequence of two second-price sealed-bid auctions. They show that, in this setting, bidders shade their bid, strategically concealing their valuation in the first auction to prevent price discrimination in the second one. See also Caillaud and Mezzetti (2004)

[4] Wang, Zhang, Yuan, et al. (2017) discuss the computational challenges in online advertising, from both the advertiser (selection of bidding strategies in real-time) and the publisher (dynamic pricing and ad fraud detection) perspectives.

[5] For a more detailed description of the market see McAfee (2011).

In this market, sellers are publishers that own digital ad spaces on a website. Publishers sell impressions (with identical or very similar characteristics) through an intermediary, called Ad Exchange. The Ad Exchange sets reserve prices, collects bids from potential advertisers, and implements allocations and payments according to the auction rules.[6] Both first-price and second-price sealed-bid auctions have been used in this market.[7] Bidders are advertisers who wish to buy impressions; they typically bid through another intermediary, called Demand Side Platform (or DSP). The role of the DSP is twofold: (i) it allows advertisers to manage large and complex advertising campaigns, that typically require to spend a fixed budget to purchase a variety of impressions with certain characteristics over a specific time interval; and (ii) it implements automatic, real-time, bidding across a large number of auctions. There is also a large number of small, and often unsophisticated, advertisers that bid for impressions using much simpler and less customizable buying tools, called Advertiser Ad Networks. The vast majority of transactions in this market (around 90%) use some type of automated technology.

Our model is a stylized version of this market that only includes two types of agents: a single seller and many, ex-ante identical, bidders. Our seller can be interpreted as an Ad Exchange that runs repeated auctions for impressions shown to the same user in different moments. Bidders should be interpreted as different Demand Side Platforms and Advertiser Ad Networks bidding on behalf of an advertiser that wants to show repeatedly its ad to the user.

Preview of Results In order to isolate the key forces at play, we first focus on a static auction where the seller receives exogenous information about the valuation of one of the bidders — which we call the incumbent — who participates in the auction with exogenous probability.[8] We show that, in choosing the reserve price, the seller faces a trade-off between exploiting the possible presence of the incumbent and targeting new bidders. In both the FPA and the SPA, the optimal reserve price has the following features (Proposition 4). If the incumbent's value is low, the reserve price *excludes the incumbent* — i.e., it is higher than his value; if the incumbent's value is intermediate, it *tracks the incumbent* — i.e., it is equal to his value; if the incumbent's value is high, it *tails the incumbent* — i.e., it is lower than his value.

The intuition for this pattern of the optimal reserve price is the following. Tracking the incumbent guarantees that the seller extracts the whole bidders' surplus in case the incumbent participates in the auction and wins, but may result in a reserve price that is too far from the optimal one for new bidders. We show that tracking is optimal except in cases where the incumbent's value is too low (as the surplus that can be extracted is limited) or too high (as it would lead to an excessive risk of no trade).

The pricing strategy in the tailing region differs across auction formats over two dimensions. First, the reserve price in the SPA is lower than the one in the FPA for every value of the incumbent. The reason is that the bid submitted by the incumbent is a perfect substitute for the seller's reserve price when the incumbent loses in the SPA, but not in the FPA. Hence, the seller has a lower incentive to raise the reserve price in the SPA.

Second, the tailing reserve price does not depend on the incumbent's value in the SPA, but it is *decreasing* in

[6] Ad Exchanges typically have programs designed to increase publisher's revenue by setting appropriate reserve prices — see, e.g., https://support.google.com/admanager/answer/11385824?hl=en.

[7] Because of time constraints, it would be infeasible to run dynamic auctions, like the ascending auction, in this context. Moreover, each auction sells a single impression, since each impression is a unique event available at a different point in time.

[8] While the primary application of our model is the market of display advertising, the insights of the analysis of the static model hold relevance in a variety of contexts where a seller has information about one of the bidders, and chooses the reserve price in a one-shot auction. Such situations might arise, e.g., in procurement auctions, where the government may know the cost of the current provider, but not whether he is willing to maintain the relationship.

the incumbent's value in the FPA. The intuition for this surprising result is that, in the FPA, the marginal benefit of increasing the reserve price is decreasing in the incumbent's value, because bids of high-value bidders are less responsive to the reserve price than bids of low-value bidders;[9] while the cost of increasing the reserve price (which reflects the risk that no new bidder bids above the reserve price) is independent of the incumbent's value, because this cost materializes only if the incumbent does not participate. In the SPA, by contrast, both the cost and the benefit of increasing the reserve price are independent of the incumbent's value, because bids are independent of the reserve price.

The revenue ranking between FPA and SPA depends on the incumbent's value (Proposition 6). When the incumbent's value is low, the seller's revenue is higher in the FPA than in the SPA. The reason is that, when the seller either excludes or tracks the incumbent, the expected payment of myopic bidders is higher in the FPA. In this case myopic bidders overestimate competition in the auction (because they neglect that the incumbent never bids more than the reserve price), which induces them to bid more aggressively in the FPA, which yields higher revenues. When the incumbent's value is high, by contrast, the seller's revenue is higher in the SPA than in the FPA. The reason is, again, that the bid by a high-value incumbent is a substitute for the seller's reserve price in the SPA. Therefore, when the seller tails a high-value incumbent, in the SPA the incumbent's bid increases the expected seller's revenue if the incumbent stays but loses. Moreover, the seller's reserve price is closer to the optimal reserve price for new bidders and allows her to obtain a higher revenue when the incumbent leaves.[10]

We then move to the analysis of the dynamic environment, where the seller's information arises endogenously. The seller's dynamic problem is stated in a recursive form, with a state equal to the valuation of the winner of the previous auction, if there was one, or to "no trade." In addition to its static effects, the reserve price also has a dynamic effect: a higher reserve price implies a higher probability of transitioning to "no trade."

Theorem 7 shows that, despite this additional dynamic effect, the optimal reserve price in repeated auctions has the same qualitative features as in the static environment. In particular, the seller tracks incumbent with intermediate values and tails high-value ones. Moreover, in the tailing region the reserve is independent of the incumbent's value in the SPA, while it is decreasing in the incumbent's value in the FPA. However, in the repeated auctions, the seller never excludes an incumbent on path. The reason is that the seller excludes low-value bidders through the initial reserve price and, hence, she never encounters an incumbent she wants to exclude. Theorem 7 has immediate implications for the dynamics of the repeated auctions. For example, it implies that we can observe declining reserve prices along a trading history: as winners with higher values overcome each other, they induce lower reserve (Fig. 4.3). Moreover, it allows us to characterize the stationary distribution of winner values' (Section 5) and thus infer the properties of the repeated auction.

The revenue ranking between SPA and FPA depends on bidder's capacity (Proposition 10): if capacity is high then the FPA yields higher dynamic value, otherwise the SPA dominates. This ranking is understood in light of the static ranking. Since high capacity induces the seller to track all incumbents (as they will likely participate to the subsequent auction), the FPA yields higher expected revenue period by period. With tailing, instead, the static revenue is higher in the SPA when the incumbent has a relatively high valuation. Since the stationary distribution of incumbents is skewed towards high valuations (as those incumbents represent winners of previous auctions), the SPA has a mechanical advantage that overcomes the overbidding, driven by myopia, in the FPA under tracking.

[9]Formally, the bidding function in the FPA is sub-modular in the bidder's value and the reserve price, reflecting substitutability between the two.

[10]Carannante, Pagnozzi, and Sartori (2024) further explore the difference in revenues across auction formats (including SPA and FPA) with a common reserve price, given the valuation of one of the bidders.

We also study trade — i.e., the long run probability that an auction results in a sale — and how trade varies with bidders' capacity. We show that trade is non-monotonic in bidders' capacity (Proposition 11), possibly decreasing when capacity is high. This is counterintuitive, since trade can fail only if the incumbent reaches capacity and leaves. However, higher capacity induces the seller to increase the reserve price, possibly offsetting the direct effect. The ranking across auction formats in frequency of trade is also ambiguous (Proposition 11). Since the reserve price in the tailing region is lower in the SPA than in the FPA, the SPA typically trades more often. However, when there is no tailing region (i.e., when capacity is high), the FPA trades more often, as this auction formats has a lower initial reserve.

Finally, we analyze the information that can be inferred by only observing the tenure of the current winner — i.e., the number of consecutive periods in which a bidder has won. A higher tenure implies a higher expected valuation of the winner and, in the SPA, a higher expected revenue for the seller and a higher probability of future trade. In the FPA, by contrast, higher tenure might predict lower seller's revenue (Proposition 13). The reason for this counterintuitive result is that the valuation of a tenured winner affects its transfer both directly and through the reserve price, and higher incumbents are tailed with lower reserve in the FPA.

Related Literature Myerson (1981) and Riley and Samuelson (1981) show that the revenue-maximizing (optimal) static auction requires a reserve price. A vast literature further explores the role of reserve prices in static auctions, accounting for risk-averse bidders (Hu, Matthews, and Zou 2010), common values (Crémer and McLean 1988; McAfee, McMillan, and Reny 1989; Levin and Smith 1996), endogenous entry (McAfee and Vincent 1993; Levin and Smith 1994; Peters and Severinov 1997), taste projection (Gagnon-Bartsch, Pagnozzi, and Rosato 2021), and loss-averse bidders (Balzer, Rosato, and von Wangenheim 2022). We contribute to this literature by analyzing how the optimal reserve price in static FPA and SPA depends on the seller's information.

Optimal reserve prices in sequential auctions are studied in Caillaud and Mezzetti (2004) and Pagnozzi and Sartori (2023), who analyze two sequential second price auctions where the seller observes transaction prices and sets reserve prices in each period, accounting for bidder's optimal response.[11] Hummel 2018 analyzes an environment where the seller attempts to learn the distribution of bidders' valuations to inform her choice of reserve prices in the future.

The literature on the online advertising market includes both "display" and "sponsored-search" advertising. Sponsored-search advertising is based on the generalized second-price auctions, where advertisers bid for specific positions of the search engine results when specific keywords are queried (Aggarwal, Goel, and Motwani (2006), Edelman, Ostrovsky, and Schwarz (2007), Varian (2007) and Athey and Ellison (2011)). In this market, as in the one for display ad, advertisers typically delegate their bidding to intermediaries (Decarolis, Goldmanis, and Penta (2020), Decarolis and Rovigatti (2021), Decarolis, Goldmanis, Penta, and Shakhgildyan (2023)). Moreover, Ostrovsky and Schwarz (2023) conduct a large-scale field experiment and show the relevance of reserve prices in sponsored search auctions. Their evidence is suggestive that reserve price adjustments can substantially increase revenues in selling mechanism similar to those used in display ad auctions. For a review on how the auction mechanisms for online ad sales evolved over time, see Decarolis, Goldmanis, and Penta (2018).

Markets for display advertising have received less attention from researchers, as impressions were typically sold through guaranteed contracts. Recently, the display advertising market switched to real-time auction mechanisms.

[11] Ortega-Reichert (1968), Hausch (1988) and Luton and McAfee (1986) analyze bidders' strategies in sequential auctions. Several papers consider the "declining price anomaly" in sequential auctions — see, e.g., Black and De Meza (1992), McAfee and Vincent (1993) and Gale and Hausch (1994).

Arnosti, Beck, and Milgrom (2016) introduce the "modified second-bid auctions" as a mechanism that overcomes disadvantages for advertisers that cannot measure the value of individual impressions.[12] Choi and Mela (2023) analyze a one-shot second price auction, where the publisher sets the reserve price and bidders, upon observing the reserve, decide how much to bid, maximizing expected utility given the minimum impression level for their ad campaign. By contrast, we study seller's optimal strategy while accounting for the dynamic nature of this market. Finally, Kobayashi and Alcobendas (2023) focus on the problem of the DSP that has to select the optimal bid in each period, subject to a budget constraint that changes as a function of past bids: in each period, if the bidder wins, the remaining budget is reduced by the amount paid; if he doesn't win, the budget remains unchanged. The results suggest that dynamic incentives significantly increase bid shading in the FPA; nevertheless, by simulating a counterfactual scenario where the market uses SPA instead of FPA, the authors find that the FPA yields higher revenue and higher bidder surplus than the SPA. [13]

A different strand of literature, with applications to display-ad, lies at the intersection of mechanism and information design. Bonatti, Bergemann, and Wu (2023) compare selling mechanisms of digital platforms in the presence of off-platform markets and show that the mechanism affects prices both on and off the platform. Bergemann, Heumann, Morris, Sorokin, and Winter (2022) consider the optimal information structure in a SPA with no reserve under different objectives (seller's or bidders' surplus, efficiency). Bergemann, Heumann, and Morris (2023) focus on the information structure that maximizes bidders' surplus when the seller chooses the revenue-maximizing auction format. Their motivating example is the DSP choosing the information structure to attract advertisers (by revealing user characteristics that affect the advertisers' values) and the Ad-Exchange choosing the auction format. The common premise of this literature is that the seller has control over the accuracy of the bidders' information about their valuation (Bergemann and Pesendorfer (2007), Eső and Szentes (2007)) and discloses it optimally as part of the mechanism. By contrast, we conduct our analysis in the IPV setting, thus assuming that the seller has no control over bidders' information.[14]

An alternative way for the seller to exploit the information about a potential buyer is to use personalized pricing (Ali, Lewis, and Vasserman 2020, Bonatti and Cisternas 2020). This customization is made by publishers, that discriminate advertisers in bilateral negotiations. We focus, instead, on a setting where the seller uses a fixed auction mechanism and sets a unique reserve price for all bidders, consistent with the behavior of Ad Exchanges towards DSPs.

The rest of the paper is organized as follows. Section 2 describes the model. In Section 3 we analyze optimal reserve prices and the seller's revenue in a static environment with exogenous information. Section 4 characterizes the optimal reserve prices in the dynamic setting, describes typical paths of reserve prices and winners' values in repeated auctions, and compares the seller's revenue in the two auction formats. In Section 5, we characterize the stationary distribution of winners' values, and use it to analyze the long-run probability of trade. Finally, Section 6 studies how the tenure of the auction winner affects the probability of trade and the seller's expected revenue.

[12]These are advertisers that bid for an impression with the goal of enhancing brand recognition or promote a sale event, rather than to elicit an immediate response from the consumer, such as the click on a link.

[13]A related literature in Computer Science and Operation Research has explored the optimal algorithms that a seller can use to set reserve prices that maximize their profits (Chen, Berkhin, Anderson, and Devanur (2011); Buchbinder, Jain, and Naor (2007)). Amin, Rostamizadeh, and Syed (2014) and Golrezaei, Javanmard, and Mirrokni (2019) focus on algorithms that use historical bids to determine optimal reserve prices. For a comprehensive overview, Choi, Mela, Balseiro, and Leary (2020) provides a survey.

[14]Assuming IPV when the intermediary (DSP) has information about the objects sale (the impression) that she might strategically disclose (list characteristics) amounts to looking at the problem after the information is disclosed. The impression is interpreted as a list of individual characteristics of the user based on which the bidders form their willingness to pay.

Section 7 concludes. All proofs are in the Appendix.

2 Model

In each period $t \in \{0, 1, \ldots\}$, a (female) seller sells an object through an auction. The seller's value is normalized to 0 and she maximizes her total expected profits, discounted at rate β. We consider both first-price sealed-bid auctions (FPA) and second-price sealed-bid auctions (SPA).

Dynamics of Bidders' Values There is a constant number n of (male) bidders, indexed by i.[15] Bidder i, at time t, has valuation $\theta_{i,t}$ for the object. Valuations evolve as follows. At $t = 0$, valuations are drawn from an absolutely continuous distribution on $[0,1]$, with CDF $F(\cdot)$, pdf f, and increasing virtual value $\psi(x) := x - \frac{1-F(x)}{f(x)}$. For $t \geq 1$, valuations depend on the sequential allocation. If the object is not allocated to bidder i at time t, then $\theta_{i,t+1}$ is drawn from $F(\cdot)$. If the object is allocated to bidder i at time t, then $\theta_{i,t+1}$ is equal to $\theta_{i,t}$ with probability $1-\eta$, and is a new draw from $F(\cdot)$ with complementary probability η. All draws from $F(\cdot)$ are independent across bidders and over time (and of the event that a winner keeps his value).

Throughout the paper, and consistent with the microfoundation in the next section, we refer to the event that the winner draws a new value (resp. keeps his value) as "the winner leaves" (resp. "the winner stays"). The discount factor β, the number of bidders n, the persistence parameter η and the distribution of values $F(\cdot)$ are the primitives of our model and are common knowledge.

Period-t Auction At the beginning of each period t, the seller observes the history of bids $b^t := \{b_{i,s}\}_{i \leq n, s<t}$ and chooses a public reserve price R_t. The following assumption characterizes bidders' behavior.

Assumption 1. ***Myopic Bidding*** *In every period t, bidders submit the equilibrium bid of a static auction with n symmetric bidders and reserve price R_t.*

In other words, bidders bid $b_{i,t} = \theta_{i,t}$ in the SPA and $b_{i,t} = b(\theta_{i,t}, R_t, n)$ in the FPA, where

$$b(\theta, R, n) := \frac{1}{F(\theta)^{n-1}} \int_0^\theta \max\{y, R\} \, \mathrm{d}F(y)^{n-1}. \tag{2.1}$$

Given bids $b_t = \{b_{i,t}\}_{i \leq n}$, the allocation and transfers of period t are determined according to the rules of the auction (FPA or SPA). In period $t+1$ the bid history is updated to $b^{t+1} = \{b^t, b_t\}$ and bidders' values evolve as described above.

Recursive Structure The seller's problem can be written in recursive form with the state given by the value of the winner in the previous period (if such winner exists). This is because (i) since losers draw new values, only the value of the winning bidder in $t-1$ affects the optimal reserve price in period t and (ii) the seller can infer this value from the winning bid, as bids are strictly monotonic.

[15] We use the term "bidder" loosely as we do not model explicitly the identity of the agents who participate in the auction in a specific period. In Section 2.1, we show that the evolution of values we are about to describe results from the dynamic participation decision of myopic advertisers who choose among many (infinitely repeated) auctions. Therefore, bidder i at time t is a (potentially) different agent from bidder i at time $t+1$.

Hence, the information contained in b^t can be summarized by the (one-dimensional) state space $\Omega = \emptyset \cup [0,1]$. State $\emptyset$ indicates that in the previous period the auction had no winner, while state θ indicates that in the previous period the auction was won by a bidder with value θ. Given the state, the seller infers the composition of the current bidders' valuations as follows

1. State $\emptyset$: there are n bidders with values drawn i.i.d from F;

2. State θ: with probability η there are n bidders with values drawn i.i.d from F; with probability $(1-\eta)$ there are $(n-1)$ bidders with values drawn i.i.d from F and a bidder with value θ.

2.1 Interpretation

In this section we present an interpretation of our model, and discuss how this interpretation captures the main features of the market for online display advertising.

At each time t, the seller runs an auction with n slots to allocate an object.[16] Each bidder enters the repeated auctions (i.e., sits for the first time in a slot) with private value θ, representing his value for consuming the object at time t and in *any subsequent period*. Winners stay in the auction (i.e., keep sitting in their slot) until a satiation shock, which is geometrically distributed with parameter η. A winner that hits satiation, and all losers, leave the auction and are replaced by new bidders, whose values are drawn i.i.d from F. Prior setting the reserve price R_t, the seller observes all bids of the auction in $t-1$ and infers the value of the winner (if any). The seller, however, does not know whether the previous winner stayed in the current auction. At all histories, bidders play according to the equilibrium of the static auction with n symmetric bidders and reserve price R_t.

In our application to display advertising, the repeated auctions represent the sale of impressions displayed to a specific user u.[17] The seller is the Ad-Exchange in charge of allocating ad-spaces on the websites user u is browsing; bidders are advertisers who want to show their advertisement to u.

Dynamics of Values An advertiser enters the repeated auctions for user u with a value $\tilde{\theta}$, that represents the benefit he obtains in the event that user u clicks on his ad; η represents the probability that the user clicks on the ad, which is independent across time and advertisers.[18] Hence, the value for the advertiser from winning the auction is $\theta = \tilde{\theta}\eta$. When a click occurs, the value of displaying the ad to the same user in subsequent periods drops to zero, and the advertiser leaves. Hence, the value is constant until it expires. A bidder who loses the current auction leaves because, with probability $1-\eta$, the higher-value competitor who outbid him participates in the subsequent auction; a loser therefore prefers an alternative auction for an identical user (i.e. where he carries his value θ) without this higher-value competitor.[19]

[16] In the Online Appendix, we show that our results extend to a setting where the number of bidders evolves stochastically provided that, in each period, (i) the seller and all bidders know the realization of n and (ii) the realization of n conveys no information about whether the winner stayed or left.

[17] Our analysis focuses on a single repeated auction, though we rationalize the participation dynamics with the presence of many users with identical characteristics, who represent substitute repeated auctions and bidding constraints that limit the number of auctions in which each bidder can simultaneously participate.

[18] Advertisers are different only in the payoff they obtain *after* the user clicks on the ad. An appropriate reinterpretation of the satiation shock and the participation dynamic can allow for heterogeneity in the value from *just showing* the ad and in the click-through rate.

[19] This explains why losers leave when the object was allocated to a competitor, but it does not explain why bidders leave when no one has met the reserve price. In the Appendix we show that our qualitative results extend to a setting where losers leave only when the object is allocated, though the analysis is more complicated.

Bidding in the Static Auction Myopic bidding reflects the assumption that bidders are unaware that the reserve price set by the seller results from her dynamic optimization problem. Myopic bidders neglect both that (i) their bids reveal information that will be used to set future reserve prices (*forward myopia*) and that (ii) the current reserve price is informative about the valuation of a competitor (*backward myopia*).[20]

We think that myopic bidding is a reasonable practical assumption in auctions for display advertising, where the majority of bidders use automated bidding systems. If these systems are calibrated using simulated data from static auctions (e.g., for feasibility constraints), they disregard the flow of information across successive auctions and treat variations of the reserve price as exogenous, e.g. driven by heterogeneity in sellers' outside option. Indeed, in auctions for online display advertising reserve prices depend on (an are often equal to) the price that the seller would obtain through a direct sale of the impression — i.e., by allocating the impression to guaranteed contracts that publishers directly negotiate with advertisers to supply a given number of generic impressions with some predefined characteristics.

We also notice that myopic bidding imposes relatively mild restriction on the sophistication of bidders, as it is satisfied in an environment where bidders disregard the information contained in the reserve price only in the first period they participate (*one-shot myopia*). The reason is that winning an auction reveals (by monotone bidding) all the private information of the bidder, so that forward myopia is irrelevant after the first period. As for backward myopia, notice that incumbents best respond to the strategies of their (myopic) competitors. This hold vacuously for the SPA (see footnote 20); in the FPA it holds because myopic bidders best respond to the equilibrium strategy of $n-1$ unconditional bidders, which indeed they are facing. Hence, our dynamics is consistent with bidders that believe they are playing a static auction the first period they participate, but are rational thereafter.

2.2 Properties of b

In our analysis, we are going to use the function $b(\theta, R, n)$, defined in (2.1), extensively. Before proceeding we first establish some analytical properties of this function and how it can be used to determine seller's revenue in the FPA and the SPA.

Recall that (2.1) denotes the equilibrium bid of type θ in the FPA with n symmetric bidders and reserve price price R. By construction, it coincides with the seller's revenue in the FPA when a bidder with value θ wins the auction. By the Revenue Equivalence Theorem, $b(\theta, R, n)$ is also the expected payment of a bidder with value θ conditional on winning in a SPA with n total bidders (and hence $(n-1)$ competitors) and reserve price R.[21]

Moreover, since it is a dominant strategy for all bidders to bid their value in the SPA, when an incumbent with value θ participates in the auction the expected payment of a new bidder with value θ' can be written as $b(\theta', \max\{\theta, R\}, n-1)$. The reason is that the bid equal to θ by the incumbent is equivalent to a reserve price, so that the expected payment by a new bidder is the same as in a standard auction with only $(n-2)$ real competitors and a reserve price equal to $\max\{\theta, R\}$.

This also implies that $b(\theta', \max\{\theta, R\}, n-1)$ is the seller's expected revenue in the SPA when a new bidder with value θ' outbids the incumbent and wins the auction.

Fact 1. *$b(\theta, R, n)$ is increasing in θ, R and n.*

[20]Backward myopia is irrelevant in the SPA, because bidders have a dominant strategy. Therefore, our analysis of the SPA is equivalent to a setting with rational bidders who fully discount the future.

[21]The RHS of (2.1) clarifies this observation: a bidder with value $\theta \geq R$ bids indeed the expectation of the maximum between the reserve price and the highest value among his $(n-1)$ competitors, conditional on his competitors having values lower than θ. The latter is exactly the expected transfer of a winner θ in a SPA.

Fact 1 implies that, when there is an incumbent θ and the seller sets a reserve price equal to his value, myopic new bidders bid "more aggressively" — in the sense that they have a higher expected payment — in the FPA than in the SPA. The reason is the following. In the SPA, the expected payment of a new bidder with value θ' conditional on winning is $b(\theta', \theta, n-1)$. When there is an incumbent in the FPA, by contrast, myopic new bidders bid as in an auction with n symmetric bidders, even if in reality there are only $(n-1)$ new bidders. Therefore, the expected payment of a new bidder with value θ' conditional on winning in the FPA is $b(\theta', \theta, n)$, which is higher than in the SPA. As we are going to show, this effect of myopic bidding tends to increase revenue in the FPA.

Fact 2. $\frac{\partial^2}{\partial R \partial n} b(\theta, R, n) < 0$, $\frac{\partial^2}{\partial R \partial \theta} b(\theta, R, n) < 0$, *and* $\frac{\partial^2}{\partial \theta \partial n} b(\theta, R, n) > 0$.

Fact 2 implies that in the FPA bidders with higher values are less sensitive to changes in the reserve price: increasing the reserve price induces all bidders (with values higher than the reserve price) to bid more aggressively, but less so the higher is the value of the bidder.

3 Static Auction

In this section, we start our analysis by considering a static environment with exogenous information about one of the n bidders: before an auction, the seller learns that a bidder with value θ (the incumbent) is going to participate in the auction with probability $(1-\eta)$. We first characterize the optimal reserve price for the seller in Section 3.1, and then compare the seller's revenue in the FPA and SPA in Section 3.2.

Notice preliminary that a seller who has no information — which we refer to as information $\emptyset$ — faces a familiar problem. By the Revenue Equivalence Theorem, we can express her revenue in both the FPA and the SPA — as well as in any efficient auction — as a function of the reserve price R by integrating the expected payment of the bidder with the highest valuation:

$$\pi_n(R) := \int_R^\infty b(\theta, R, n) \mathrm{d}F(\theta)^n . \tag{3.1}$$

Maximizing this revenue yields the standard monopoly price condition (Myerson, 1981; Riley and Samuelson, 1981).

Fact 3. *Independently of n, η and of the auction format, the optimal reserve price in state $\emptyset$ is $R_\emptyset = r^M$, where r^M solves $\psi(r^M) = 0$.*

In the next section we characterize the optimal reserve price in the FPA and SPA given any possible θ, and then compare the expected seller's revenue in the two auction formats.

3.1 Static Reserve Prices with an Incumbent

When the seller has information θ, she knows that with probability $(1-\eta)$ a bidder with valuation θ will participate to the auction. The seller then chooses the reserve price R to maximize a weighted sum of her expected revenues in auction format $i = F, S$ when the incumbent participates and when he does not — i.e., she maximizes

$$\pi^i(\theta, R) := \eta \pi_n(R) + (1-\eta) \pi^i_{n-1,\theta}(R), \tag{3.2}$$

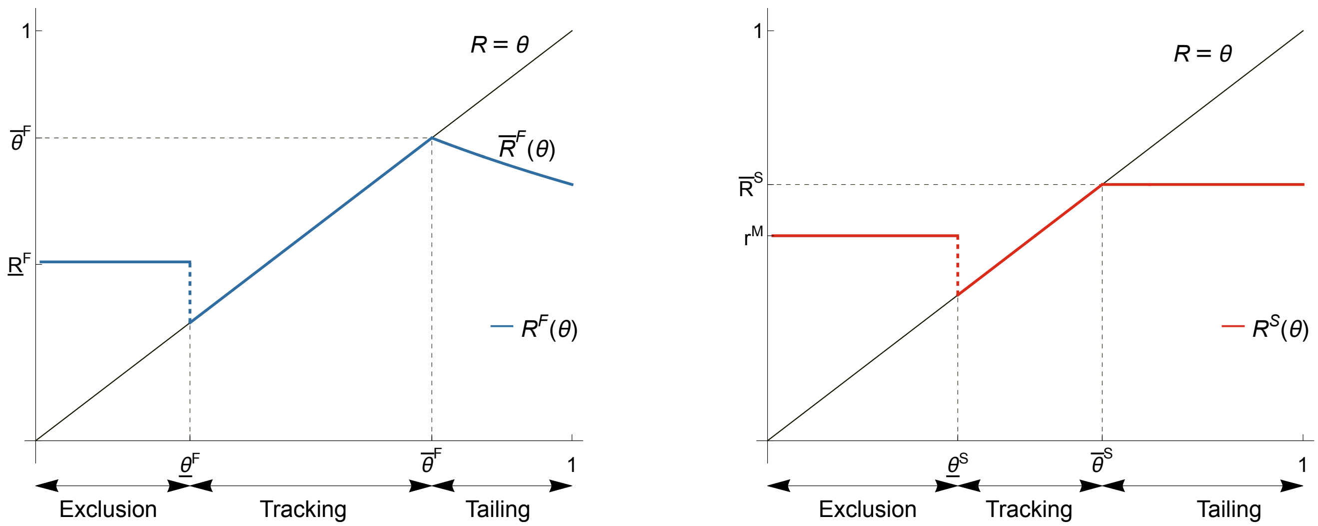

Figure 3.1: Optimal static reserve prices in FPA (left) and SPA (right) as a function of the seller's information, for $\theta \sim \mathcal{U}[0,1]$, $\eta = 0.5$, and $n = 4$.

where $\pi^i_{n-1,\theta}(R)$ represents the seller's revenue in the iPA, $i = F, S$, with reserve price R and an incumbent with value θ (as well as $n-1$ symmetric bidders).[22] The seller faces a trade off between setting the optimal reserve price for the standard symmetric bidders and setting a reserve price aimed at extracting surplus from the incumbent with value θ if he participates. The next proposition characterizes the optimal solution to this trade off.

Proposition 4. *The optimal reserve price in auction format $i = F, S$ with information θ is*

$$R^i(\theta) = \begin{cases} \underline{R}^i & if\, 0 \le \theta < \underline{\theta}^i & Exclusion \\ \theta & if\, \underline{\theta}^i \le \theta < \overline{\theta}^i & Tracking \\ \overline{R}^i(\theta) & if\, \theta \ge \overline{\theta}^i & Tailing \end{cases}$$

where $\underline{R}^i > \underline{\theta}^i$, $\underline{R}^S = r^M > \underline{R}^F$, and

- $\overline{R}^S(\theta) = \overline{\theta}^S$ *solves*

$$\psi(\overline{\theta}^S) = \frac{1-\eta}{n\eta f(\overline{\theta}^S)}, \tag{3.3}$$

- $\overline{R}^F(\theta)$ *solves*

$$\psi(\overline{R}^F(\theta)) = \frac{1-\eta}{n\eta f(\overline{R}^F(\theta))} - \frac{(1-\eta)(n-1)\log(F(\theta))}{n\eta f(\overline{R}^F(\theta))} \tag{3.4}$$

  and is strictly decreasing in θ and $\overline{R}^F(\overline{\theta}^F) = \overline{\theta}^F$.

Figure 3.1 displays reserve prices in the FPA and SPA for $\theta \sim \mathcal{U}[0,1]$.[23] First, if θ is relatively small, in both auction formats the seller excludes the incumbent from the auction by setting a reserve price higher than his value. We define this policy *exclusion*. Second, for intermediate values of θ (between $\underline{\theta}^i$ and $\overline{\theta}^i$), in both auction formats the seller sets a reserve price exactly equal to the incumbent's value. This policy, that we define *tracking*, allows the seller to capture the whole surplus generated by the incumbent, when he participates in the auction.

[22] See equations (7.6) and (7.7) in the Appendix for an analytical definition.
[23] All figures in the paper assume that θ is uniformly distributed.

Third, for high values of θ, the seller sets a reserve price lower than the incumbent's value and higher than r^M.[24] This policy that we define *tailing*, however, requires a qualitatively different reserve price in the two auction formats. The tailing reserve price is constant in the SPA: the seller sets the same reserve price for all high types that she does not want to track. By contrast, in the FPA the reserve price is decreasing in θ: the higher the value of the incumbent, the lower the reserve price set by the seller. Finally, the tailing reserve price is higher in the FPA than in the SPA, with the difference being decreasing in θ, and it is equal in the two auction formats when $\theta = 1$.[25]

The optimal reserve price is shaped by the following effects. Setting R above θ reduces the seller's revenue discretely by

$$(1 - \eta) F(\theta)^{n-1} \theta \tag{3.5}$$

i.e., by the valuation of the incumbent times the probability that he wins (because he participates and no other bidder exceeds his valuation). By contrast, if the incumbent does not participate, the seller's revenue is maximized at r^M. Hence, tracking an incumbent θ pushes the seller away from the static optimum: when $\theta > r^M$, the reserve price is too high for the standard bidders and the seller bears an excessive risk of no trade; when $\theta < r^M$, the reserve price is too low for standard bidders and the seller loses revenue from infra-marginal bidders.[26] Moreover, the cost of tracking increases as the incumbent's value is further away from r^M, therefore *tracking* is optimal for intermediate values of θ, while the optimal policy is different for high and low values of θ.

Exclusion is optimal if θ is low enough. The discrete revenue loss that arises when setting R above θ, expressed by (3.5), becomes negligible for θ low, so that the benefit of setting a reserve price that is optimal for the standard bidders eventually dominates. In this case, the seller sets the same reserve price regardless of the value of the incumbent that she wants to exclude, and this reserve price maximizes a convex combination of her revenues with n and $(n-1)$ standard bidders.

Finally, *tailing* is optimal if θ is very high. In this case, tracking the incumbent is too costly: if the incumbent does not participate, the reserve price is excessively higher than r^M. The seller still takes into account the possible presence of an incumbent, but only follows him at a distance by setting a reserve price lower than θ. In the SPA, $\overline{R}^S$ is independent of θ because, conditional on tailing, (i) the cost of marginally increasing the reserve price is independent of θ, as this cost arises only when the incumbent does not participate (and no other bidder meets the reserve price) and (ii) the benefit of marginally increasing the reserve price is also independent of θ, as it is equal to the probability that all standard bidders are lower than the reserve price, so that the reserve binds when the incumbent participates.

By contrast, in the FPA $\overline{R}^F(\theta)$ is decreasing in θ because, conditional on tailing, (i) the cost of marginally increasing the reserve price is independent of θ, as in the SPA, but (ii) the benefit does depend on θ, as it represents the effect of a higher reserve price on the winning bid (either by the incumbent, or by the standard bidder with a higher value) when the incumbent participates. By Fact 2 bids of higher-value bidders are less sensitive to the reserve price, so that the marginal benefit of increasing the reserve price in the FPA is decreasing in θ.

The fact that $\overline{R}^F(1) = \overline{R}^S(1)$, combined with the fact that the tailing reserve is decreasing in the FPA (and constant in the SPA) allows us to conclude that $\overline{\theta}^S < \overline{\theta}^F$: the FPA tracks some incumbents which the SPA tails.

[24] In the equations characterizing the reserve price in the tailing region (equations (3.3) and (3.4)) there is a positive wedge relative to the monopoly price condition.

[25] Formally, the difference between conditions (3.3) and (3.4) is proportional to $-\log F(\theta)$, which is positive, decreasing in θ, and converging to zero as $F(\theta) \to 1$. Hence, $\overline{R}^F(1) = \overline{R}^S(1) = \overline{\theta}^S$.

[26] Notice that the loss for $\theta < r^M$ is present even if the incumbent participates, while the no trading loss for $\theta > r^M$ materializes only if the incumbent does not participate. As a consequence, when $\eta = 0$, we have perfect tracking for high θ but not for low θ (see Proposition 5).

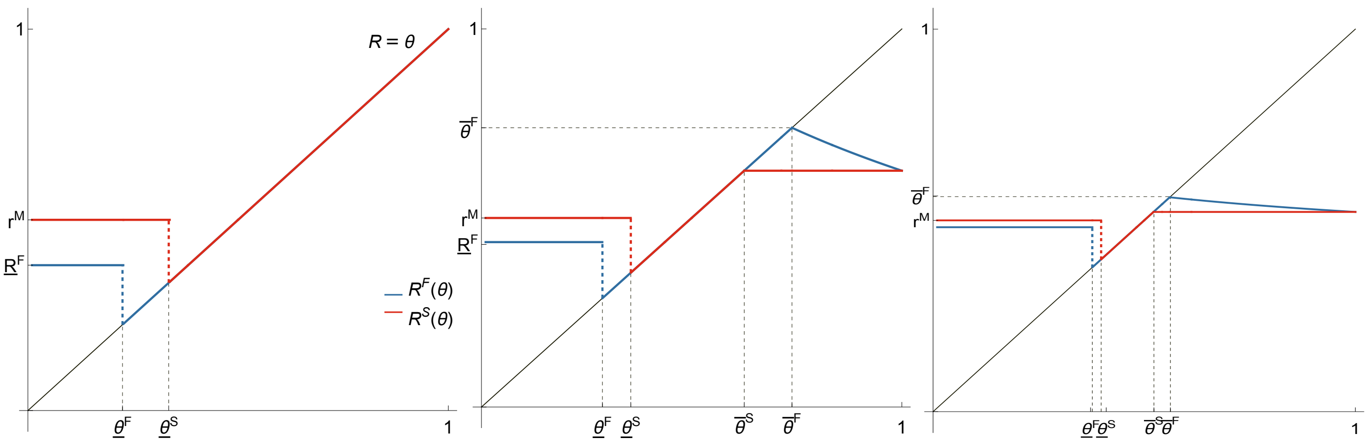

Figure 3.2: Effect of η on the static reserve prices in FPA (blue) and SPA (red) for $\theta \sim \mathcal{U}[0,1]$, and $n = 2$.

The next proposition shows that the FPA does more tracking overall — i.e., the FPA also tracks some incumbents which the SPA excludes — and performs some comparative statics.

Proposition 5. *The following results hold:*

- *There is a (common) threshold $\bar{\eta}$ such that both auction formats track all incumbents if and only if $\eta < \bar{\eta}$:* $\overline{\theta}^i = 1 \Leftrightarrow \eta < \overline{\eta}$
- *In both FPA and SPA, the set of incumbent's values that the seller tracks is decreasing in η and n:* $\frac{\partial \underline{\theta}^i}{\partial \eta}, \frac{\partial \underline{\theta}^i}{\partial n} > 0$ and $\frac{\partial \overline{\theta}^i}{\partial \eta}, \frac{\partial \overline{\theta}^i}{\partial n} \leq 0$, *with strict inequality if $\eta > \overline{\eta}$.*
- *The tracking region is larger in the FPA than in the SPA: $\underline{\theta}^F < \underline{\theta}^S < \overline{\theta}^S \leq \overline{\theta}^F$, with strict inequality if $\eta > \overline{\eta}$.*
- $\lim_{\eta \to 1} R^i(\theta) = r^M$ *and* $\lim_{n \to \infty} R^S(\theta) = r^M$, *while* $\lim_{n\to\infty} \underline{\theta}^F < r^M < \lim_{n \to \infty} \overline{\theta}^F \; \forall \eta < 1$.

Figure 3.2 shows the optimal reserve price in the FPA and SPA for different values of η. In both auction formats, for sufficiently low values of η, the seller never tails an incumbents: the low valuations are excluded, the high valuations are tracked.[27] This is natural: if the risk that the incumbent leaves the auction is small the seller sets a reserve that fully expropriates even the high valuation incumbents that induce a significant (upward) distortion from the optimal reserve in a symmetric auction.

Increasing either η or n reduces the seller's incentives to track the incumbent, because he is more likely to lose or not participate in the auction. Intuitively, the higher is η, the riskier it is for the seller to set a reserve price that departs from the optimal reserve price for standard bidders. As $\eta \to 1$, the incumbent never participates and the seller disregards the information and sets the reserve price equal to r^M in both auction formats. As $n \to \infty$, the probability of the incumbent winning goes to zero, and the seller always sets the reserve price equal to r^M in the SPA. In the FPA, instead, even when $n \to \infty$ the reserve price depends on the value of the incumbent and there exists a region of incumbent's types that the seller tails.[28]

[27] The threshold $\overline{\eta}$ is independent of the auction format and solves $\frac{1-\overline{\eta}}{\overline{\eta}} = n\psi(1) f(1)$. That the highest possible incumbent is tracked in the FPA if and only if it is tracked in the SPA is another consequence of the interim equivalence of the two formats when the incumbent is 1 (see discussion after Proposition 6): If he leaves, then the two formats are the same by the RET; if he stays he wins for sure and makes the same expected payment (as a function of the reserve).

[28] In this case, the tailing reserve price $R_{lim}(\theta)$ solves $\psi(R_{lim}(\theta)) = -\frac{(1-\eta)\log(F(\theta))}{\eta f(R_{lim}(\theta))}$.

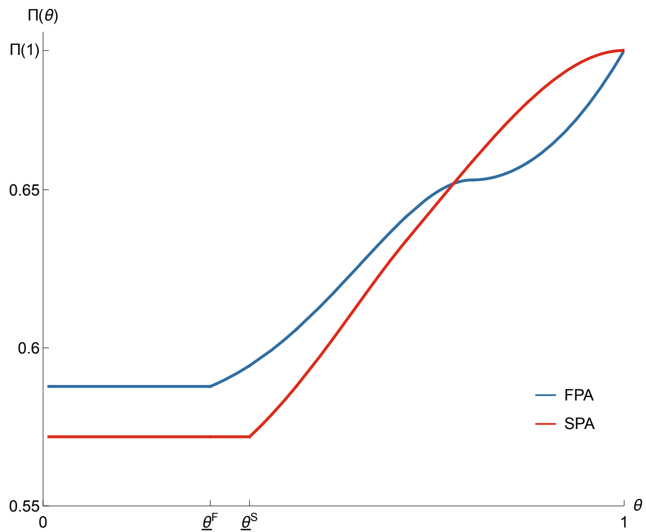

Figure 3.3: Seller's revenues in FPA and SPA as a function of the incumbent's value, with $\theta \sim \mathcal{U}[0,1]$, $\eta = 0.5$, and $n = 4$.

The differences among auction formats hinge on the fact that, conditional on the incumbent participating, the seller's reserve price plays a stronger role in the FPA than in the SPA. In fact, in the FPA, the reserve price affects the bids of standard bidders and hence the seller's revenue when the incumbent loses.[29] By contrast, in the SPA the seller's reserve price is irrelevant when the incumbent loses, since the incumbent's bid (which is equal to his value and hence higher than the seller's reserve price with tailing) has the same role as a seller's chosen reserve price.[30] Intuitively, in the FPA the effect of the reserve price hinges on the reserve price being observable by bidders, while in the SPA any bid is equivalent to a reserve price, from the seller's point of view. Hence, the seller has a stronger incentive to increases the reserve price in the FPA than in the SPA, whenever the reserve price tails the incumbent.

3.2 Seller's Revenue in a Static Auction

We now analyze how $\Pi^i(\theta) := \max_R \pi^i(\theta, R)$ — i.e., the value of problem (3.2) — depends on the valuation of the incumbent θ and on the auction format. By monotonicity, this value is increasing in θ in either format, strictly unless θ is excluded. Because we are in a setting with asymmetric bidders, the Revenue Equivalence Theorem does not hold and the two formats might differ in the (optimized) revenue. The next Proposition — and its graphic representation in Figure 3.3 — proves that equivalence holds only at $\theta = 1$ and that in general the ranking is ambiguous, with the FPA (resp., SPA) dominating if the incumbent has low (resp., high) valuation.

Proposition 6. *If $\theta \leq \overline{R}^S$, then $\Pi^F(\theta) > \Pi^S(\theta)$. If $\overline{R}^S < 1$, then $\Pi^F(1) = \Pi^S(1)$ and there exists a cutoff $\tilde{\theta} > \bar{R}^S$ such that $\Pi^S(\theta) > \Pi^F(\theta)$ if $\theta > \tilde{\theta}$.*

The intuition for the revenue ranking in Proposition 6 is the following.[31] When the incumbent does not participate in the auction, then we are in a standard setting with n symmetric bidders so the RET applies and the FPA and SPA yield the same expected revenue (as a function of the reserve price). Moreover, when the incumbent θ

[29]Notice that an increase in either θ or n reduces this effect. The reason is that, if either the incumbent's value or the number of new bidders increase, so does the value of a wining entrant, and the reserve price has a weaker effect on higher-value bidders.

[30]Of course, in a standard auction without incumbents the role of a reserve price in the two formats is the same.

[31]The analytical specifications of the revenue function in the two formats, equations (7.6)-(7.7) in the Appendix, provide a complementary tool to guide this intuition.

participates and wins, for all R his expected transfer is $b(\theta, R, n)$ in both formats. This is because myopia has no effect on incumbents' behavior, so transfers conditional on winning are the same by the RET. Therefore, the revenue ranking is driven by how the expected transfer differs between the two auction formats in the event that the incumbent participates but loses. If $\theta = 1$ this event has zero probability, so seller's revenue is the same function of the reserve price in two formats, yielding the same policy and the same value.

When the incumbent θ loses an auction with reserve R to a bidder with value θ', the expected seller's revenue is $b(\theta', R, n)$ in the FPA — i.e., the bid submitted by the winner — and $b(\theta', \theta, n-1)$ in the SPA — i.e., the expected payment of the winner in an auction with $(n-1)$ standard bidders and reserve θ since the losing incumbent acts as a reserve price rather than as a "real" competitor.[32] Mechanically, the losing bid by the incumbent (equal to his value) overrules the reserve price in the SPA, but not in the FPA where reserve price needs to be announced to bidders in order to affect their bids. Likewise, losing an active competitor affects the expected transfer of each winner in the SPA, while it is of no consequence in the FPA where the bid is sufficient to determine the realized transfer. Analytically, we see two differences in these expressions: *i*) the last argument is n in the FPA and $n-1$ in the SPA, and *ii*) the second argument is R in the FPA, which is lower (and strictly so in the tailing region) then θ in the SPA. We call the first *myopic bidding* effect and the second *reserve price* effect. Since — per Fact 1 — b is increasing in all its arguments, the two effects go in opposite directions and ranking the formats requires a quantitative evaluation of the two.

If there is no tailing, i.e. if $\eta < \bar{\eta}$ defined in point *i*) of Proposition 5, then the reserve price effect is neutralized because $R = \theta$ for all (non-excluded) incumbents.[33] Therefore, the only difference in revenues across the two auction formats is driven by the myopic bidding effect, which makes *any* new bidder θ' who wins against a tracked incumbent θ pay $b(\theta', \theta, n)$ in the FPA, strictly more than his expected payment $b(\theta', \theta, n-1)$ in the SPA. By static myopia, standard bidders fail to realize that a reserve price equal to the value of a competitor means that there are only $(n-1)$ real competitors in the auction, and hence they overestimate competition. This has no effect in the SPA, while it induces overbidding in the FPA.

In the tailing region the reserve price effect adds a counteracting force: when the incumbent loses to an entrant with value θ', the seller's revenue is $b(\theta', \theta, n-1)$ in the SPA (independently of the reserve price that she sets) and $b(\theta', \overline{R}^F(\theta), n)$ in the FPA, because in the SPA the incumbent's bid equal to θ acts as a reserve price which is higher than $\overline{R}^F(\theta)$. Since this reserve price effect is increasing in θ, overall the seller's revenue is higher in the SPA if and only if the incumbent's value is sufficiently high.[34]

Although revenue in the static environment is higher in the FPA for a larger set of incumbent's values than in the SPA, as displayed in Figure 3.3, in repeated auctions an incumbent's value is relatively more likely to be high than low, since incumbents are winners of previous auctions. And the static revenue is higher in the SPA precisely

[32] Allowing for exclusion, the expected transfer is $b(\theta', \max\{\theta, R\}, n-1)$. We focus for brevity on the interesting case where the incumbent is not excluded by the reserve, but actually overrules it in case he is defeated.

[33] Again, the discussion focuses on the regions — which will be relevant for the dynamic model — where the seller does not exclude the incumbent. If both auctions exclude the incumbent it is easy to see that the FPA dominates: in the FPA the new bidders bid as if they had $n-1$ competitors while in reality they only have $n-2$ since the incumbent is excluded by the reserve; by myopic bidding any winner makes therefore a strictly larger transfer in the FPA than in the SPA for any (exclusion) reserve price and *a fortiori* at the optimum. When the SPA excludes and the FPA tracks, i.e. $\theta \in [\underline{\theta}^S, \underline{\theta}^F]$ the FPA dominates because, by the same argument, it would dominate if she chose to exclude using $\underline{R}^S$.

[34] In Carannante et al. we focus on a similar revenue comparison holding the reserve price fixed. In that case, we show a similar result. Proposition shows that the result is robust to letting the seller reoptimize and set a reserve price. Indeed, the result will be useful in the dynamic setting when optimally the seller sets the same reserve for each given incumbent, which happens when $\eta \approx 1$ so that $R(\theta) = r^M$ in both formats.

when the incumbent's value is high. So it is not surprising that, in the dynamic setting, the seller may obtain a higher revenue in the SPA than in the FPA, as we going to show in Section 4.

4 Repeated Auctions

In this section, we consider our dynamic model, where the seller's information about the incumbent arises endogenously, since the incumbent in an auction is the winner of a previous auction. In Section 4.1, we start by analyzing the optimal reserve prices chosen by the seller in the FPA and SPA, and then study the characteristics of the minimum reserve price (Section 4.1.1). Section 4.2 describes typical paths of reserve prices and winners' values in repeated auction and Section 4.3 analyzes and compares the seller's revenue in the FPA and SPA.

In every period, the seller maximizes the sum of current and discounted future profits in the repeated auctions, by choosing a reserve price as a function of the whole history of bids. By assumption 1, the seller's problem can be described as a dynamic control problem, in which the state is either the valuation of the incumbent who won the auction in the previous period, or state $\emptyset$ if there was no bid above the reserve price in the previous period. We are going to show that, in contrast to the static environment, by choosing a reserve price the seller affects not only her revenue in the current period, but also the transition dynamics between the possible states and therefore her future payoff.

4.1 Optimal Reserve Prices

In Theorem 7 we will prove that, on path, the seller never excludes an incumbent — i.e., she never chooses a reserve price that is higher than the value of a bidder that may win an auction.[35] Anticipating this result, in the iPA, $i = F, S$, we can write the value function of the seller's problem in state θ as

$$V^i(\theta) = \max_{R \leq \theta} \pi^i(\theta, R) + \beta \left[\eta \left(F(R)^n V^i_\emptyset + \int_R^1 V^i(\theta') \, \mathrm{d}F(\theta')^n \right) + \right. \\ \left. + (1-\eta) \left(F(\theta)^{n-1} V^i(\theta) + \int_\theta^1 V^i(\theta') \, \mathrm{d}F(\theta')^{n-1} \right) \right], \tag{4.1}$$

and in state $\emptyset$ as

$$V^i_\emptyset = \max_R \pi_n(R) + \beta \left[F(R)^n V^i_\emptyset + \int_R^1 V^i(\theta') \, \mathrm{d}F(\theta')^n \right], \tag{4.2}$$

where in both expressions the first terms are the static revenues defined in Section 3. In state $\emptyset$, the expected continuation value reflects the fact that the next-period state is: (i) $\emptyset$ if no bidder has a value above R or (ii) the value of the auction winner θ'. By contrast, when there is an incumbent with value θ, the next-period state is: (i) $\emptyset$ if the incumbent leaves and no new bidder has a value above R; (ii) the incumbent's value θ if he stays and wins again; (iii) the value of the highest new bidder θ', if he outbids the incumbent or if the incumbent leaves.

Notice that the expression for the static revenue in auction format i, given by (3.2), represents the only difference between the value functions in the FPA and SPA. The reason is that, given any initial state and reserve price, the

[35] Intuitively, this is because in order to sell today to a type that the seller wants to exclude tomorrow, she has to set a reserve price that is suboptimally low: this reduces today's expected profits and does not provide any benefit tomorrow. The Appendix describes the general expressions for the value functions that take into account the possibility of exclusion.

transition between states is identical in both auction formats, since the seller always learns the value of the highest bidder if this is higher than the reserve price and this is the only thing that affects the transition dynamic.[36]

The expressions for the value functions show that the reserve price affects the transition dynamics only when there is no incumbent (i.e., either in state θ when the incumbent leaves or in state $\emptyset$), because in this case a new bidder has to bid above the reserve price to become the incumbent.[37] Therefore, a marginal increase in the reserve price R shifts the continuation value from $V(R)$ to $V_\emptyset$ in case there is no incumbent and the highest new bidder's value is exactly equal to R. This dynamic effect is independent of the incumbent's value and represents either a benefit or a cost for the seller, depending on whether $V(R)$ is higher or lower than $V_\emptyset$.

Our main result characterizes the dynamically optimal reserve price in the two auction formats.[38]

Theorem 7. *Consider auction format $i = F, S$. In state $\emptyset$, the optimal reserve price $R^i_\emptyset$ solves*

$$\psi(R^i_\emptyset) = -\beta\left(V^i\left(R^i_\emptyset\right) - V^i_\emptyset\right). \tag{4.3}$$

The optimal reserve price in state θ is

$$R^i(\theta) = \begin{cases} \theta & \text{if } R^i_\emptyset \leq \theta < \overline{\theta}^i \\ \overline{R}^i(\theta) & \text{if } \theta \geq \overline{\theta}^i \end{cases}$$

where $\overline{R}^i(\overline{\theta}^i) = \overline{\theta}^i$ and

- *$\overline{R}^S(\theta) = \overline{\theta}^S$ such that*

$$\psi(\overline{\theta}^S) = \frac{1-\eta}{n\eta f(\overline{\theta}^S)} - \beta(V^S(\overline{\theta}^S) - V^S_\emptyset), \tag{4.4}$$

- *$\overline{R}^F(\theta)$ is strictly decreasing in θ and solves*

$$\psi(\overline{R}^F(\theta)) = \frac{(1-\eta)(1-(n-1)\log(F(\theta)))}{n\eta f(\overline{R}^F(\theta))} - \beta(V^F(\overline{R}^F(\theta)) - V^F_\emptyset). \tag{4.5}$$

Figure 4.1 shows the optimal reserve prices in the FPA and SPA, and compares them with the ones in the static environment. The optimal reserve prices characterized by Theorem 7 have the same qualitative characteristics as in the static setting, with three notable differences. First, the reserve price is always weakly greater than the initial reserve price $R^i_\emptyset$ that the seller chooses in state $\emptyset$, which implies that the seller never observes an incumbent with a value lower than $R^i_\emptyset$. Second, the seller never excludes an incumbent. Third, in the tailing region, the reserve price is lower than in the static setting because of dynamic concerns.

The intuitive reason why there is no exclusion on path is that, if the seller wants to exclude an incumbent in a period, then in the previous period she has to set a reserve price that does not exclude him, which cannot be optimal. Formally, in the proof of Theorem 7 we show that the seller only wants to exclude types that are lower than the minimum reserve price $R^i_\emptyset$, and hence that can never win an auction and become incumbents.

[36] Clearly, this does not imply that the two auction formats have the same dynamics, because optimal reserve prices (as a function of the incumbent's value) are different in the FPA and SPA.

[37] The reserve price has no dynamic role when the incumbent stays, because without exclusion the incumbent's bid is never lower than the reserve price and, hence, the reserve price never binds for new bidders.

[38] Abusing notation, we use the same notation for the optimal reserve prices in a static auction and in the repeated auctions. This reflects the fact that, as we are going to show, these two types of reserve price share the same qualitative features.

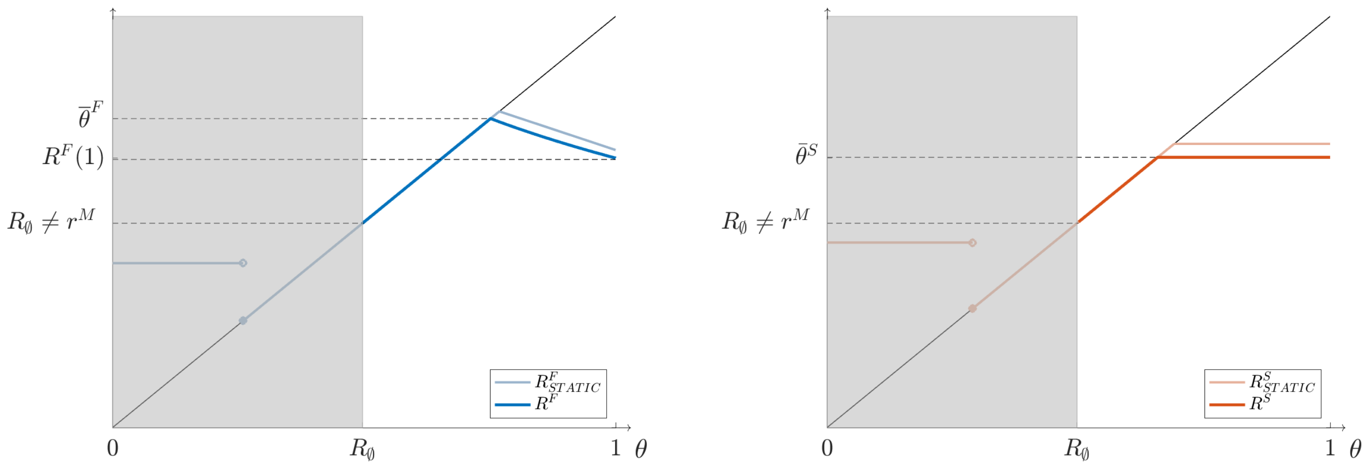

Figure 4.1: Optimal dynamic (bright color) and static (shaded color) reserve prices in FPA (left) and SPA (right) as a function of the state θ, with $\theta \sim \mathcal{U}[0,1]$, $\beta = 0.5$, $\eta = 0.5$, and $n = 2$.

In the tailing region — i.e., when $\theta \geq \overline{\theta}^i$ — the conditions that characterize the optimal reserve prices $\overline{R}^i(\theta)$ in Theorem 7 include the dynamic wedge $\beta(V^i(\overline{R}^i(\theta)) - V^i_\emptyset)$, $i = F, S$, in addition to the static wedges of Proposition 4. This wedge reflects the effect of increasing the reserve price on the transition dynamics, that we have discussed above. Since this dynamic wedge does not depend on θ, as in the static setting the tailing reserve price is decreasing in θ in the FPA and independent of θ in the SPA. Moreover, in the proof of Theorem 7 we show that the dynamic wedge is always positive (since $V^i(\overline{R}^i(\theta)) > V^i_\emptyset$) and hence that the reserve price is strictly lower than in the static setting.

The dynamic wedge vanishes if either $\beta = 0$ or $\eta = 1$. When $\beta = 0$, the repeated auctions are equivalent to a sequence of static auctions where the seller maximizes her profit given the information of the previous auction (as analyzed in Section 3). When $\eta = 1$, the incumbent always leaves and we have a sequence of static auctions with symmetric bidders. In this case the seller to chooses a reserve price equal to r^M, regardless of the state, yielding a constant value function equal to $\frac{\pi_n(r^M)}{1-\beta}$.[39] Of course, exactly as in the static setting, when $\eta = 0$ the seller always tracks the incumbent.

The same dynamic wedge also determines the optimal reserve price $R^i_\emptyset$ that the seller chooses when there is no incumbent (see condition (4.3)). In this case, however, $V^i(R^i_\emptyset) - V^i_\emptyset$ might be either positive or negative, as we discuss in Section 4.1.1.

Notice that Theorem 7 allows to significantly simplify the analysis of our dynamic model.[40] In the SPA, the whole dynamics of the repeated auctions is completely pinned down by the two extreme values of the tracking region. Hence studying these two values (and their comparative statics) allows us to fully characterize all the effects of changes in the model primitives. For the FPA, by contrast, since the tailing reserve price is not constant, a complete analysis requires the whole schedule of reserve prices in $[\overline{\theta}^F, 1]$. For many qualitative results, however, since the tailing reserve price is strictly decreasing for those incumbent's values, we only need the three thresholds

[39] In the limit $n \to \infty$, the seller can sell to a bidder with value 1 in every period, so that the value function is also constant (and equal to $\frac{1}{1-\beta}$). But while both $\eta = 1$ and $n \to \infty$ yield the same reserve price r^M when there is no incumbent, they induce different reserve prices when there is an incumbent. The reason is that the static wedge vanishes in the SPA but not in the FPA (where the reserve price depends on the incumbent's value even when $n \to \infty$ by Proposition 5).

[40] The theorem also has direct implications of economic relevance that we detail in Section 4.2.

$R_{\emptyset}^{F}$ and $\overline{\theta}^{F}$ and $\overline{R}^{F}(1)$.

4.1.1 Minimum Reserve Price

Since the seller never chooses a reserve price higher than the incumbent's value, $R_{\emptyset}^{i}$ represents the lowest possible bidder that the seller excludes from the repeated auctions. This induces the seller to distort the initial reserve price away from its static benchmark r^M (see Fact 3). From condition (4.3), $R_{\emptyset}^{i} < r^M$ if and only if $V^i\left(R_{\emptyset}^{i}\right) > V_{\emptyset}^{i}$: the seller reduces the reserve price relative to r^M if and only if she prefers an auction with the lowest possible incumbent rather than one with no incumbent at all.

As discussed previously, the dynamic distortion on the right-hand-side of condition (4.3) vanishes when either $\beta = 0$ or $\eta = 1$ and, hence, $R_{\emptyset}^{i} = r^M$. The following proposition shows that, for specific values of the model's parameters, the sign of the dynamic wedge $V^i\left(R_{\emptyset}^{i}\right) - V_{\emptyset}^{i}$ only depends on static revenue functions. In particular, an important role is played by the static seller's revenue when he tracks the incumbent θ, which is given by $\pi_{n-1,\theta}^{i}(\theta)$ (using the notation of Section 3.1). To simplify notation, we denote this revenue by $\widetilde{\pi}_{n-1}^{i}(\theta)$.

Proposition 8. *Local to $\beta = 0$ or $\eta = 1$, $R_{\emptyset}^{i} > r^M$ if and only if*

$$\widetilde{\pi}_{n-1}^{i}\left(r^M\right) - \pi_n\left(r^M\right) < 0. \tag{4.6}$$

When $\eta = 0$, $R_{\emptyset}^{i} > r^M$ if and only if

$$\widetilde{\pi}_{n-1}^{i}\left(r^M\right) - \pi_n\left(r^M\right) < \beta \int_{r^M}^{1} \frac{\mathrm{d}}{\mathrm{d}\theta'}\widetilde{\pi}_{n-1}^{i}(\theta')\, \frac{F(\theta')^{n-1}(1-F(\theta'))}{\left(1-\beta F(\theta')^{n-1}\right)}\mathrm{d}\theta'. \tag{4.7}$$

Condition (4.6) has a clear economic interpretation. Local to $\beta = 0$ or $\eta = 1$, $R_{\emptyset}^{i}$ is lower than r^M if and only if, in a static auction with $n-1$ new bidders and reserve price r^M, the seller prefers to insure against no trade (i.e., to obtain r^M when all bidders are below r^M), rather than having one additional bidder. We show in the Appendix that this comparison depends on the number of bidders: when $\eta = 1$, there exists a threshold $\bar{n}^i$ such that $R_{\emptyset}^{i} > r^M$ if and only if $n > \bar{n}^i$. This is displayed in Figure 4.2, that shows the reserve price $R_{\emptyset}^{i}$ in the FPA and SPA for different values on n.

At $\eta = 0$, instead, the sign of the difference between $R_{\emptyset}$ and r^M is determined by condition (4.7), which has an additional term. Since $\frac{\mathrm{d}}{\mathrm{d}\theta'}\widetilde{\pi}_{n-1}^{i}(\theta') > 0$, this term is always positive and, if condition (4.6) is satisfied, then condition (4.7) is satisfied too. Therefore, $R_{\emptyset}^{i} > r^M$ local to $\eta = 1$ implies $R_{\emptyset}^{i} > r^M$ local to $\eta = 0$.[41] When condition (4.7) is satisfied but condition (4.6) is not, $R_{\emptyset} > r^M$ at $\eta = 0$ and converges to r^M from below for $\eta \to 1$, implying that $R_{\emptyset}$ is non-monotone in η. The central panel of Figure 4.2 shows an example of this.

Figure 4.2 also provides insight on how the minimum reserve price varies across auction formats, suggesting that it is lower in the FPA then in the SPA. The next Proposition establishes this result for extreme values of η.

Proposition 9. *Local to $\eta = 0$ and $\eta = 1$, $R_{\emptyset}^{F} < R_{\emptyset}^{S}$.*

In the Appendix, we prove that the difference between $R_{\emptyset}^{i}$ and r^M at $\eta = 0$ and $\eta = 1$ is proportional to the expressions in conditions (4.6) and (4.7) and is larger in the FPA than in the SPA (because tracking revenues are

[41] The converse is not true because, when condition (4.7) is satisfied, $\widetilde{\pi}_{n-1}^{i}\left(r^M\right) - \pi_n\left(r^M\right)$ can be either positive or negative.

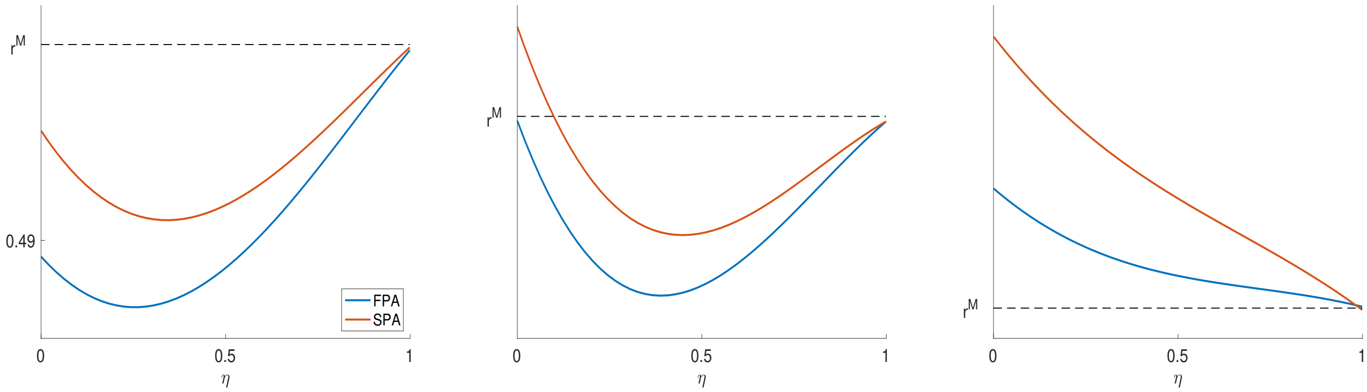

Figure 4.2: Initial reserve price $R_\emptyset$ as a function of η in FPA and SPA with $\theta \sim \mathcal{U}[0,1]$, $\beta \approx 0.6$ and, from the left, $n = 2$, $n = 3$ and $n = 6$.

higher but *marginal* revenues are lower in the FPA than in the SPA, see Fact 14). Since the only instrument used by the seller to exclude bidders in the dynamic setting is the initial reserve price, the fact that this reserve price is lower in the FPA than in the SPA (Figure 4.2 and Proposition 9) is consistent with the result that, in the static setting, the seller tracks some incumbents in the FPA that she excludes in the SPA. For example, at $\eta = 1$ the reserve price is the same in both auction formats, but slightly reducing η has a higher effect on FPA than on SPA, therefore $R^F_\emptyset < R^S_\emptyset$.

4.2 Paths of Reserve Prices and Auction Winners

Theorem 7 allows us to describe the dynamics of reserve prices and incumbents that can be observed in repeated auctions, when the seller sets the optimal reserve prices. Figure 4.3 displays an example of typical paths of reserve prices and winners' values in the FPA and SPA. Notice that the paths restart with the initial reserve price $R_\emptyset$ and no incumbent after any period in which there was no trade. The paths are often increasing after a period of trade, but sometimes decreasing because of the optimal reserve price schedule in repeated auctions.

In each period, there is a lower bound on the possible value of a winner, which is either the value of the previous winner (if he stays), or the reserve price (if the previous winner leaves). In the tracking region, these two bounds coincide and, when there is trade, the value of the auction winner can only increase. In the tailing region, instead, the value of the winner can decrease if the incumbent leaves, because the reserve price is lower than the incumbent's value.

The dynamics of the winner's values gives only a partial view of the dynamics of the reserve prices, as the two move in the same direction only with tracking. In the tailing region, the reserve price is unresponsive to the value of the winner in the SPA, and moves *against* it in the FPA. For this reason, in the SPA all downward movements in the winner's value have no consequence on the reserve price, which either increases until it reaches $\bar{\theta}_S$ and then remains constant, or reverts back to $R_\emptyset$ after a period of no trade. In FPA, by contrast, the reserve price is strictly increasing only up to $R^F(1)$, while it can both increase and decrease in the interval $[R^F(1), \overline{\theta}_F]$. This depends on whether a previous winner does not win again because he faces a new bidder with a higher value (in which case the reserve price decreases) or because he leaves the auction (in which case the reserve price can both increase or decrease, depending on the whether the value of the new winner is higher or lower than the value of the previous one).

Finally, notice that although the two auction formats have the same transition dynamic for a given reserve price, they will generically induce different realizations of paths of winners' values and reserve prices, because the seller chooses different reserve prices in the two formats. For example, consider the simple case where $\beta = 0$ (so that the dynamic distortions discussed in Section 4.1 vanish). In this case, for a given incumbent, the reserve price is higher in FPA than in SPA (and strictly so when there is tailing in the SPA). Therefore, for any incumbent's value, the FPA results in lower trade and efficiency than the SPA. Of course, total efficiency depends also on the stationary distribution of incumbents, which in turn depends on the reserve prices chosen by the seller and therefore differs in the two auction formats. We analyze this stationary distribution and its effects on trade in Section 5.

4.3 Seller's Revenue: FPA vs. SPA

In this section, we compare the seller's revenue in the FPA and SPA. Recall from, the analysis of the static environment (Proposition 5), that the revenue ranking of the two auction formats depends on the incumbent's value. Specifically, the FPA benefits from myopic bidding with low-value incumbents, while the SPA exploits high-value incumbents acting as implicit reserve prices. These dynamics carry over to a dynamic setting, but are further influenced by the stationary distribution of incumbents' values.

Since incumbents are winners of previous auctions, they tend have relatively high valuations, which skews the distribution of realized incumbents' valuations towards regions where the SPA dominates in the static setting. Consistent with this intuition, Figure 4.4 shows that, when values are uniformly distributed, the seller's revenue is higher in the SPA than in the FPA for a wide range of values of η.[42]

Moreover, given that η affects both the expected valuation of an incumbent and seller's reserve prices, it is not surprising that the revenue ranking between FPA and SPA depends on η. In particular, the higher is η, the higher is seller's incentive to tail the incumbent by setting a reserve price lower than his valuation. In this case, the SPA allows the seller to choose a reserve price closer to the optimal one for new bidders than in the FPA (as we have discussed in Section 3), which tends to increase revenue in the SPA. By contrast, a sufficiently low η induces the seller to track the incumbent independently of his valuation, which favors the FPA because of myopic bidding (as we have discussed in Section 3).

The next proposition formalizes this intuition and compares the seller's revenue in the FPA and SPA when either η is very small, or very large.

Proposition 10. *There exists a neighborhood $\mathcal{N}_0$ of $\eta = 0$ such that $V_\emptyset^F(\eta) > V_\emptyset^S(\eta)$, $\forall \eta \in \mathcal{N}_0$. Suppose that there is a unique $x^\star$ such that $\psi(x^\star) = b\left(x^\star, r^M, n\right)$. Then there exists a neighborhood $\mathcal{N}_1$ of $\eta = 1$ such that $V_\emptyset^S(\eta) > V_\emptyset^F(\eta)$, $\forall \eta \in \mathcal{N}_1 \setminus \{1\}$.*

It is not surprising that the expected revenue is higher in the FPA than in the SPA when $\eta = 0$, because in this case both auction formats always track the incumbent and they only differ in the initial reserve price $R_\emptyset^i$. But recall that myopia induces new bidders to overpay in the FPA when the reserve price is equal to the value of the incumbent (see the discussion following Proposition 6). Hence, if the seller chooses an initial reserve price equal to $R_\emptyset^S$ in the FPA, she obtains exactly the same path of winners as in the SPA and a higher revenue. Choosing the optimal initial

[42] Figure 4.4 also shows that the seller's revenue in both auction formats is decreasing in η, because an increase in η reduces bidders' capacity and hence worsens the distribution of incumbents' values. More precisely, for a given reserve price policy, a higher η results in a distribution of bidders' values that FOSD the distribution with a lower η, in every period, and therefore a higher expected revenue. The expected revenue is even higher with a higher η when the seller chooses the optimal reserve price policy.

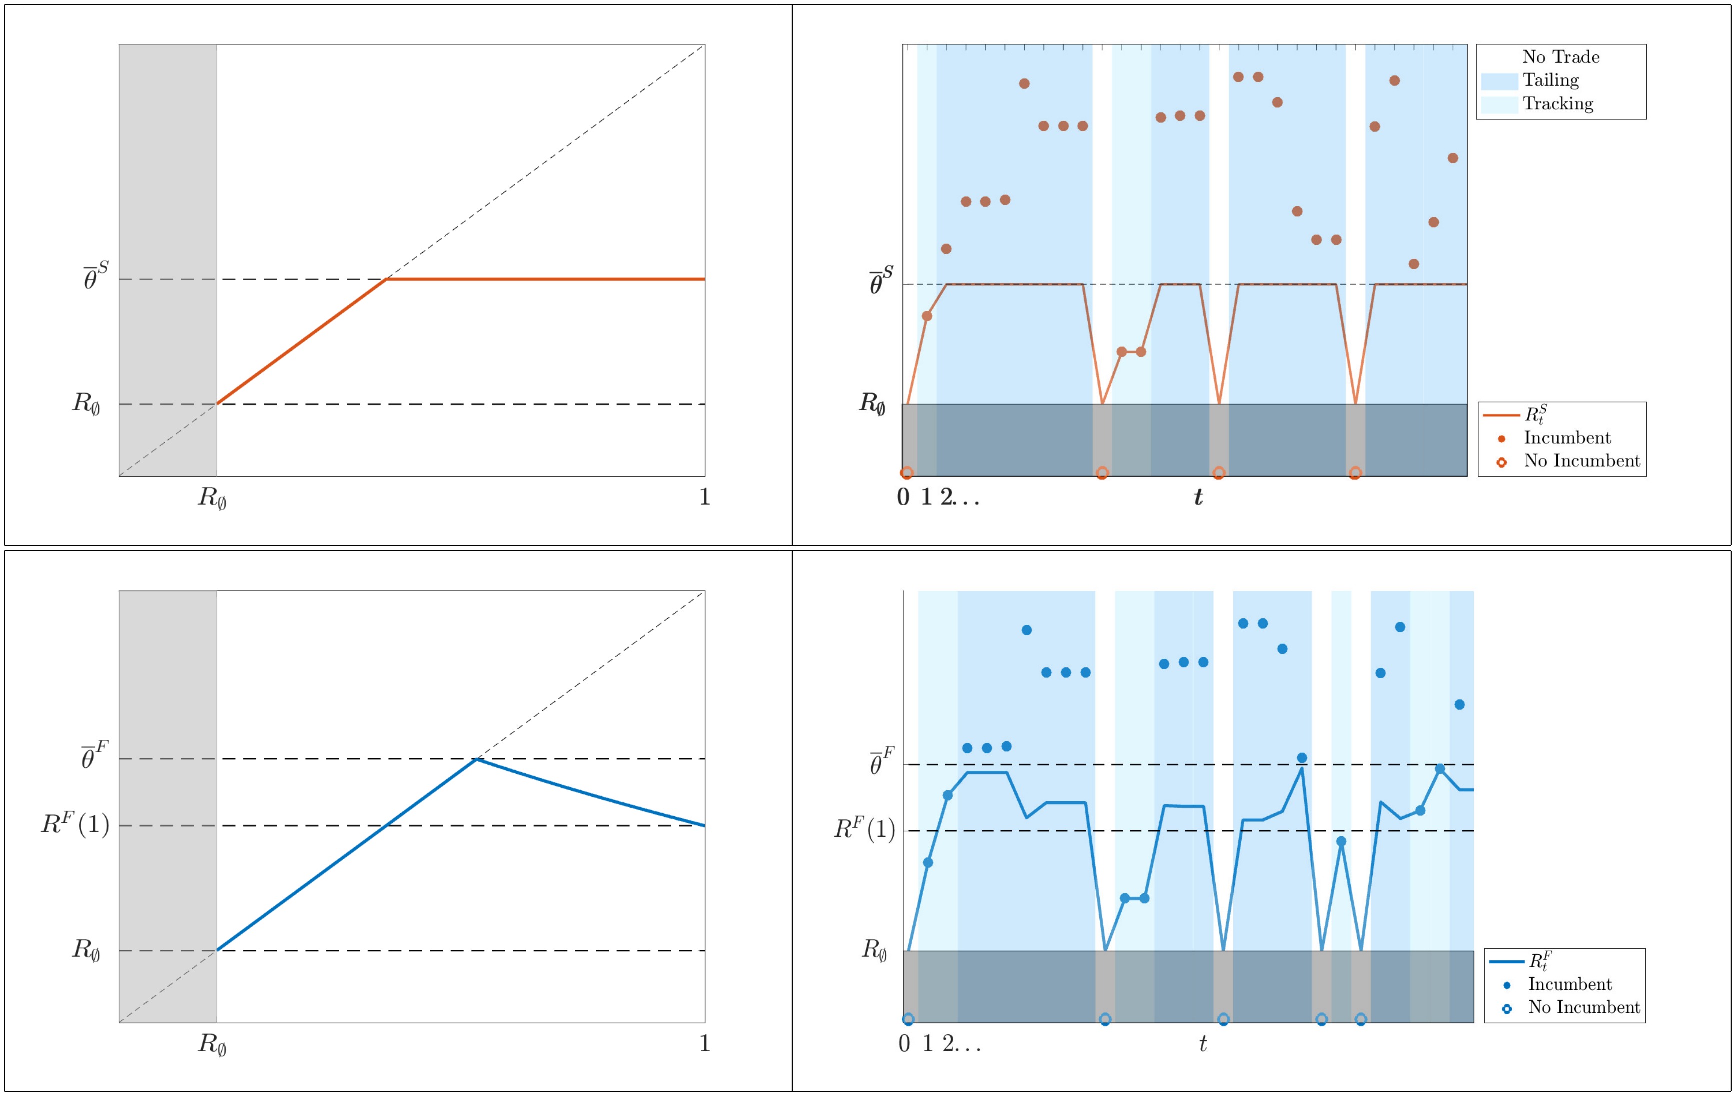

Figure 4.3: Reserve price schedules (left) and typical paths of incumbent's values (dots) and reserve prices (continuous line) in the repeated auctions (right), for SPA (top) and FPA (bottom). In the paths plots, the white areas highlight periods with no incumbent, the dark shaded areas periods with tracking, and the light shaded areas period with tailing.

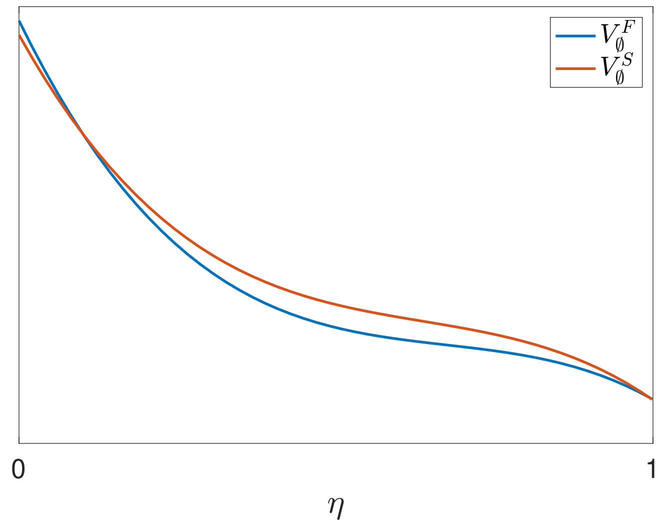

Figure 4.4: Seller's value in FPA and SPA as a function of η, with $\beta = 0.25$ and $n = 6$.

reserve price $R_\emptyset^F$ can only further increase revenue in the FPA compared to the SPA.[43]

At $\eta = 1$, by contrast, myopic bidding has no effect because the incumbent always leaves and the seller chooses the same reserve price in the FPA and the SPA, so that the two auction formats are revenue equivalent. Local to $\eta = 1$, since changes in the reserve prices and the incumbent's distribution are second order, we can interpret a marginal reduction in η as replacing an arbitrary bidder in an optimal static auction with a "special" bidder.

If this special bidder has a value which is drawn from the unconditional distribution $F(\cdot)$, then by the Revenue Equivalence Theorem the seller obtains the same expected revenue in the FPA and SPA. Conditional on the *actual value* of the special bidder, however, we show that the seller's revenue is higher in the FPA if this value is low, and is higher in the SPA if this value is high.[44] Now, in our environment of repeated auctions, the special bidder introduced in the static auction is actually an incumbent and, hence, his value is drawn form a distribution that FOSD F — the incumbent's value is the maximum of n draws from the unconditional distribution F, truncated above the reserve price. This shifts weight towards higher realizations of the special bidder's values, which results in a higher revenue in the SPA than in the FPA.

The previous argument hinges on the fact that the two auctions have the same reserve price and the same distribution of incumbents, which requires $\eta \approx 1$. For values of η outside this neighborhood, reserve prices differ in the FPA and SPA, which influence the seller's revenue when the incumbent leaves. As we have discussed in Section 3.2, the fact that the seller typically sets a lower reserve price in the SPA (closer to r^M) favors this auction format in this case. This additional force explains the fact that revenue in the SPA can often be higher than in the FPA — as is apparent from inspection of Figure 4.4 — despite the effect of myopic bidding that may induce to believe that the FPA should be preferred by the seller. Therefore, our quantitative exercise suggests that the effect

[43] Notice that this argument holds for any η that is sufficiently low to induce the seller to always track the incumbent in both auction formats (which is the case, for example, when $\beta = 0$ and $\eta < \overline{\eta}$, where $\overline{\eta}$ is defined in the proof of Proposition 5).

[44] The assumption in Proposition 10 of single crossing between the virtual value and the bidding function in the FPA ensures that there is a unique threshold such that revenue in the SPA is higher than in the FPA if and only if the special bidder's value is higher than the threshold. This is a sufficient condition for our result that only depends on the distribution F. As we discuss in the proof of Proposition 10, however, the assumption is not necessary since the result also holds for distributions that do not satisfy the assumption.

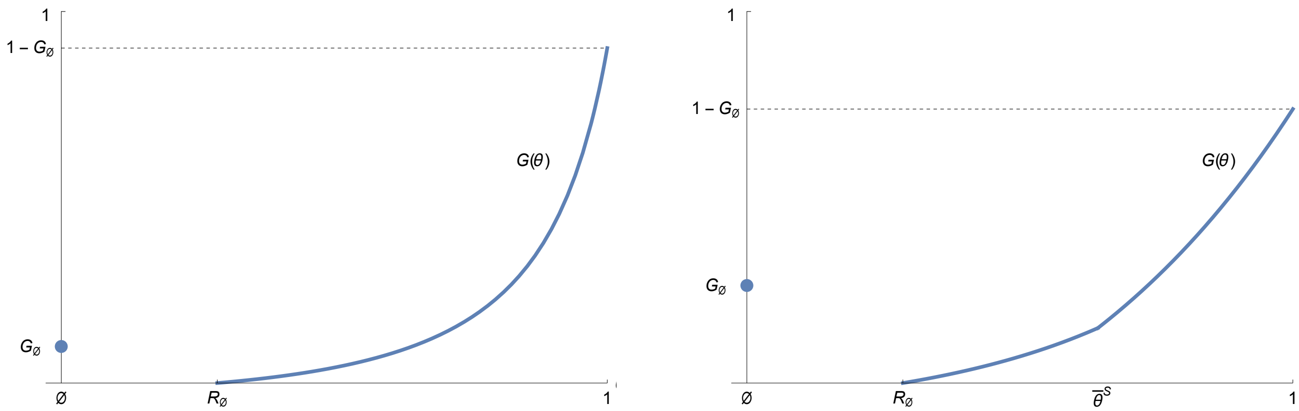

Figure 5.1: Stationary distributions over incumbent's types with η low (left) and high (right).

of myopic bidding is small relative to the role that incumbents play as implicit reserve prices in the SPA, which allows the seller to tailor the reserve price to the event that the incumbent leaves.

5 Stationary Distribution of Winners' Values and Trade

In this section, we derive the stationary distribution of incumbent's values in the repeated auctions, combining the possible transitions between states described in Section 4.1 with the optimal reserve prices characterized by Theorem 7. We then use this distribution to analyze trade — i.e., the probability that the object is sold to one of the bidders. All results refer to both the FPA and SPA and, to simplify notation, we suppress the superscript i that refers to the auction format.

Figure 5.1 shows the stationary distributions of incumbent's values for different values of η. This distribution has a point mass $G_\emptyset \in [0,1]$ — representing the probability that the auction has no incumbent — and is absolutely continuous on $[R_\emptyset, 1]$, with CDF $G(\theta)$ such that $G(R_\emptyset) = 0$ and $G(1) = 1 - G_\emptyset$. Notice that the distribution with lower η assigns more weight to higher incumbent's values.

An auction has no incumbent if: (i) the previous auction has no incumbent and no bidder bids above $R_\emptyset$, or (ii) the previous-auction incumbent leaves and no new bidder bids above the reserve price. Therefore, $G_\emptyset$ satisfies

$$G_\emptyset = G_\emptyset F(R_\emptyset)^n + \eta \int_{R_\emptyset}^1 F(R(\theta))^n \,\mathrm{d}G(\theta). \tag{5.1}$$

An auction has incumbent $\theta \geq R_\emptyset$ if: (i) the previous-auction incumbent θ stays and wins again, or (ii) the previous auction has no incumbent and the highest new bidder has value θ, or (iii) the previous auction has an incumbent $\theta' \neq \theta$, and the highest new bidder has value θ above θ' (when θ' stays) or above $R(\theta')$ (when θ' leaves). Therefore, $\mathrm{d}G(\theta)$ satisfies

$$\begin{aligned}\mathrm{d}G(\theta) =& \mathrm{d}G(\theta)(1-\eta)F(\theta)^{n-1} + G_\emptyset \mathrm{d}F(\theta)^n + \left[\eta \mathrm{d}F(\theta)^n + (1-\eta)\mathrm{d}F(\theta)^{n-1}\right]G(\theta) \\ &+ \eta \mathrm{d}F(\theta)^n \int_{\{\theta' : R(\theta') \leq \theta < \theta'\}} \mathrm{d}G(\theta'). \end{aligned} \tag{5.2}$$

The last term in the previous expression indicates a transition to a new incumbent θ that is lower than the previous one θ', which requires that the seller tails θ' with a reserve price $R(\theta') \leq \theta$ (see Figure 4.3). If $\theta < R(1)$, this condition is never satisfied; if $\theta > \overline{\theta}$, this condition is satisfied by all $\theta' > \theta$.[45] Finally, if $\theta \in [R(1), \overline{\theta}]$ (which is non-empty only in the FPA), θ can be reached from *high* incumbents $\theta' > R(\theta)^{-1}$, but not from *intermediate* incumbents $\theta' \in (\theta, R(\theta)^{-1})$ that induce a reserve price above θ.

In the Appendix, we derive the stationary distribution of incumbent's values following a two-step procedure. First, we fix $G_\emptyset$ and solve (5.2) as a function of θ. Because θ determines the specification of the set $\{\theta' : R(\theta') \leq \theta < \theta'\}$, as discussed above, this solution is characterized by three different functions. Second, we pin down $G_\emptyset$ imposing continuous pasting of the solutions obtained in step 1.

5.1 Trade

We define trade as the long-run frequency with which the object is sold in the repeated auctions — i.e.,[46]

$$T = 1 - G_\emptyset.$$

In this section, we analyze how long-run trade depends on the persistence of bidders and compare trade in the FPA and SPA.

When $\eta = 0$, since incumbents always stay in the auctions (and the seller chooses a reserve price lower than their value), the object is sold in every period following the first period with trade and, therefore, $T = 1$. When $\eta = 1$, the seller runs a sequence of (optimal) static auctions, implying that $T = 1 - F\left(r^M\right)^n$. It may be expected that trade decreases smoothly between these two extrema, because a lower η makes it more likely that incumbents stay in the auctions, which is a sufficient condition for trade. But this is only a partial intuition: Figure 5.2 shows an example where trade increases (i.e., $G_\emptyset$ decreases) in η, when η is high (see the Appendix for details).

In fact, total differentiation of (5.1) reveals that an increase in η affects $G_\emptyset$ through three different channels. First, there is a *direct effect* that reflects the change in $G_\emptyset$ implied by the incumbent leaving more often. This effect is always positive (consistent with our partial intuition) and is equal to

$$\int_{R_\emptyset}^{1} F(R(\theta))^n \, \mathrm{d}G(\theta). \tag{5.3}$$

Second, there is a *reserve price effect* that reflects the change in $G_\emptyset$ implied by the fact that the seller adjusts the reserve price when η changes:

$$G_\emptyset \frac{\mathrm{d}R_\emptyset}{\mathrm{d}\eta} \mathrm{d}F(R_\emptyset)^n + \eta \int_{R_\emptyset}^{1} \frac{\mathrm{d}R(\theta)}{\mathrm{d}\eta} \mathrm{d}F(R(\theta))^n \, \mathrm{d}G(\theta). \tag{5.4}$$

While the sign of $\frac{\mathrm{d}R_\emptyset}{\mathrm{d}\eta}$ is ambiguous (see Section 4.1.1), $\frac{\mathrm{d}R(\theta)}{\mathrm{d}\eta} \leq 0$ with a strict inequality in the tailing region. This reflects the fact that an increase in η induces the seller to reduce the tailing reserve price, which has a positive effect on trade.

Third, there is a *distribution effect* that reflects the change in $G_\emptyset$ resulting from a change in the stationary

[45] Since the tailing reserve price is (weakly) decreasing: (i) $\theta < R^i(1)$ implies that $\theta < R^i(\theta')$ all θ' and (ii) $\theta > \overline{\theta}^i$ (i.e., θ is tailed and $R(\theta) < \theta$) implies that $R(\theta') < \theta$ for all $\theta' > \theta$.

[46] In the market for online display advertising, impressions that are not sold through the auction are often allocated to guaranteed contracts of the publisher. In this case, $G_\emptyset$ reflects the percentage of impressions allocated to guaranteed contracts.

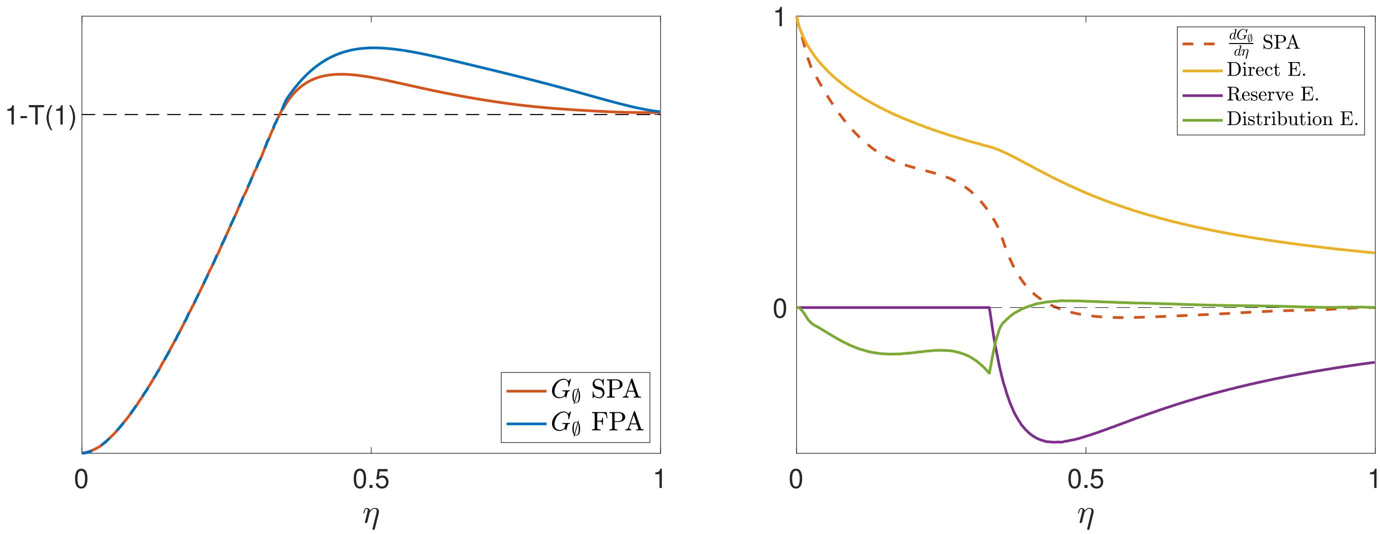

Figure 5.2: Left: $G_\emptyset$ as a function of η in the FPA and SPA. Right: Decomposition of $\frac{\mathrm{d}G_\emptyset}{\mathrm{d}\eta}$ in the SPA. In both plots, $\beta = 0$ and $n = 2$.

distribution of incumbent's values:

$$\frac{\mathrm{d}G_\emptyset}{\mathrm{d}\eta}F\left(R_\emptyset\right)^n - \eta\frac{\mathrm{d}R_\emptyset}{\mathrm{d}\eta}g\left(R_\emptyset\right)F\left(R_\emptyset\right)^n + \eta\int_{R_\emptyset}^{1}F\left(R\left(\theta\right)\right)^n\frac{\mathrm{d}}{\mathrm{d}\eta}\mathrm{d}G\left(\theta\right). \tag{5.5}$$

The total effect of η on $G_\emptyset$ is the sum of (5.3), (5.4) and (5.5).

Figure 5.2 shows an example of the three effects of η on $G_\emptyset$ in the SPA. Notice that the distribution effect is relatively small while the reserve price effect is zero when the seller tracks all incumbents (i.e., for η low) and strictly negative when there is tailing, which overcomes the direct effect when η is high.

The decomposition of the effects of η can also be used to investigate the transition dynamic induced by a permanent change in bidders' persistence.[47] The direct effect reflects the short-term consequence of an increase in bidders' persistence, when neither the reserve prices nor the stationary distribution of winners' value change. This effect always reduces trade. The additional reserve price effect reflects a medium-term response, when the seller reacts by adjusting the reserve prices in response to changes in bidders' persistence. This effects tends to be negative and may overcome the direct effect, resulting in an overall increase of trade. Finally, in the long term the stationary distribution of winners' values also changes as reflected by the distribution effect.

The next result compares trade in the two auction formats, and shows that there is no uniform ranking.

Proposition 11. *If $\overline{\theta}^F = \overline{\theta}^S = 1$, then $G^F_\emptyset < G^S_\emptyset$. If $\beta = 0$, then $G^F_\emptyset \geq G^S_\emptyset$, with a strict inequality whenever there is tailing.*

If $\overline{\theta}^i = 1$, the seller tracks all incumbents and trade only depends on the initial reserve price. In this case, since $R^F_\emptyset < R^S_\emptyset$ (by Proposition 9), trade is higher in the FPA. Thus, for parameter values that prevent tailing (e.g., η small), the FPA dominates the SPA both in terms of seller's revenue (Proposition 10) and in terms of trade.[48] The

[47]For example, a reduction in users' attention may induce advertisers to show their ad for a higher number of times before inducing a purchase, which results in a permanent reduction in η.

[48]Both results are driven by tracking, but for different reasons: the revenue result arises because the FPA is the best static format under tracking, while the trade result depends on the fact that tracking is relatively more desirable than exclusion in the FPA, which tends to reduce

presence of tailing introduces an important additional effect: the tailing reserve price is typically higher in the FPA, which tends to reduce trade compared to the SPA (see the left panel of Figure 5.2). However, even when $\beta = 0$ so that the initial reserve is the same in the two auction formats and the tailing reserve price is uniformly higher in the FPA, there are realizations of bidders' values such that trade is higher in the FPA than in the SPA. In the Appendix, we prove the second statement of Proposition 11 through a quantitative evaluation that accounts for the stationary distribution of incumbents.

6 Winners' Tenure

So far, our analysis examined either the ex-ante or the long-run properties of repeated auctions. Yet, given its Markovian structure, our model provides precise interim predictions on a variety of outcomes, such as dynamic revenue or probability of trade, based on the valuation of the incumbent. These predictions can be empirically tested only if the available dataset includes information on the (evolution of) incumbent's value, in our context equivalent to the history of bids. Although obtaining such dataset might be unfeasible, due to its proprietary nature, some characteristics of incumbents might still be observable. What predictions can be drawn based on such characteristics?

In this section, we analyze information that can be inferred from the tenure of the winner — i.e., how long he has won the repeated auctions. In our application to display advertisement, tenure is given by the number of consecutive times that a specific advertisement has been displayed to the same user.[49] First, we show that winners with higher tenure tend to have higher value (Lemma 12). Then we show that, depending on the auction format, tenure has different predictions on important auction moments (Proposition 13 and subsequent discussion); in particular (and somehow paradoxically), due to the interaction with the reserve price in the FPA, expected revenues can be decreasing in winners' tenure.

Denote the number of consecutive periods in which the current winner has won the repeated auctions, i.e. his *tenure*, with $\tau \in \{1, 2, \dots\}$, and let $g(\theta, \tau)$ be the joint distribution of incumbent's value and tenure.[50] An incumbent with value θ and tenure τ moves to tenure $\tau + 1$, i.e. wins for one additional period, if *i*) he stays in the auction and *ii*) he outbids all the $(n-1)$ new bidders. This event is independent of τ and has probability $(1-\eta)F(\theta)^{n-1}$. Hence, for $\tau \geq 1$ the joint distributions follow the recursion

$$g(\theta, \tau + 1) = (1-\eta)F(\theta)^{n-1} g(\theta, \tau) \tag{6.1}$$

while $g(\theta, 1) = g(\theta)\left(1 - (1-\eta)F(\theta)^{n-1}\right)$ is pinned down by the marginalization $g(\theta) = \sum_{\tau \geq 1} g(\theta, \tau)$. Notice that incumbents with higher valuation are more likely to increase their tenure, as $(1-\eta)F(\theta)^{n-1}$ is strictly increasing in θ. This suggests that high tenure predicts high-value incumbents. Formally,

Lemma 12. *Let $g_\tau(\theta) \equiv \frac{g(\theta,\tau)}{\int_{R_\emptyset}^1 g(\theta,\tau)d\theta}$ be the conditional distribution of the values of winners with tenure τ. $g_\tau(\theta)$ is increasing in τ in the FOSD order and converges in probability to 1.*

Our initial aim was to formulate predictions about the dynamic auction using observable attributes of the in-

the initial reserve price.

[49] This is observable even without accessing proprietary data from the Ad-Exchange.

[50] Notice that the stationary distribution $g(\theta)$ characterized in the previous section represents the marginal distribution of $g(\tau, \theta)$ with respect to θ, that is $g(\theta) = \sum_{\tau \geq 1} g(\theta, \tau)$.

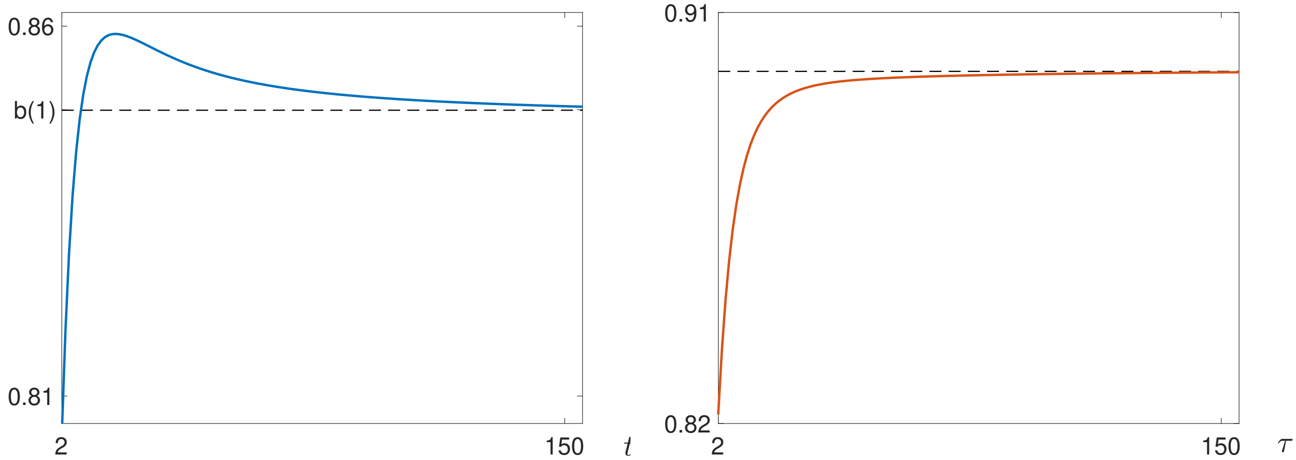

Figure 6.1: Winner's expected payment in FPA (left) and SPA (right) as a function of his tenure τ.

cumbents. Lemma 12 is an important intermediate step towards this goal, as it establishes that tenure is a good predictor of our sufficient statistic.

Proposition 13. *In the SPA, a higher winner's tenure implies a higher expected seller's revenue. In the FPA, an increase in the winner's tenure eventually reduces the expected seller's revenue if and only if*

$$\frac{d}{d\theta} b\left(\theta, R^F(\theta), n\right)\Big|_{\theta=1} < 0. \tag{6.2}$$

Figure 6.1 shows how the seller's revenue varies with the tenure of a persistent winner ($\tau \geq 2$) in the two auction formats. Notice that the transfer that a persistent winner makes to the seller depends on his valuation both directly and indirectly, since he "set his own reserve" In the SPA, the payment of a persistent winner with value θ is $e^S(\theta) = b\left(\theta, R^S(\theta), n\right)$. This is an increasing function of θ, since $R^S(\theta)$ is (weakly) increasing in θ and b is increasing in all his arguments. Hence, the expected payment of a winner with tenure τ in the SPA, $E^S(\tau) = \mathbb{E}_{\theta|\tau}\left[e^S(\theta)\right]$, is increasing in τ as a direct consequence of the FOSD order among tenure (Lemma 12).

In the FPA the payment of a persistent winner with value θ is $e^F(\theta) = b\left(\theta, R^F(\theta), n\right)$. Contrary to the SPA, this payment is not necessarily increasing in θ because, in the tailing region, the reserve price is decreasing in θ, which counteracts the direct effect of value on the winner's bid. Condition (6.2) requires that, local to 1, the reserve price effect dominates, so that the payment of a persistent winner decreases with his value.[51] Because the value of a high-tenure winner converges (in probability) to 1, the payment local to the upper bound determines the eventual effect of tenure, yielding the counterintuitive result that higher tenures eventually predict *lower* revenues for the seller (see the right panel of Fig. 6.1).

Following the same logic, we can use the tenure of a persistent winner to predict the probability of trade in the subsequent auction (re-trade).[52] This probability is a decreasing function of the reserve price, which depends on the persistent winner's value θ. Since in the SPA the reserve price is weakly increasing in θ, the probability of

[51] Since $\frac{\mathrm{d}}{\mathrm{d}\theta} b = \frac{\partial}{\partial\theta} b + \frac{\mathrm{d}R}{\mathrm{d}\theta}\frac{\partial}{\partial R} b$ and both partials are positive (Fact 1), the condition $\frac{\mathrm{d}R}{\mathrm{d}\theta}\Big|_{\theta=1} < 0$ (equivalent in the FPA to active tailing) is necessary but not sufficient for (6.2). For example, with two uniform bidders and $\beta = 0$, we have $\frac{\mathrm{d}}{\mathrm{d}\theta} b\Big|_{\theta=1} = -\frac{3-\eta(17\eta-2)}{32\eta^2}$ which is positive for $\eta > 0.48$, while to get tailing it is sufficient that $\eta > \frac{1}{3}$.

[52] It would be meaningless to use winner's tenure to predict trade in the current auction since there is a winner (of any tenure) if and only if the current auction allocated the object. So our exercise is to use winner's tenure to predict the likelihood that the subsequent auction will

re-trade decreases in tenure by Lemma 12. By contrast, in the FPA the reserve price is decreasing in θ in the tailing region and, hence, shifting mass from lower to higher θ within this region increases the probability of re-trade. By Lemma 12, since winners with sufficiently high tenure are almost surely tailed, the probability of re-trade is eventually increasing in tenure.

7 Conclusions

We have analyzed a stylized model of the market for online display advertising, focusing on how the seller adjusts reserve prices using the information conveyed by past bids. The seller auctions identical objects for infinitely many periods, reflecting the high frequency of unique transactions in this market. After a successful auction, the seller distorts the reserve price to extract surplus from the winning bidder, who has revealed his valuation and might bid again in future auctions. This surplus extraction is risky, however, as it requires raising the reserve price above the monopoly price, which is optimal in the event the incumbent leaves. If the probability of this event is high enough, the seller adopts a *tailing* strategy, leaving some surplus to incumbents with high values.

The two auction formats used to sell online display advertisements — the FPA and the SPA — differ sharply in the optimal tailing strategy. This reflects both differences in how each auction format exploits bidders' myopia and a mechanical difference in how the presence of an incumbent interacts with the reserve price in determining the seller's revenue. In the SPA, the incumbent is a substitute for the reserve price, allowing the seller to set a lower reserve price than in the FPA. Moreover, the reserve price in the FPA depends negatively on the incumbent's value, implying that the winners' tenure has counter-intuitive predictions for the seller's revenue and the probability of re-trade. The optimal tailing strategies also determine the minimum reserve price (which is the one following a period without trade) and, overall, drive the differences in the dynamics, trade and revenue across the two auction formats.

Surprisingly, given that myopia induces static overbidding only in the FPA, the SPA yields a higher dynamic value to the seller, unless the incumbents' persistence is extremely high. This is because the SPA better exploits the possible presence of a bidder with a relatively high value — i.e. one drawn from the stationary distribution of incumbents. We also characterize the long-run probability of trade and show that incumbents' persistence affects this probability both directly (since trade can fail only if the incumbent leaves) and indirectly through the response of the seller's reserve price policy. Because of these contrasting effects, trade may be non-monotonic in incumbents' persistence.

The results we derived and the very tractability of our dynamic setting rely on our assumptions about the stochastic process for bidders' preferences and on myopic bidding. We believe that both assumptions are reasonable descriptions of the online display ad market, where many (often small) buyers aim to run user-specific campaigns that require showing their advertisement for a finite number of times, which is unknown to the seller. Recall that myopia is restrictive only in the first period bidders participate in the auction since incumbents possess no residual private information and best respond to their (myopic) new competitors. Therefore, bidders' behavior in our model is consistent with any learning dynamics whose starting point is the static auction — say because the algorithm used by bidders to choose their initial strategy is trained in simulated auctions where different reserve prices simply

allocate the object. A related question is to compute the survival rate of a winner, i.e. the likelihood that he will win the subsequent auction and increase his tenure. Since the survival rate of an incumbent θ is the increasing function $S(\theta) = (1-\eta) F^{n-1}(\theta)$ in either auction format, Lemma 12 yields immediately that the survival rate is increasing in tenure.

reflect different sellers' valuations or outside options.[53]

Relaxing myopic bidding is a natural extension of our analysis. There are two avenues for addressing the tractability challenges presented by this extension. The first is to focus on an empirical investigation of how automatic bidding systems learn and respond to seller discrimination through reserve prices. The recent literature at the intersection of market design and machine learning provides tools to explore this question (e.g., Amin, Rostamizadeh, and Syed (2014) and Golrezaei, Javanmard, and Mirrokni (2019)). The second avenue requires to simplify the model to obtain a tight characterization of the equilibrium even in a context where bidders internalize sellers' discrimination. This fits in a more standard strand of the literature related to the ratchet effect and durable goods markets. For example, Bonatti and Cisternas (2020) recently analyzed a dynamic setting with (myopic) sellers and a single sophisticated buyer who reduces the quantity demanded to hide his valuation and lower discrimination. Relatedly, allowing for competition among multiple buyers, Caillaud and Mezzetti 2004 and Pagnozzi and Sartori (2023) study the optimal reserve price in English auctions where bidders are sophisticated, but restricted to a simpler (two-period) environment than the one analyzed in this paper.

[53]Notice that backward and forward myopia in the first period originate from the same misspecified model (i.e., the static auction with exogenous reserve price); if bidders do not realize that their bids will be used to price discriminate them in the future, there is no reason why they should think that the *current* reserve is used to discriminate — and hence is informative of — a potential competitor, and vice versa.

References

[1] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. 2006. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 1–7.

[2] S Nageeb Ali, Greg Lewis, and Shoshana Vasserman. 2020. Voluntary disclosure and personalized pricing. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 537–538.

[3] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. 2014. Repeated contextual auctions with strategic buyers. *Advances in Neural Information Processing Systems*, 27.

[4] Nick Arnosti, Marissa Beck, and Paul Milgrom. 2016. Adverse selection and auction design for internet display advertising. *American Economic Review*, 106, 10, 2852–2866.

[5] Susan Athey and Glenn Ellison. 2011. Position auctions with consumer search. *The Quarterly Journal of Economics*, 126, 3, 1213–1270.

[6] Benjamin Balzer, Antonio Rosato, and Jonas von Wangenheim. 2022. Dutch vs. first-price auctions with expectations-based loss-averse bidders. *Journal of Economic Theory*, 205, 105545.

[7] Dirk Bergemann, Tibor Heumann, and Stephen Morris. 2023. Bidder-optimal information structures in auctions.

[8] Dirk Bergemann, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter. 2022. Optimal information disclosure in classic auctions. *American Economic Review: Insights*, 4, 3, 371–388.

[9] Dirk Bergemann and Martin Pesendorfer. 2007. Information structures in optimal auctions. *Journal of economic theory*, 137, 1, 580–609.

[10] Jane Black and David De Meza. 1992. Systematic price differences between successive auctionsare no anomaly. *Journal of Economics & Management Strategy*, 1, 4, 607–628.

[11] Alessandro Bonatti, Dirk Bergemann, and Nicholas Wu. 2023. How do digital advertising auctions impact product prices?

[12] Alessandro Bonatti and Gonzalo Cisternas. 2020. Consumer scores and price discrimination. *The Review of Economic Studies*, 87, 2, 750–791.

[13] Niv Buchbinder, Kamal Jain, and Joseph Naor. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *European Symposium on Algorithms*. Springer, 253–264.

[14] Bernard Caillaud and Claudio Mezzetti. 2004. Equilibrium reserve prices in sequential ascending auctions. *Journal of Economic Theory*, 117, 1, 78–95. ISSN: 0022-0531.

[15] Federica Carannante, Marco Pagnozzi, and Elia Sartori. 2024. Interim information and seller's revenue in standard auctions. *Mimeo*.

[16] Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R Devanur. 2011. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1307–1315.

[17] Hana Choi and Carl F. Mela. 2023. Optimizing reserve prices in display advertising auctions. *mimeo*.

[18] Hana Choi, Carl F Mela, Santiago R Balseiro, and Adam Leary. 2020. Online display advertising markets: a literature review and future directions. *Information Systems Research*, 31, 2, 556–575.

[19] Jacques Crémer and Richard P McLean. 1988. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, 1247–1257.

[20] Francesco Decarolis, Maris Goldmanis, and Antonio Penta. 2020. Marketing agencies and collusive bidding in online ad auctions. *Management Science*, 66, 10, 4433–4454.

[21] Francesco Decarolis, Maris Goldmanis, and Antonio Penta. 2018. Recent developments in online ad auctions. *ECONOMIC ANALYSIS OF THE DIGITAL REVOLUTION*, 143.

[22] Francesco Decarolis, Maris Goldmanis, Antonio Penta, and Ksenia Shakhgildyan. 2023. Bid coordination in sponsored search auctions: detection methodology and empirical analysis. *The Journal of Industrial Economics*, 71, 2, 570–592.

[23] Francesco Decarolis and Gabriele Rovigatti. 2021. From mad men to maths men: concentration and buyer power in online advertising. *American Economic Review*, 111, 10, 3299–3327.

[24] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97, 1, 242–259.

[25] Péter Eső and Balazs Szentes. 2007. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74, 3, 705–731.

[26] Tristan Gagnon-Bartsch, Marco Pagnozzi, and Antonio Rosato. 2021. Projection of private values in auctions. *American Economic Review*, 111, 10, 3256–3298.

[27] Ian L Gale and Donald B Hausch. 1994. Bottom-fishing and declining prices in sequential auctions. *Games and Economic Behavior*, 7, 3, 318–331.

[28] Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. 2019. Dynamic incentive-aware learning: robust pricing in contextual auctions. *Advances in Neural Information Processing Systems*, 32.

[29] Donald B Hausch. 1988. A model of sequential auctions. *Economics Letters*, 26, 3, 227–233.

[30] Audrey Hu, Steven A Matthews, and Liang Zou. 2010. Risk aversion and optimal reserve prices in first-and second-price auctions. *Journal of Economic Theory*, 145, 3, 1188–1202.

[31] Patrick Hummel. 2018. Reserve prices in repeated auctions. *International Journal of Game Theory*, 47, 1, 273–299.

[32] Shunto J Kobayashi and Miguel Alcobendas. 2023. Dynamic auctions with budget-constrained bidders: evidence from the online advertising market. *mimeo*.

[33] Dan Levin and James L Smith. 1994. Equilibrium in auctions with entry. *The American Economic Review*, 585–599.

[34] Dan Levin and James L Smith. 1996. Optimal reservation prices in auctions. *The Economic Journal*, 106, 438, 1271–1283.

[35] Richard Luton and R Preston McAfee. 1986. Sequential procurement auctions. *Journal of Public Economics*, 31, 2, 181–195.

[36] R Preston McAfee. 2011. The design of advertising exchanges. *Review of Industrial Organization*, 39, 169–185.

[37] R Preston McAfee, John McMillan, and Philip J Reny. 1989. Extracting the surplus in the common-value auction. *Econometrica: Journal of the Econometric Society*, 1451–1459.

[38] R Preston McAfee and Daniel Vincent. 1993. The declining price anomaly. *Journal of Economic Theory*, 60, 1, 191–212.

[39] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research*, 6, 1, 58–73.

[40] Armando Ortega-Reichert. 1968. *Models for competitive bidding under uncertainty*. Stanford University.

[41] Michael Ostrovsky and Michael Schwarz. 2023. Reserve prices in internet advertising auctions: a field experiment. *Journal of Political Economy*.

[42] Marco Pagnozzi and Elia Sartori. 2023. Rating bidders in repeated auctions. *Mimeo*.

[43] Michael Peters and Sergei Severinov. 1997. Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory*, 75, 1, 141–179.

[44] John G Riley and William F Samuelson. 1981. Optimal auctions. *The American Economic Review*, 71, 3, 381–392.

[45] Hal R Varian. 2007. Position auctions. *International Journal of Industrial Organization*, 25, 6, 1163–1178.

[46] Jun Wang, Weinan Zhang, Shuai Yuan, et al. 2017. Display advertising with real-time bidding (rtb) and behavioural targeting. *Foundations and Trends in Information Retrieval*, 11, 4-5, 297–435.

Appendix

This Appendix contains proofs of all results in the main text. The omitted proofs of additional lemmata and results are in the Online Appendix.

Proof of Facts 1, 2 and 3

Consider the function

$$b(\theta,R,n)=\frac{\int_0^R R\mathrm{d}F(y)^{n-1}+\int_R^\theta y\mathrm{d}F(y)^{n-1}}{F(\theta)^{n-1}}.$$

Partial and cross differentiation yield

$$\frac{\partial}{\partial R}b(\theta,R,n)=\frac{R\mathrm{d}F(R)^{n-1}-R\mathrm{d}F(R)^{n-1}+\int_0^R \mathrm{d}F(y)^{n-1}}{F(\theta)^{n-1}}=\left(\frac{F(R)}{F(\theta)}\right)^{n-1}>0 \tag{7.1}$$

$$\frac{\partial}{\partial \theta}b(\theta,R,n)=\frac{(n-1)F(\theta)^{n-2}f(\theta)\int_R^\theta F(y)^{n-1}\mathrm{d}y}{F(\theta)^{2(n-1)}}=\frac{(n-1)f(\theta)\int_R^\theta F(y)^{n-1}\mathrm{d}y}{F(\theta)^n}>0 \tag{7.2}$$

$$\frac{\partial^2}{\partial R\partial\theta}b(\theta,R,n)=-(n-1)\,f(\theta)\frac{F(R)^{n-1}}{F(\theta)^n}<0 \tag{7.3}$$

$$\frac{\partial^2}{\partial R\partial n}b(\theta,R,n)=\left(\frac{F(R)}{F(\theta)}\right)^{n-1}\log\left(\frac{F(R)}{F(\theta)}\right)<0 \tag{7.4}$$

where the last inequality follows from $\frac{F(R)}{F(\theta)}<1$. The fact that $b(\theta,R,n)>b(\theta,R,n-1)$ follows from the facts that $\max\{y,R\}$ is increasing in y and that the maximum of $n-1$ identical draws stochastically dominates the maximum of $n-2$ identical draws.

Finally, direct differentiation of the objective function

$$\pi_n(R)=\int_R^1 b(\theta,R,n)\,\mathrm{d}F(\theta)^n$$

gives the critical condition

$$-b(R,R,n)\,\mathrm{d}F(R)^n+\int_R^1\frac{\partial}{\partial R}b(\theta,R,n)\,\mathrm{d}F(\theta)^n=0.$$

Since $b(x,x,n)=x$, substituting (7.1) and integrating we obtain

$$\begin{aligned}-R\mathrm{d}F(R)^n+\int_R^1\left(\frac{F(R)}{F(\theta)}\right)^{n-1}\mathrm{d}F(\theta)^n&=-R\mathrm{d}F(R)^n+\int_R^1 nF(R)^{n-1}f(\theta)\,\mathrm{d}\theta\\&=-R\mathrm{d}F(R)^n+nF(R)^{n-1}(1-F(R))\\&=nF(R)^{n-1}(1-F(R)-Rf(R))=-\mathrm{d}F(R)^n\psi(R),\end{aligned}$$

which is satisfied if and only if $\psi(R) = 0$. Notice moreover that we obtain the equivalent rewriting

$$\pi_n(R) = \pi_n(1) - \int_R^1 \frac{\mathrm{d}}{\mathrm{d}x}\pi_n(x)\,\mathrm{d}x = \int_R^1 \psi(x)\,\mathrm{d}F(x)^n. \tag{7.5}$$

Proof of Proposition 4

In either auction format and for all θ, the function $\pi^i(\theta, R)$ has a different specification when $R \le \theta$ and when $R > \theta$, with a downward discontinuity at $R = \theta$. Formally, [54]

$$\pi^F(\theta,R) = \eta\pi_n(R) + (1-\eta)\cdot \begin{cases} F(\theta)^{n-1} b(\theta,R,n) + \int_\theta^1 b(x,R,n)\,\mathrm{d}F(x)^{n-1} & \text{if } R \le \theta \\ \int_R^1 b(x,R,n)\,\mathrm{d}F(x)^{n-1} & \text{if } R > \theta \end{cases} \tag{7.6}$$

$$\pi^S(\theta,R) = \eta\pi_n(R) + (1-\eta)\cdot \begin{cases} F(\theta)^{n-1} b(\theta,R,n) + \int_\theta^1 b(x,\theta,n-1)\,\mathrm{d}F(x)^{n-1} & \text{if } R \le \theta \\ \int_R^1 b(x,R,n-1)\,\mathrm{d}F(x)^{n-1} & \text{if } R > \theta \end{cases} \tag{7.7}$$

Throughout, let $\pi^{+,i}(\theta,R)$ and $\pi^{-,i}(R)$ be the the upper and lower branch of $\pi^i(\theta,R)$ respectively,[55] and let $\widetilde{\pi}^i(\theta) = \pi^i(\theta,\theta)$ be the revenue with tracking and, hence, $R = \theta$. Notice the following

Lemma 14. *The tracking revenues satisfy* $\widetilde{\pi}^S(\theta) < \widetilde{\pi}^F(\theta)$ *but* $\frac{d}{d\theta}\widetilde{\pi}^S(\theta) > \frac{d}{d\theta}\widetilde{\pi}^F(\theta)$.

Moreover, let $\Pi^{-,i} = \max_R \pi^{-,i}(R)$ and $R^{-,i}$ be the maximizer. Finally, let $R^{i,+}(\theta)$ be the maximizer of $\pi^{+,i}(\theta,R)$. Because $\pi^{+,i}(\theta,R) \ge \pi^{-,i}(R)$, whenever feasible $R^{+,i}(\theta)$ also maximizes $\pi^i(\theta,R)$, that is $R^{i,+}(\theta) \le \theta \Rightarrow R^{i,+}(\theta) = \arg\max \pi^i(\theta,R)$. Otherwise, to solve the seller's problem we need to compare $\Pi^{-,i}$ with $\widetilde{\pi}^i(\theta)$.

We begin by characterizing the tailing reserve price $R^{i,+}(\theta)$ and the region where it is feasible. In the SPA, $R^{S,+}(\theta)$ is obtained by setting to zero the following function:

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}R}\pi^{+,S}(\theta,R) &= \frac{\mathrm{d}}{\mathrm{d}R}\left[\eta\pi_n(R) + (1-\eta)\left(F(\theta)^{n-1} b(\theta,R,n) + \int_\theta^1 b(x,\theta,n-1)\,\mathrm{d}F(x)^{n-1}\right)\right] \\
&= \eta\frac{\mathrm{d}}{\mathrm{d}R}\pi_n(R) + (1-\eta)F(\theta)^{n-1}\frac{\mathrm{d}}{\mathrm{d}R}b(\theta,R,n) \\
&= F(R)^{n-1}\left(\eta n(1-F(R)-Rf(R)) + (1-\eta)\right) \\
&= F(R)^{n-1}\left((1-\eta) - \eta n f(R)\psi(R)\right).
\end{aligned} \tag{7.8}$$

This yields the condition $\psi\left(R^{S,+}(\theta)\right) = \frac{1-\eta}{\eta n f\left(R^{S,+}(\theta)\right)}$, which is independent of θ and feasible if and only if $\theta \ge R^{S,+}$.

In the FPA, we preliminary show that $R^{F,+}(\theta)$ is decreasing in θ by proving that the objective $\pi^{+,F}(\theta,R)$ is submodular and appealing to the Topkis theorem. We have

$$\begin{aligned}
\frac{\partial^2}{\partial R\partial\theta}\pi^{+,F}(\theta,R) &= (1-\eta)\left(\frac{\partial^2 b(\theta,R)}{\partial\theta\partial R}\cdot F(\theta)^{n-1} + \frac{\partial b(\theta,R)}{\partial R}\cdot \mathrm{d}F(\theta)^{n-1} - \frac{\partial b(\theta,R)}{\partial R}\cdot \mathrm{d}F(\theta)^{n-1}\right) \\
&\propto \frac{\partial^2 b(\theta,R)}{\partial\theta\partial R},
\end{aligned}$$

[54] The functions $\pi^i_{\theta,n-1}(R)$ in the main text are the expressions in curly brackets. The expression for the SPA uses the fact that the bid by an incumbent θ is equivalent to a reserve price and hence the expected payment of a new bidder with value x can be written as $b(x,\theta,n-1)$ (see Section 2.2).

[55] For example, $\pi^{+,F}(\theta,R) = \eta\pi_n(R) + (1-\eta)\left[F(\theta)^{n-1} b(\theta,R,n) + \int_\theta^1 b(x,R,n)\,\mathrm{d}F(x)^{n-1}\right]$. Notice that $\pi^{-,i}$ does not depend on the incumbent's value because it presumes that the incumbent is excluded.

which is negative per Fact 2. Analytically, $R^{F,+}(\theta)$ is obtained by setting the following derivative equal to zero:

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}R}\pi^{+,F}(\theta,R) &= \frac{\mathrm{d}}{\mathrm{d}R}\left[\eta\pi_n(R)+(1-\eta)\left(F(\theta)^{n-1}b(\theta,R,n)+\int_\theta^1 b(x,R,n)\,\mathrm{d}F(x)^{n-1}\right)\right]\\ &= F(R)^{n-1}\left((1-\eta)-\eta n f(R)\psi(R)\right)+(1-\eta)\int_\theta^1 \frac{\mathrm{d}}{\mathrm{d}R}b(x,R,n)\,\mathrm{d}F(x)^{n-1}\\ &= F(R)^{n-1}\left((1-\eta)-\eta n f(R)\psi(R)+(1-\eta)\int_\theta^1\frac{\mathrm{d}F(x)^{n-1}}{F(x)^{n-1}}\right)\\ &= F(R)^{n-1}\left((1-\eta)-\eta n f(R)\psi(R)-(1-\eta)(n-1)\log(F(\theta))\right).\end{aligned} \tag{7.9}$$

Rearranging yields the expression in the statement.

The feasibility condition $R^{F,+}(\theta)\le\theta$ is satisfied as long as $\theta\ge\overline{\theta}^F$ such that

$$(1-\eta)-\eta n f\left(\overline{\theta}^F\right)\psi\left(\overline{\theta}^F\right)-(1-\eta)(n-1)\log\left(F\left(\overline{\theta}^F\right)\right).$$

By concavity, whenever $R^{F,+}(\theta)$ is not feasible then the profit function conditional on no exclusion is decreasing in the relevant domain $R\ge\theta$, so the seller compares the tracking revenue $\widetilde{\pi}^i$ with the exclusion revenue Π^{-i}.

We now proceed to characterize $R^{-,i}$. For the SPA, $\pi^{-,s}(R)=\eta\pi_n(R)+(1-\eta)\pi_{n-1}(R)$ is a convex combination of functions which — per Fact 3 — are both maximized at r^M. It then follows that $R^{-,s}=r^M$ and $\Pi^{-,s}=\eta\pi_n\left(r^M\right)+(1-\eta)\pi_{n-1}\left(r^M\right)$.

Now, $\underline{\theta}^i$ solves $\Pi^{-,i}=\widetilde{\pi}^i\left(\underline{\theta}^i\right)$. For the SPA, we get

$$\eta\pi_n\left(r^M\right)+(1-\eta)\pi_{n-1}\left(r^M\right)=\eta\pi_n\left(\underline{\theta}^s\right)+(1-\eta)\pi_{n-1}\left(\underline{\theta}^s\right)+(1-\eta)F\left(\underline{\theta}^s\right)^{n-1}\underline{\theta}^s,$$

$$\Leftrightarrow\qquad \eta\left[\pi_n\left(r^M\right)-\pi_n\left(\underline{\theta}^s\right)\right]+(1-\eta)\left[\pi_{n-1}\left(r^M\right)-\pi_{n-1}\left(\underline{\theta}^s\right)\right]=(1-\eta)F\left(\underline{\theta}^s\right)^{n-1}\underline{\theta}^s. \tag{7.10}$$

This condition has a clear economic interpretation: since — per Fact 3 — $\pi_n\left(r^M\right)\ge\pi_n(x)$ for every n and x, the seller sets a sub-optimally low reserve in exchange for the insurance to trade because, in the event that none of the $(n-1)$ standard bidders bids above the reserve price $\underline{\theta}^s$, the (surviving) tracked incumbent guarantees revenue $\underline{\theta}^s$.

Notice that the LHS of (7.10) is decreasing (as a function of $\underline{\theta}^S$) because $\pi_n(x)$ and $\pi_{n-1}(x)$ are both increasing in $\left[0,r^M\right]$, while the RHS is increasing. Moreover, the LHS is positive when $\underline{\theta}^S=0$ and equal to 0 when $\underline{\theta}^S=r^M$, while the RHS is equal to 0 when $\underline{\theta}^S=0$ and positive when $\underline{\theta}^S=r^M$. Hence, they cross in a unique point in $\left[0,r^M\right]$, implying that $\underline{\theta}^S<r^M$.

Similar arguments and computations also hold in the FPA and yield the optimal reserve price $R^{-,F}<r^M$ and the exclusion threshold $\underline{\theta}^F$. For the former notice that $\underline{R}^F$ solves

$$\max_R \eta\int_R^1\psi(x)\mathrm{d}F(x)^n+(1-\eta)\int_R^1 b(x,R,n)\mathrm{d}F(x)^{n-1}$$

and the FOC gives

$$-\eta\psi(R)\mathrm{d}F(R)^{n}+(1-\eta)\left[-R\mathrm{d}F(R)^{n-1}-\int_{R}^{1}\frac{\mathrm{d}b(x,R,n)}{\mathrm{d}R}\mathrm{d}F(x)^{n-1}\right]>$$
$$-\eta\psi(R)\mathrm{d}F(R)^{n}+(1-\eta)\left[-R\mathrm{d}F(R)^{n-1}-\int_{R}^{1}\frac{\mathrm{d}b(x,R,n-1)}{\mathrm{d}R}\mathrm{d}F(x)^{n-1}\right]=$$
$$-\psi(R)\left(\eta\mathrm{d}F(R)^{n}+(1-\eta)\mathrm{d}F(R)^{n-1}\right),$$

where the inequality uses (7.4). This means that the FOC at r^{M} is positive, hence $\underline{R}^{F}<r^{M}$. The same manipulations as for the SPA yield that $\underline{\theta}^{F}$ solves

$$\eta\left[\pi_{n}\left(\underline{R}^{F}\right)-\pi_{n}\left(\underline{\theta}^{F}\right)\right]+(1-\eta)\left[\pi_{n-1}\left(\underline{R}^{F}\right)-\pi_{n-1}\left(\underline{\theta}^{F}\right)\right]=(1-\eta)F\left(\underline{\theta}^{F}\right)^{n-1}\underline{\theta}^{F}. \tag{7.11}$$

Using the same argument as for the SPA — i.e., LHS decreasing to 0 in $\left[0,\underline{R}^{F}\right]$ and RHS increasing from 0 in $\left[0,\underline{R}^{F}\right]$ — yields that there is a unique solution to (7.11) in $\left[0,\underline{R}^{F}\right]$.

Proof Proposition 5

Tracking across Auction Formats We start by showing that $\underline{\theta}^{F}<\underline{\theta}^{S}<r^{M}<\overline{\theta}^{S}<\overline{\theta}^{F}$. If $\overline{\theta}^{S}<1$, then $\overline{\theta}^{S}=R^{F}(1)<\overline{\theta}^{F}$ because the tailing reserve is strictly decreasing in the FPA. The fact that $\overline{\theta}^{S}>r^{M}$ follows from $\psi\left(\overline{\theta}^{S}\right)=\frac{1-\eta}{\eta n f\left(\overline{\theta}^{S}\right)}>0=\psi\left(r^{M}\right)$. The fact that $\underline{\theta}^{S}<r^{M}$ was established in Proposition 4.

Finally, notice from equation (7.10) and (7.11) that $\underline{\theta}^{i}$ solves $g^{i}(x)=h(x)$, where $h(x)=(1-\eta)F(x)^{n-1}x$, which is increasing and $g^{i}(x)=\eta\left[\pi_{n}\left(\underline{R}^{i}\right)-\pi_{n}(x)\right]+(1-\eta)\left[\pi_{n-1}\left(\underline{R}^{i}\right)-\pi_{n-1}(x)\right]$ which is decreasing and such that

$$g^{S}(x)-g^{F}(x)=\eta\left[\pi_{n}\left(\underline{R}^{F}\right)-\pi_{n}\left(r^{M}\right)\right]+(1-\eta)\left[\pi_{n-1}\left(\underline{R}^{F}\right)-\pi_{n-1}\left(r^{M}\right)\right],$$

which is negative by Fact 3. The claim then follows.

Comparative Statics of Tracking Recall that in the SPA $\overline{\theta}^{S}$ solves

$$g(x)=\frac{1-\eta}{n\eta},$$

where we assume that $g(x)=\psi(x)f(x)$ is increasing in $[0,1]$. Then

$$\frac{\mathrm{d}\overline{\theta}^{S}}{\mathrm{d}\eta}\propto\frac{\mathrm{d}}{\mathrm{d}\eta}\frac{1-\eta}{n\eta}<0$$

and

$$\frac{\mathrm{d}\overline{\theta}^{S}}{\mathrm{d}n}\propto\frac{\mathrm{d}}{\mathrm{d}n}\frac{1-\eta}{n\eta}<0.$$

For the FPA, $\overline{\theta}^{F}$ solves

$$g(x)=\frac{(1-\eta)}{\eta n}\left[1-(n-1)\log(F(x))\right],$$

where the RHS is positive but decreasing in x. To show the result, it is then sufficient to notice the RHS shifts down (as a function of x) because

$$\frac{\mathrm{d}}{\mathrm{d}\eta}\frac{(1-\eta)}{\eta n}\left[1-(n-1)\log\left(F\left(x\right)\right)\right] \propto \frac{\mathrm{d}}{\mathrm{d}\eta}\frac{(1-\eta)}{\eta n} < 0.$$

Full Tracking The threshold $\overline{\theta}^{i}$ solves $g(x) = z^{i}(x)$, where g is (strictly) increasing and z^{i} is weakly decreasing (constant if $i = S$). A necessary and sufficient condition for there to be a (unique) solution is that $g(1) = f(1) > z^{i}(1) = \frac{1-\eta}{n\eta}$, which is equivalent to

$$\eta > \overline{\eta} = \frac{1}{f(1)\,n+1}.$$

The threshold $\overline{\eta}$ is non-trivial if $f(1) > 0$. The idea is that there must be positive density around 1, otherwise it is suboptimal to track at the upper bound.

Limit Results Since $\lim_{\eta\to 1}\frac{1-\eta}{n\eta} = 0$ then $g\left(\overline{\theta}^{i}\right) \to 0$ for both $i \in \{F, S\}$. The fact that $\overline{\theta}^{i} \to r^{M}$ then follows from continuity of g. Likewise, as $\lim_{n\to\infty}\frac{1-\eta}{n\eta} = 0$, we also get $\overline{\theta}^{S} \to r^{M}$ as $n \to \infty$. On the contrary, $\lim_{n\to\infty} z^{F}(x) = -\frac{1-\eta}{\eta}\log\left(F(x)\right)$ and, from continuity of both functions, we get that $\overline{\theta}^{F} \to \overline{\theta}^{F,\infty}$ where $\overline{\theta}^{F,\infty} > r^{M}$ is such that

$$\psi\left(\overline{\theta}^{F,\infty}\right) f\left(\overline{\theta}^{F,\infty}\right) = -\frac{1-\eta}{\eta}\log\left(F\left(\overline{\theta}^{F,\infty}\right)\right).$$

For all θ, $R^{F}(\theta) \to R^{F,\infty}(\theta)$ that solves

$$\psi\left(R^{F,\infty}(\theta)\right) f\left(R^{F,\infty}(\theta)\right) = -\frac{(1-\eta)}{\eta}\log\left(F(R^{F,\infty}(\theta))\right).$$

Proof Proposition 6

First we show that, if $\theta \leq \overline{\theta}^{S}$, then $\Pi^{S}(\theta) < \Pi^{F}(\theta)$. Since $\underline{\theta}^{F} \leq \underline{\theta}^{S} \leq \overline{\theta}^{S}$ by Proposition 5, we distinguish three cases. First, if $\theta \leq \underline{\theta}^{F}$, then $\pi^{-,F}(R) - \pi^{-,S}(R) = \int_{R}^{1} b(x,R,n) - b(x,R,n-1)\,\mathrm{d}F(x)^{n-1} > 0$ for all R so that $\Pi^{-,F} = \max_{R}\pi^{-,F}(R) > \max_{R}\pi^{-,S}(R) = \Pi^{-,S}$. Second, if $\theta \in \left(\underline{\theta}^{F}, \underline{\theta}^{S}\right)$, then the FPA tracks but the SPA excludes and $\Pi^{F}(\theta) = \widetilde{\pi}^{F}(\theta) > \Pi^{-,F} > \Pi^{-,S} = \Pi^{S}(\theta)$. Third, if $\theta \in (\underline{\theta}^{S}, \overline{\theta}^{S})$, then both auction formats track and — per Fact 14 — $\Pi^{F}(\theta) = \widetilde{\pi}^{F}(\theta) > \widetilde{\pi}^{S}(\theta) = \Pi^{S}(\theta)$.

Next we show that, if $\overline{\theta}^{S} < 1$, then the SPA dominates in a neighborhood of 1and the FPA and SPA yield the same revenue when $\theta = 1$. Notice that

$$\pi^{+,F}(\theta,R) - \pi^{+,S}(\theta,R) = (1-\eta)\left[\int_{\theta}^{1}\left(b(x,R,n) - b(x,\theta,n-1)\right)\mathrm{d}F(x)^{n-1}\right],$$

which implies that $\pi^{+,F}(1,R) = \pi^{+,S}(1,R)$ for all R. Therefore, the objective functions in the FPA and SPA coincide when $\theta = 1$ and, hence, they have the same maximizer $\overline{\theta}^{S} = \overline{R}^{S} = \overline{R}^{F}(1)$ (which is feasible by the assumption $\overline{\theta}^{S} < 1$) and the same maximum $\Pi^{S}(1) = \Pi^{F}(1)$.

Since $\overline{\theta}^{S} < 1 \Rightarrow \overline{\theta}^{F} < 1$, the interval $[\overline{\theta}^{F}, 1]$ is non-empty. Suppose that θ belongs to such interval. Then θ is tailed in both the FPA and SPA and $\Pi^i(\theta) = \pi^{+,i}\left(\theta, R^{+,i}(\theta)\right)$ for $i \in \{S, F\}$. By the envelope theorem,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\theta}\Pi^F(\theta) &= (1-\eta)\frac{\partial}{\partial\theta}\left[F(\theta)^{n-1} b\left(\theta, R^{+,F}(\theta), n\right) + \int_\theta^1 b\left(x, R^{+,F}(\theta), n\right)\mathrm{d}F(x)^{n-1}\right] \\ &= (1-\eta)\frac{\partial}{\partial\theta} b(\theta, R^{+,F}(\theta), n) F(\theta)^{n-1} > 0,\end{aligned}$$

and

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\theta}\Pi^S(\theta) &= (1-\eta)\frac{\partial}{\partial\theta}\left[F(\theta)^{n-1} b\left(\theta, R^{+,S}, n\right) + \int_\theta^1 b(x,\theta,n-1)\,\mathrm{d}F(x)^{n-1}\right] \\ &= (1-\eta)\frac{\partial}{\partial\theta}\left[\int_0^{R^{+,S}} R^{+,S}\mathrm{d}F(x)^{n-1} + \int_{R^{+,S}}^{\theta} x\mathrm{d}F(x)^{n-1} + \int_\theta^1 b(x,\theta,n-1)\,\mathrm{d}F(x)^{n-1}\right] \\ &= (1-\eta)\left[\theta\mathrm{d}F(\theta)^{n-1} - b(\theta,\theta,n-1)\mathrm{d}F(\theta)^{n-1} + \int_\theta^1 \frac{\partial}{\partial R} b(x,\theta,n-1)\mathrm{d}F(x)^{n-1}\right] \\ &= (1-\eta)\int_\theta^1 \left(\frac{F(\theta)}{F(x)}\right)^{n-2}\mathrm{d}F(x)^{n-1} \\ &= (1-\eta)(n-1)F(\theta)^{n-2}(1-F(\theta)).\end{aligned}$$

It follows that $\lim_{\theta\to 1}\frac{\mathrm{d}}{\mathrm{d}\theta}\Pi^S(\theta) = 0 < \lim_{\theta\to 1}\frac{\mathrm{d}}{\mathrm{d}\theta}\Pi^F(\theta)$. Moreover,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\theta}\Pi^F(\theta) > \frac{\mathrm{d}}{\mathrm{d}\theta}\Pi^S(\theta) \quad &\Leftrightarrow \quad \frac{\partial}{\partial\theta} b(\theta, R^{+,F}(\theta), n) F(\theta) > (n-1)(1-F(\theta)) \\ &\Leftrightarrow \quad f(\theta)\int_{R^{+,F}(\theta)}^{\theta} F(y)^{n-1}\mathrm{d}y > (1-F(\theta))F(\theta)^{n-1} \\ &\Leftrightarrow \quad \int_{R^{+,F}(\theta)}^{\theta}\left(\frac{F(y)}{F(\theta)}\right)^{n-1}\mathrm{d}y > \frac{(1-F(\theta))}{f(\theta)},\end{aligned} \tag{7.12}$$

which is satisfied (at least) in a left neighborhood of 1 (because when $\theta = 1$, $R^{+,F}(1) < 1$ and hence the LHS of condition (7.12) is positive, and the RHS is equal to 0). Integrating backwards, for any θ in this neighborhood,

$$\begin{aligned}\Pi^S(\theta) - \Pi^F(\theta) &= \Pi^S(1) - \int_\theta^1 \frac{\mathrm{d}}{\mathrm{d}\theta'}\Pi^S(\theta') - (\Pi^F(1) - \frac{\mathrm{d}}{\mathrm{d}\theta'}\Pi^F(\theta')\,\mathrm{d}\theta') \\ &= \int_\theta^1 \frac{\mathrm{d}}{\mathrm{d}\theta'}\Pi^F(\theta') - \frac{\mathrm{d}}{\mathrm{d}\theta'}\Pi^S(\theta')\,\mathrm{d}\theta' > 0.\end{aligned}$$

Proof of Theorem 7

We can write the seller's problem in recursive form where the state is the winner of the previous period auction, the incumbent. The initial value is given by (3.1). Accounting for the possibility that $R > \theta$, the value with incumbent θ is

$$V^i(\theta) = \max\left\{V^{-,i}, V^{+,i}(\theta)\right\},$$

where

$$V^{-,i} = \max_R \pi^{-,S}(R) + \beta \left\{ \eta \left(F(R)^n V_\emptyset^i + \int_R^1 V^i(y)\, \mathrm{d}F(y)^n \right) + (1-\eta) \left(F(R)^{n-1} V_\emptyset^i + \int_R^1 V^i(y)\, \mathrm{d}F(y)^{n-1} \right) \right\} \tag{7.13}$$

is the value conditional on exclusion[56] and

$$V^{+,i}(\theta) = \max_{R \le \theta} \pi^{+,i}(\theta, R) + \beta \left\{ \eta \left(F(R)^n V_\emptyset^i + \int_R^1 V^i(y)\, \mathrm{d}F(y)^n \right) + (1-\eta) \left(F(\theta)^{n-1} V^i(\theta) + \int_\theta^1 V^i(y)\, \mathrm{d}F(y)^{n-1} \right) \right\}. \tag{7.14}$$

Lemma 15. *$V^i(\theta)$ is increasing in θ.*

The rest of the proof is organized as follows. First, Lemmas 16 and 17 derive conditions (4.3), (4.4), (4.5) by direct differentiation of the value functions.[57] Second, Lemmas 18 and 19 establish that, on path, the seller does not exclude any incumbent and, hence, that $V^i(\theta) = V^{+,i}(\theta)$. Recall that in the exposition we have implicitly assumed that this is the case — see (4.1).

Lemma 16. *The initial reserve price $R_\emptyset^i$ solves $\psi\left(R_\emptyset^i\right) = -\beta\left(V^i\left(R_\emptyset^i\right) - V_\emptyset^i\right)$.*

Lemma 17. *If $R^i(\theta) < \theta$, then $R^i(\theta)$ satisfies conditions (4.4)-(4.5). $R^S(\theta)$ is flat and $R^F(\theta)$ is decreasing in θ.*

Finally, we prove that the seller does not exclude any incumbent by showing that the optimal reserve price conditional on exclusion is weakly below $R_\emptyset^i$ (Lemma 18) and that the optimal reserve price conditional on tailing is higher than $R_\emptyset^i$ (Lemma 19), which makes it impossible to have an incumbent with value lower than $R_\emptyset^i$.

Lemma 18. *Let R_{Ex}^i be the optimal reserve price conditional on excluding an incumbent. It holds that $R_{Ex}^S = R_\emptyset^S$ and $R_{Ex}^F < R_\emptyset^F$*

Proof. R_{Ex}^i solves

$$\max_R \pi^{-,i}(R) + \beta \left\{ \eta \left(F(R)^n V_\emptyset^i + \int_R^1 V^i(y)\, \mathrm{d}F(y)^n \right) + (1-\eta) \left(F(R)^{n-1} V_\emptyset^i + \int_R^1 V^i(y)\, \mathrm{d}F(y)^{n-1} \right) \right\}.$$

For the SPA, the FOC of this problem is

$$\frac{\mathrm{d}}{\mathrm{d}R}\pi^{-,S} + \beta\left(V_\emptyset^S - V^S(R)\right)\left[\eta \mathrm{d}F(R)^n + (1-\eta)\,\mathrm{d}F(R)^{n-1}\right] = 0$$

$$\Leftrightarrow \quad \left(\beta\left(V_\emptyset^S - V^S(R)\right) - \psi(R)\right)\left[\eta \mathrm{d}F(R)^n + (1-\eta)\,\mathrm{d}F(R)^{n-1}\right] = 0.$$

Since the second factor is never 0, this condition is equivalent to (4.3), yielding that $R_{Ex}^S = R_\emptyset^S$.

For the FPA, the FOC is

$$\zeta(R) + \left(\beta\left(V_\emptyset^F - V^F(R)\right) - \psi(R)\right)\left[\eta \mathrm{d}F(R)^n + (1-\eta)\,\mathrm{d}F(R)^{n-1}\right] = 0,$$

[56] Notice it is immaterial to restrict to cases in which the reserve price actually excludes ($R < \theta$) because in that case $V^{-,i} < V^{+,i}(\theta)$.

[57] All the expressions for the static marginal are from the proof of Proposition 4.

where

$$\zeta(R) := \frac{\mathrm{d}}{\mathrm{d}R}\pi^{-,F} - \frac{\mathrm{d}}{\mathrm{d}R}\pi^{-,S} = (1-\eta)\frac{\mathrm{d}}{\mathrm{d}R}\left[\int_R^1 b(x,R,n)\,\mathrm{d}F(x)^{n-1} - \int_R^1 b(x,R,n-1)\,\mathrm{d}F(x)^{n-1}\right]$$
$$= (1-\eta)\int_R^1 \underbrace{\frac{\mathrm{d}}{\mathrm{d}R}b(x,R,n) - \frac{\mathrm{d}}{\mathrm{d}R}b(x,R,n-1)}_{<0 \text{ by } (7.4)}\mathrm{d}F(x)^{n-1} < 0.$$

Since $\beta\left(V_\emptyset^F - V^F\left(R_\emptyset^F\right)\right) = \psi\left(R_\emptyset^F\right)$ by Lemma 16, the FOC evaluated at $R_\emptyset^F$ is negative. Therefore, the optimum R_{Ex}^F must be lower than $R_\emptyset^F$. □

Lemma 19. *Let R^i be a optimal reserve price conditional on tailing an incumbent. It holds that $R^i > R_\emptyset^i$. Moreover, $V^i\left(R^i\right) > V_\emptyset^i$.*

Proof. For the SPA, the (unique) tailing reserve $\overline{R}^S$ solves (7.33). Substituting (7.8) and rearranging yields

$$F(R)^{n-1}\left((1-\eta) - \eta n f(R)\psi(R)\right) + \beta\eta \mathrm{d}F(R)^n\left(V_\emptyset^i - V^i(R)\right) = 0$$

$$F(R)^{n-1}(1-\eta) + \eta \mathrm{d}F(R)^n\left[\beta\left(V_\emptyset^i - V^i(R)\right) - \psi(R)\right] = 0. \tag{7.15}$$

By Lemma 16, the second addendum is zero at $R_\emptyset^S$ which implies that $\overline{R}^S$ must be higher than $R_\emptyset^S$.

For the FPA, since $R^F(\theta)$ is decreasing (Lemma 17), it is sufficient to show that $R^F(1) > R_\emptyset^F$. Because $\pi^{+,F}(1,R) = \pi^{+,S}(1,R)$, (7.15) also characterizes $R^F(1)$ and the same argument used for the SPA establishes the desideratum. Finally, it is straightforward to show that $V^i\left(R^i\right) > V_\emptyset$. □

Proof of Proposition 8

We first show that, local to $\beta = 0$ or $\eta = 1$, $R_\emptyset^i > r^M \Leftrightarrow \widetilde{\pi}_{n-1}^i\left(r^M\right) < \pi_n\left(r^M\right) \Leftrightarrow n$ is large. Since $R_\emptyset^i = r^M$ when $\beta = 0$ or $\eta = 1$ (by Lemma 16), this is equivalent to proving that the derivatives of $R_\emptyset^i$ w.r.t. β local to $\beta = 0$ and w.r.t. η local to $\eta = 1$, evaluated at r^M, are positive when condition (4.6) is met.

Differentiating w.r.t. β the characterizing equation — Lemma 16 — for $R_\emptyset^i$ we get

$$\left[\psi'\left(R_\emptyset^i\right) + \beta\frac{\mathrm{d}}{\mathrm{d}\theta}V^i\left(R_\emptyset^i\right)\right]\frac{\mathrm{d}}{\mathrm{d}\beta}R_\emptyset^i = -\left(V^i\left(R_\emptyset^i\right) - V_\emptyset^i\right) - \beta\frac{\partial}{\partial\beta}\left(V^i\left(R_\emptyset^i\right) - V_\emptyset^i\right).$$

Lemma 15 and $\psi' > 0$ yield that

$$\frac{\mathrm{d}}{\mathrm{d}\beta}R_\emptyset^i \propto -\left(V^i\left(R_\emptyset^i\right) - V_\emptyset^i\right) - \beta\frac{\partial}{\partial\beta}\left(V^i\left(R_\emptyset^i\right) - V_\emptyset^i\right).$$

Recall that at $\beta = 0$, $R_\emptyset^i = r^M$ and the seller's values are equal to the static profits. Therefore,

$$\left.\frac{\mathrm{d}}{\mathrm{d}\beta}R_\emptyset^i\right|_{\beta=0} \propto \widetilde{\pi}_{n-1}^i\left(r^M\right) - \pi_n\left(r^M\right), \tag{7.16}$$

which implies that $\left.\frac{\mathrm{d}}{\mathrm{d}\beta}R_\emptyset^i\right|_{\beta=0} > 0$ if and only if $\widetilde{\pi}_{n-1}^i\left(r^M\right) - \pi_n\left(r^M\right) < 0$.[58]

[58] Still need to justify the claim (in the text) that this condition is met for n small enough, namely that $\widetilde{\pi}_{n-1}^i(r^M) - \pi_n(r^M)$ is decreasing

Differentiating the expression in Lemma 16 w.r.t. η instead yields

$$\left[\psi'\left(R^i_\emptyset\right)+\beta\frac{\mathrm{d}}{\mathrm{d}\theta}V^i\left(R^i_\emptyset\right)\right]\frac{\mathrm{d}}{\mathrm{d}\eta}R^i_\emptyset=-\beta\frac{\partial}{\partial\eta}\left(V^i\left(R^i_\emptyset\right)-V^i_\emptyset\right).$$

By the same argument used above for the derivative w.r.t. β,

$$\frac{\mathrm{d}}{\mathrm{d}\eta}R^i_\emptyset\propto-\frac{\partial}{\partial\eta}\left(V^i\left(R^i_\emptyset\right)-V^i_\emptyset\right). \tag{7.17}$$

By the envelope theorem,

$$\begin{aligned}\frac{\partial}{\partial\eta}V^i(R^i_\emptyset)=&\ \pi_n\left(R^i_\emptyset\right)-\widetilde{\pi}^i_{n-1}\left(R^i_\emptyset\right)+\beta\left(F\left(R^i_\emptyset\right)^nV^i_\emptyset+\int_{R^i_\emptyset}^1V^i\left(\theta'\right)\mathrm{d}F\left(\theta'\right)^n\right)\\&-\beta\left(F\left(R^i_\emptyset\right)^{n-1}V^i\left(R^i_\emptyset\right)+\int_{R^i_\emptyset}^1V^i\left(\theta'\right)\mathrm{d}F\left(\theta'\right)^{n-1}-\eta\left(F\left(R^i_\emptyset\right)^n\frac{\partial}{\partial\eta}V^i_\emptyset+\int_{R^i_\emptyset}^1\frac{\partial}{\partial\eta}V^i\left(y\right)\mathrm{d}F\left(y\right)^n\right)\right)\\&+\beta\left(1-\eta\right)\left(F\left(R^i_\emptyset\right)^{n-1}\frac{\partial}{\partial\eta}V^i_\emptyset+\int_{R^i_\emptyset}^1\frac{\partial}{\partial\eta}V^i\left(y\right)\mathrm{d}F\left(y\right)^{n-1}\right).\end{aligned} \tag{7.18}$$

At $\eta=1$, $V^i_\emptyset=V^i(\theta)=\frac{\pi_n(r^M)}{1-\beta}$ for all θ and $R^i_\emptyset=r^M$. So the previous expression simplifies to

$$\begin{aligned}\frac{\partial}{\partial\eta}V^i(R^i_\emptyset)=&\ \frac{\partial}{\partial\eta}\widetilde{\pi}^i\left(R^i_\emptyset\right)+\beta\frac{\pi_n\left(r^M\right)}{1-\beta}\left(F\left(R^i_\emptyset\right)^n+\int_{R^i_\emptyset}^1\mathrm{d}F\left(\theta'\right)^n-F\left(R^i_\emptyset\right)^{n-1}-\int_{R^i_\emptyset}^1\mathrm{d}F\left(\theta'\right)^{n-1}\right)\\&+\beta\left(F\left(R^i_\emptyset\right)^n\frac{\partial}{\partial\eta}V^i_\emptyset+\int_{R^i_\emptyset}^1\frac{\partial}{\partial\eta}V^i\left(y\right)\mathrm{d}F\left(y\right)^n\right).\end{aligned}$$

Moreover,

$$\frac{\partial}{\partial\eta}V^i_\emptyset=\beta\left(F\left(R^i_\emptyset\right)^n\frac{\partial}{\partial\eta}V^i_\emptyset+\int_{R^i_\emptyset}^1\frac{\partial}{\partial\eta}V^i\left(y\right)\mathrm{d}F\left(y\right)^n\right). \tag{7.19}$$

Therefore, at $\eta=1$,

$$\begin{aligned}\left.\frac{\mathrm{d}}{\mathrm{d}\eta}R^i_\emptyset\right|_{\eta=1}\propto&-\left[\pi_n\left(r^M\right)-\widetilde{\pi}^i_{n-1}\left(r^M\right)\right]\\&+\beta\frac{\pi_n\left(r^M\right)}{1-\beta}\left(F\left(r^M\right)^{n-1}+1-F\left(r^M\right)^{n-1}-F\left(r^M\right)^n-1+F\left(r^M\right)^n\right)\\\propto&\ \widetilde{\pi}^i_{n-1}\left(r^M\right)-\pi_n\left(r^M\right),\end{aligned} \tag{7.20}$$

which establishes the claim.

in n.

In addition, at $\eta = 0$, $\frac{\mathrm{d}}{\mathrm{d}\eta} R^i_\emptyset \propto \widetilde{\pi}^i_{n-1}\left(R^i_\emptyset\right) - \pi_n\left(R^i_\emptyset\right)$, indeed from (7.17) $\frac{\mathrm{d}}{\mathrm{d}\eta} R^i_\emptyset \propto -\frac{\partial}{\partial\eta}\left(V^i\left(R^i_\emptyset\right) - V^i_\emptyset\right)$. At $\eta = 0$,

$$\begin{aligned}
\frac{\partial}{\partial\eta}\left(V^i\left(R^i_\emptyset\right) - V^i_\emptyset\right)\Big|_{\eta=0} &\propto \pi_n\left(R^i_\emptyset\right) - \widetilde{\pi}^i_{n-1}\left(R^i_\emptyset\right) + \beta\left[\left(F\left(R^i_\emptyset\right)^n V^i_\emptyset + \int_{R^i_\emptyset}^1 V^i(\theta')\,\mathrm{d}F(\theta')^n\right)\right. \\
&\quad \left. - \left(F\left(R^i_\emptyset\right)^{n-1} V^i\left(R^i_\emptyset\right) + \int_{R^i_\emptyset}^1 V^i(\theta')\,\mathrm{d}F(\theta')^{n-1}\right)\right] \\
&= \pi_n\left(R^i_\emptyset\right) - \widetilde{\pi}^i_{n-1}\left(R^i_\emptyset\right) + \beta\left[\left(F\left(R^i_\emptyset\right)^n V^i + V^i \int_{R^i_\emptyset}^1 \mathrm{d}F(\theta')^n\right)\right. \\
&\quad \left. - \left(F\left(R^i_\emptyset\right)^{n-1} V^i + V^i \int_{R^i_\emptyset}^1 \mathrm{d}F(\theta')^{n-1}\right)\right] \\
&= \pi_n\left(R^i_\emptyset\right) - \widetilde{\pi}^i_{n-1}\left(R^i_\emptyset\right),
\end{aligned} \tag{7.21}$$

where the second equality uses that at $\eta = 0$ $V^i_\emptyset = V^i(\theta) = V^i$.

It can also be shown that, at $\eta = 0$, $R^i_\emptyset > r^M$ if and only if

$$\widetilde{\pi}^i_{n-1}\left(r^M\right) - \pi_n\left(r^M\right) < \beta \int_{r^M}^1 \frac{\mathrm{d}}{\mathrm{d}\theta'}\widetilde{\pi}^i(\theta')\frac{F(\theta')^{n-1}(1-F(\theta'))}{\left(1-\beta F(\theta')^{n-1}\right)}\mathrm{d}\theta'. \tag{7.22}$$

Proof of Proposition 9

Expressions (7.16) and (7.20) also imply that the difference between $R^i_\emptyset$ and r^M is proportional to $\widetilde{\pi}^i_{n-1}\left(r^M\right) - \pi_n\left(r^M\right)$ at $\eta = 1$. By Fact 14, this difference is larger in the FPA than in the SPA. Therefore, $R^F_\emptyset < R^S_\emptyset$ local to $\eta = 1$ because, at $\eta = 1$, the two reserve prices are equal but the one in the FPA is more sensitive to change in η than the one in the SPA. A similar argument proves that $R^F_\emptyset < R^S_\emptyset$ local to $\eta = 0$, using expression (7.22) and the fact the marginal tracking revenue $\frac{\mathrm{d}}{\mathrm{d}\theta'}\widetilde{\pi}^i(\theta)$ is lower in the FPA (Fact 14).

Proof of Proposition 10

First, consider a neighborhood $\mathcal{N}_0$ of $\eta = 0$ such that there is full tracking in both auction formats, i.e. $\overline{\theta}^S = \overline{\theta}^F = 1$. We prove that $V^F_\emptyset > V^S_\emptyset$ for all $\eta \in \mathcal{N}_0$. Let $\hat{R}^S_\eta$ be the reserve price *policy* in the SPA given η. Define the value achieved by this policy in the SPA by $v^S\left(\hat{R}^S_\eta\right)$ and notice that by optimality this value is equal to $V^S_\emptyset$. Suppose the seller replicates the policy $\hat{R}^S_\eta$ in the FPA — i.e., she chooses the same initial reserve price. This policy is feasible but might be suboptimal, hence it achieves value $v^F\left(\hat{R}^S_\eta\right) \leq V^F_\emptyset$. Notice that the policy induces point wise — i.e., for any realization of the sequence of bidders' values — the same sequence of states as in the SPA.

Compare now the static revenue in the two auction formats, in each possible state. In state $\emptyset$, the Revenue Equivalence Theorem holds because bidders are symmetric and the reserve price is the same in the FPA and SPA. Hence, the expected static revenue is the same in the two auction formats. By contrasts, if there is an incumbent θ, then by Fact 14 the FPA yields a (strictly) higher expected static revenue. Therefore, the policy $\hat{R}^S_\eta$ yields a strictly higher value in the FPA than in the SPA. Summing up, $V^F_\emptyset \geq v^F\left(\hat{R}^S_\eta\right) > v^S\left(\hat{R}^S_\eta\right) = V^S_\emptyset$.

Second, consider $\eta = 1$. In this case, $R^i_\emptyset = R^i(\theta) = r^M$ and $V^i_\emptyset = V^i(\theta) = \frac{\pi_n(r^M)}{1-\beta}$ in both auction formats. Therefore, the SPA yields higher revenue than the FPA in a neighborhood of $\eta = 1$ if and only if $\frac{\partial}{\partial\eta}\left(V^S_\emptyset - V^F_\emptyset\right)\big|_{\eta=1} <$

0. The next Lemma shows that this condition depends on a comparison of the static revenues in the two auction formats with the optimal reserve price.

Lemma 20. *Let*

$$\Delta(y) = \int_y^1 \left(b\left(x, r^M, n\right) - \psi(x)\right) dF(x)^{n-1}.$$

Condition $\frac{\partial}{\partial \eta}\left(V_\emptyset^S - V_\emptyset^F\right)\big|_{\eta=1} < 0$ *is equivalent to*

$$\int_{r^M}^1 \Delta(y)\, dF(y)^n < 0. \tag{7.23}$$

To complete the proof we show that condition (7.23) is satisfied the assumption in the statement of the proposition. First, Lemma 21 uses the Revenue Equivalence Theorem to establish that an inequality similar to (7.23) holds. Finally, Lemma 22 proves that, under our assumption of single crossing between the bidding function in the FPA and the virtual value, this inequality implies condition (7.23).

Lemma 21. *It holds that* $\int_{r^M}^1 \Delta(y)\, dF(y) < 0$.

Proof. By (7.37), $\Delta(y) = \pi^F_{n-1,y}\left(r^M\right) - \pi^S_{n-1,y}\left(r^M\right)$. Recall that $\pi^i_{n-1,y}(r)$ is the expected revenue in auction i conditional on the realization y of the value of one of the bidders, and by LIE the expectation of this function with respect to y is the ex-ante expected revenue, which is the same in the FPA and SPA. Therefore, $\int_0^1 \pi^i_{y,n-1}\left(r^M\right) \mathrm{d}F(y) = \pi_n\left(r^M\right)$ for $i \in \{F, S\}$ and

$$\int_0^1 \pi^F_{n-1,y}\left(r^M\right) - \pi^S_{n-1,y}\left(r^M\right) \mathrm{d}F(y) = 0. \tag{7.24}$$

Moreover, for all $y < r^M$, $\pi^i_{y,n-1}\left(r^M\right)$ is independent of y (see (7.6) and (7.7)) and higher in the FPA than the SPA by Fact 1. Therefore, we can rewrite (7.24) as

$$F\left(r^M\right) \cdot \underbrace{\int_{r^M}^1 b(x, r, n) - b(x, r, n-1)\, \mathrm{d}F(x)^{n-1}}_{>0} + \int_{r^M}^1 \Delta(y)\, \mathrm{d}F(y) = 0,$$

from which the statement follows. □

Lemma 22. $\int_{r^M}^1 \Delta(y)\, dF(y) < 0$ *implies* $\int_{r^M}^1 \Delta(y)\, dF(y)^n < 0$.

Proof. Define

$$\delta(x) = b\left(x, r^M, n\right) - \psi(x) = \frac{1 - F(x)}{f(x)} - \frac{\int_{r^M}^x F(y)^{n-1} \mathrm{d}y}{F(x)^{n-1}}.$$

Since $\Delta(y) = \int_y^1 \delta(x)\, \mathrm{d}F(x)^{n-1}$, exchanging order of integration yields

$$\int_{r^M}^1 \Delta(y)\, \mathrm{d}F(y) = \int_{r^M}^1 \left(F(x) - F\left(r^M\right)\right) \delta(x)\, \mathrm{d}F(x)^{n-1}$$

and

$$\int_{r^M}^1 \Delta(y)\, \mathrm{d}F(y)^n = \int_{r^M}^1 \left(F(x)^n - F\left(r^M\right)^n\right) \delta(x)\, \mathrm{d}F(x)^{n-1}. \tag{7.25}$$

Using the factorization of the n^{th} power

$$F(x)^n - F\left(r^M\right)^n = \left(F(x) - F\left(r^M\right)\right)\sum_{j=0}^{n-1} F(x)^j F\left(r^M\right)^{n-1-j},$$

substituting in (7.25) and exchanging the order of integration and summation, we have

$$\begin{aligned}\int_{r^M}^1 \Delta(y)\,\mathrm{d}F(y)^n &= \int_{r^M}^1 \left(F(x) - F\left(r^M\right)\right)\sum_{j=0}^{n-1} F(x)^j F\left(r^M\right)^{n-1-j} \delta(x)\,\mathrm{d}F(x)^{n-1} \\ &= \sum_{j=0}^{n-1} F\left(r^M\right)^{n-1-j} \int_{r^M}^1 \left(F(x) - F\left(r^M\right)\right) F(x)^j \delta(x)\,\mathrm{d}F(x)^{n-1}.\end{aligned}$$

Since $\left(F(x) - F\left(r^M\right)\right) > 0$ in the relevant domain, the sign of the integrand is determined by the sign of δ. Notice that $\delta(1) < 0 < \delta\left(r^M\right)$ and, since by assumption there is a unique $x^\star$ such that $\delta(x^\star) = 0$, then $\delta(x) > 0$ for $x \in \left[r^M, x^\star\right]$ and $\delta(x) < 0$ for $x \in [x^\star, 1]$. It follows that $F(x)^j \delta(x) \le F(x^\star)^j \delta(x)$ for all $x \in \left[r^M, 1\right]$ (where both sides of the inequality are positive if $x < x^\star$ and negative if $x > x^\star$). Therefore, for all j,

$$\int_{r^M}^1 \left(F(x) - F\left(r^M\right)\right) F(x)^j \delta(x)\,\mathrm{d}F^{n-1}(x) < F\left(x^\star\right)^j \int_{r^M}^1 \left(F(x) - F\left(r^M\right)\right) \delta(x)\,\mathrm{d}F^{n-1}(x).$$

Finally, summing over j,

$$\begin{aligned}\int_{r^M}^1 \Delta(y)\,\mathrm{d}F^n(y) &< \sum_{j=0}^{n-1} F\left(r^M\right)^{n-1-j} F\left(x^\star\right)^j \int_{r^M}^1 \left(F(x) - F\left(r^M\right)\right) \delta(x)\,\mathrm{d}F(x)^{n-1} \\ &= \sum_{j=0}^{n-1} F\left(r^M\right)^{n-1-j} F\left(x^\star\right)^j \int_{r^M}^1 \Delta(y)\,\mathrm{d}F(y)\end{aligned}$$

implying the result. □

The single crossing assumption in the statement of the proposition is sufficient, but not necessary for condition (7.23) to hold. In fact, the condition only depends on primitives and it can be readily checked that it is satisfied in examples where single crossing fails. Since we have been unable to find an example where the condition is not satisfied, we conjecture that single crossing may not be actually needed for the result (although a formal proof has proven elusive).

Stationary Distribution of Winners' Values and Trade

In the section we establish results that are stated and used in Sections 5 and 5.1.

Standard arguments guarantee that, for each value of β, η and n, a stationary distribution of winners' values exists and is unique (SLP 12.13). Given the distribution $G_t \in \Delta(\emptyset \cup [R_\emptyset, 1])$ at time t and the policies from Theorem 7, the distribution $G_{t+1} \in \Delta(\emptyset \cup [R_\emptyset, 1])$ at time $t+1$ is given by

$$G_{t+1,\emptyset} = G_{t,\emptyset} F(R_\emptyset)^n + \int_{R_\emptyset}^1 \eta F(R(\theta))^n\,\mathrm{d}G_t(\theta),$$

$$\begin{aligned}\mathrm{d}G_{t+1}(\theta) =& \mathrm{d}G_t(\theta)(1-\eta)F(\theta)^{n-1} + G_{t,\emptyset}\mathrm{d}F(\theta)^n + \left[\eta \mathrm{d}F(\theta)^n + (1-\eta)\mathrm{d}F(\theta)^{n-1}\right]G_t(\theta)\\ &+\eta \mathrm{d}F(\theta)^n \int_{\{\theta':R(\theta')\le\theta<\theta'\}} \mathrm{d}G_t(\theta').\end{aligned}$$

Imposing stationarity — i.e., that the distributions at t and $t+1$ are the same — yields (5.1) and (5.2).

Expression (5.2) depends on the set of incumbents that allow a downward transition into θ:

$$\{\theta' : R(\theta') \le \theta < \theta'\} = \begin{cases} \varnothing & \text{if } \theta \in \mathcal{A} := [R_\emptyset, R(1)] \\ [R(\theta)^{-1}, 1] & \text{if } \theta \in C := [R(1), \overline{\theta}] \\ (\theta, 1] & \text{if } \theta \in \mathcal{B} := \left[\overline{\theta}, 1\right] \end{cases} \tag{7.26}$$

Substituting the three branches of (7.26) into (5.2) yields, after straightforward manipulations,

$$\mathrm{d}G^A(\theta) = \frac{[G_\emptyset + \eta G(\theta)]\,\mathrm{d}F(\theta)^n + (1-\eta)\,\mathrm{d}F(\theta)^{n-1}G(\theta)}{1-(1-\eta)F(\theta)^{n-1}} \tag{7.27}$$

$$\mathrm{d}G^B(\theta) = \frac{[G_\emptyset + \eta(1-G_\emptyset)]\,\mathrm{d}F(\theta)^n + (1-\eta)\,\mathrm{d}F(\theta)^{n-1}G(\theta)}{1-(1-\eta)F(\theta)^{n-1}} \tag{7.28}$$

$$\mathrm{d}G^C(\theta) = \frac{[G_\emptyset + \eta(1-G_\emptyset)]\,\mathrm{d}F(\theta)^n + (1-\eta)\,\mathrm{d}F(\theta)^{n-1}G(\theta) - \eta \mathrm{d}F(\theta)^n\left(G\left(R(\theta)^{-1}\right) - G(\theta)\right)}{1-(1-\eta)F(\theta)^{n-1}} \tag{7.29}$$

where $\mathrm{d}G^I(\theta)$ denotes the stationary density for $\theta \in \mathcal{I}$. Comparing (7.27), (7.28) and (7.29) (using that G is increasing and bounded by $1-G_\emptyset$) yields that $\mathrm{d}G^A(\theta) < \mathrm{d}G^C(\theta) < \mathrm{d}G^B(\theta)$.

In the Online Appendix, we derive the stationary density of θ by first treating $G_\emptyset$ as a constant and then solving for it using continuous pasting.[59]

To analyze how $G_\emptyset$ varies with η (Section 5.1), first notice that $G_\emptyset$ is a continuous function of η because the functions that characterize it (see $\hat{G}^A_x(\theta)$, $\hat{G}^B_x(\theta)$ and $\hat{G}^C_x(\theta)$ in the Online Appendix) are continuous in η for all θ. If $\eta = 0$, (5.1) yields

$$G_\emptyset = G_\emptyset F(R_\emptyset)^n$$

and hence $G_\emptyset = 0$. If $\eta = 1$, (5.1) yields

$$\begin{aligned} G_\emptyset &= G_\emptyset F\left(r^M\right)^n + \int_{r^M}^1 F\left(r^M\right)^n \mathrm{d}G(\theta) \\ &= G_\emptyset F\left(r^M\right)^n + F\left(r^M\right)^n (1-G_\emptyset) = F\left(r^M\right)^n > 0. \end{aligned}$$

[59]Continuity in θ of the stationary distribution is an immediate consequence of the properties of the stochastic process and the fact that the distribution F is continuous.

Total differentiation of (5.1) (taking into account that $R_\emptyset, R(\theta)$ and $G(\theta)$ all depend on η) yields

$$\frac{\mathrm{d}}{\mathrm{d}\eta}G_\emptyset = \frac{\mathrm{d}}{\mathrm{d}\eta}G_\emptyset F(R_\emptyset)^n + G_\emptyset \frac{\mathrm{d}}{\mathrm{d}\eta}R_\emptyset \mathrm{d}F(R_\emptyset)^n - \eta\frac{\mathrm{d}}{\mathrm{d}\eta}R_\emptyset F(R_\emptyset)^n \mathrm{d}G(R_\emptyset) + \int_{R_\emptyset}^1 F(R(\theta))^n \mathrm{d}G(\theta) + \\ + \eta\left[\int_{R_\emptyset}^1 \frac{\mathrm{d}}{\mathrm{d}\eta}R(\theta)\mathrm{d}F(R(\theta))^n \mathrm{d}G(\theta) + \int_{R_\emptyset}^1 F(R(\theta))^n \mathrm{d}\left(\frac{\mathrm{d}}{\mathrm{d}\eta}G(\theta)\right)\right].$$

This expression contains the direct effect, the reserve price effect and the distribution effect of η described in the main text.

Finally, in the Online Appendix we show a closed-form example where $G_\emptyset$ in the SPA is non-monotonic in η.

Proof of Proposition 11

We use the notation of the proof of the results in Sections 5. To make explicit the dependence of the the stationary distribution in region $\mathcal{A}$ on the initial reserve price $R_\emptyset$, we denote by $\hat{G}^A_{G_\emptyset,R_\emptyset}$ the solution to (7.27) with initial condition $G(R_\emptyset) = 0$. Notice that for all $\theta \in (R_\emptyset, R(1)]$, $\hat{G}^A_{G_\emptyset,R_\emptyset}(\theta)$ is a decreasing function of $R_\emptyset$ (since increasing $R_\emptyset$ shrinks the domain of integration and depresses G which feeds back into $\mathrm{d}G$ through (7.27)).

If $\overline{\theta}^F = \overline{\theta}^S = 1$, both auction formats track all incumbents and differ only for the initial reserve price $R^i_\emptyset$. In this case, $G^i_\emptyset$ is the (unique) solution to $\hat{G}^A_{x,R^i_\emptyset}(1) = 1 - x$ and (by Proposition 9) $R^F_\emptyset < R^S_\emptyset$. Therefore, $\hat{G}^A_{x,R^F_\emptyset}(1) > \hat{G}^A_{x,R^S_\emptyset}(1)$ and hence $G^F_\emptyset < G^S_\emptyset$.

If $\beta = 0$, $R^S_\emptyset = R^F_\emptyset = r^M$ and the reserve price coincides with the static optima on $\left[r^M, 1\right]$ (Proposition 5). Recall that in the static optima there is tailing in the FPA if and only if there is tailing in the SPA. Without tailing, $G^F_\emptyset = G^S_\emptyset$ because $R^S_\emptyset = R^F_\emptyset$ (following the same argument as above). With tailing, $\overline{\theta}^F > R^F(1) = \overline{\theta}^S = R^S(1)$, so that region $\mathcal{A} = \left[r^M, \overline{\theta}^S\right]$ is the same in the FPA and in the SPA. In this case (since $\overline{\theta}^S = R^F(1)$), $G^S_\emptyset$ is the unique solution to $\hat{G}^A_{x,r^M}\left(\overline{\theta}^S\right) = \hat{G}^B_x\left(\overline{\theta}^S\right)$, while $G^F_\emptyset$ is the unique solution to $\hat{G}^A_{x,r^M}\left(\overline{\theta}^S\right) = \hat{G}^C_x\left(\overline{\theta}^S\right)$. Notice that

$$\begin{aligned} \hat{G}^C_x\left(\overline{\theta}^S\right) &= \hat{G}^B_x\left(\overline{\theta}^F\right) - \int_{\overline{\theta}^S}^{\overline{\theta}^F} \mathrm{d}G^C(\theta')\,\mathrm{d}\theta' \\ &> \hat{G}^B_x\left(\overline{\theta}^F\right) - \int_{\overline{\theta}^S}^{\overline{\theta}^F} \mathrm{d}G^B(\theta')\,\mathrm{d}\theta' = \hat{G}^B_x\left(\overline{\theta}^S\right), \end{aligned}$$

where we used that $\mathrm{d}G^B(\theta') > \mathrm{d}G^C(\theta')$. Therefore, $G^F_\emptyset > G^S_\emptyset$.

Proof of Lemma 12

First, we solve for the joint distributions $g(\theta,\tau)$. Iterating on 6.1, we obtain

$$g(\theta,\tau) = \left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1} g(\theta,1)$$

and using

$$g(\theta) = \sum_{\tau\geq 1} g(\theta,\tau) = g(\theta,1)\sum_{\tau\geq 1}\left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1} = \frac{g(\theta,1)}{1-(1-\eta)F(\theta)^{n-1}}$$

we get $g(\theta,\tau) = \left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1}\left[1-(1-\eta)F(\theta)^{n-1}\right]g(\theta)$.

Next, we prove that $\frac{g_{\tau}(\theta)}{g_{\tau-1}(\theta)}$ is increasing in θ for all τ, that is the conditional distributions satisfy the MLR property (which implies FOSD). Using 6.1 we obtain

$$\frac{g_{\tau}(\theta)}{g_{\tau-1}(\theta)} = (1-\eta)F(\theta)^{n-1} \cdot \frac{\int_{R_{\emptyset}}^{1} g(\theta', \tau-1)\mathrm{d}\theta'}{\int_{R_{\emptyset}}^{1} g(\theta', \tau)\mathrm{d}\theta'},$$

which is increasing in θ.

To prove that the conditional distributions converge in probability to 1, let $\overline{g}(1) = \sup_{\theta} g(\theta, 1)$ which is finite. For all $\epsilon > 0$,

$$\begin{aligned}
\lim_{\tau\to\infty} \mathbb{P}_{\tau}(\theta < 1-\epsilon) &= \lim_{\tau\to\infty} \int_{R_{\emptyset}}^{1-\epsilon} g_{\tau}(\theta)\,\mathrm{d}\theta = \lim_{\tau\to\infty} \int_{R_{\emptyset}}^{1-\epsilon} \left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1} g(\theta,1)\mathrm{d}\theta \\
&< \lim_{\tau\to\infty} \int_{R_{\emptyset}}^{1-\epsilon} \left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1} \overline{g}(1)\,\mathrm{d}\theta \\
&= \overline{g}(1) \int_{R_{\emptyset}}^{1-\epsilon} \lim_{\tau\to\infty} \left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1} \mathrm{d}\theta = 0,
\end{aligned}$$

where we used that $\left[(1-\eta)F(\theta)^{n-1}\right]^{\tau-1} \to 0$ uniformly in $[R_{\emptyset}, 1-\epsilon]$.

Proof of Proposition 13

Recall that expected seller's revenue given tenure τ is $E^{i}(\tau) = \mathbb{E}_{\theta|\tau}\left(e^{i}(\theta)\right)$ where $e^{i}(\theta) = b\left(\theta, R^{i}(\theta), n\right)$ is the expected transfer of a winner θ who set his own reserve $R^{i}(\theta)$ and $\theta|\tau$ denotes the conditional distribution of persistent winners with tenure τ (that also depends on the auction format i). Because $e^{S}(\theta)$ is increasing in θ, then $\theta|\tau+1 \trianglerighteq^{\text{FOSD}} \theta|\tau$ (Lemma 12) implies that $E^{S}(\tau+1) > E^{S}(\tau)$ for all τ. To complete the proof we need to show that there is T such that for all $\tau > T$ $E^{F}(\tau+1) < E^{F}(\tau)$ when condition 6.2 holds. We first establish a useful auxiliary result.

Lemma 23. *Consider three bounded real sequences $\{p_t\}_{t\geq 1}$, $\{y_t\}_{t\geq 1}$ and $\{z_t\}_{t\geq 1}$ with the following properties*

i) Monotonicity. For all t, $p_{t+1} > p_t$ and $y_{t+1} < y_t$.

ii) Convergence. $\lim_{t\to\infty} p_t = 1$, $\lim_{t\to\infty} y_t = \bar{y} < \bar{z} = \lim_{t\to\infty} z_t$.

Then, the sequence x_t defined pointwise as $x_t = p_t y_t + (1-p_t) z_t$ is eventually decreasing.

Suppose that condition 6.2 holds, and notice that this ensures that there exists $x^{\star} < 1$ such that e^{F} is strictly decreasing in $[x^{\star}, 1]$.[60] Let $p_t = \mathbb{P}_{\theta|t}(\theta > x^{\star})$, $y_t = \mathbb{E}_{\theta|t}\left[e^{F}(\theta)|\theta > x^{\star}\right]$, $z_t = \mathbb{E}_{\theta|t}\left[e^{F}(\theta)|\theta < x^{\star}\right]$ and notice that $E^{F}(t) = \mathbb{E}_{\theta|t}\left[e^{F}(\theta)\right]$ is exactly the derived sequence $x_t = p_t y_t + (1-p_t) z_t$. Therefore, to prove that $E^{F}(t)$ is eventually decreasing it is sufficient to show that the hypothesis of Fact 23 hold in our case.

Monotonicity. Because p_t is the expectation of an increasing function (the indicator that θ is above $x^{\star}$) and y_t is the expectation of a function that is decreasing in the relevant domain, both monotonicity assumptions follow from the FOSD ranking (Lemma 12).

Convergence. As tenure grows, the distribution of persistent winners converges in probability to 1 (Lemma 12), so by continuous mapping[61] we get both $p_t \to 1$ and $y_t \to e^{F}(1)$. To see that $z_t \to e^{F}(x^{\star})$, use (6.1) to obtain that

[60] It is immediate to see that, if the inequality in (6.2) is reversed (e.g. if there is full tracking), then the limit result is also reversed.

[61] Notice that the function e is continuous in θ, while $\mathbb{I}\left[\theta > x^{\star}\right]$ is continuous for $\theta > x^{\star}$, an interval that contains the limit point 1.

$\frac{g_\tau(\theta)}{g_\tau(\theta')} \propto \left(\frac{F(\theta)}{F(\theta')}\right)^{(n-1)(\tau-1)}$ and therefore $\theta > \theta' \Rightarrow \lim_{\tau\to\infty} \frac{g_\tau(\theta)}{g_\tau(\theta')} = \infty$. Hence, the distribution conditional on being below $x^\star$ converges in probability to the upper bound and, by continuous mapping, the conditional expectation converges to the value of the function at $x^\star$. Because $e^F(x^\star) > e^F(1)$ by assumption, the limit of z_t exceeds the limit of y_t and all properties are verified.

This concludes the proof of Proposition 13. In the text we argue that the probability of re-trade is decreasing (in tenure) in the SPA, while it is eventually increasing in the FPA if there is tailing. The arguments made for revenues are readily adapted once we notice that, given a winner of valuation θ, the probability of re-trade is $t^i(\theta) = 1 - \eta F\left(R^i(\theta)\right)^n$, which is (weakly) decreasing in the SPA, and strictly increasing in the tailing region $\left[\bar{\theta}^F, 1\right]$ for the FPA.

Online Appendix

The Online Appendix is organized as follows. Sections 1 and 2 extend the main model: in Section 1 we analyze a setting where the number of bidders varies across periods, while in Section 2 we analyze a setting where bidders do not leave after a period without transaction. Section 3 provides the proofs omitted from the Appendix.

Stochastic Number of Bidders

In this section, we relax our initial assumption of having a constant number of bidders, n, who participate in the auction in every period. We show that the qualitative features of the optimal reserve prices remain consistent with those in our basic model.

Suppose that in each period the number of bidders, n, is stochastic and drawn from a distribution $N \in \Delta(\{2,3,\ldots\})$ with finite support $\mathcal{N}$.[62] Draws of n are serially independent and independent of the presence of the incumbent.[63] We assume that both the seller and the bidders observe the realization of n before choosing their actions — i.e., respectively, selecting the reserve price and bidding. The structure of the repeated auction is otherwise unchanged.

The seller solves a dynamic problem with a two-dimensional state (θ, n), where the number of bidders only affects the current period revenues. The transition dynamics is given by the product measure

$$\mu(\theta', n'|\theta, R, n) = N(n') \times \mu(\theta'|\theta, R), \tag{7.30}$$

where $\mu(\theta'|\theta, R)$ represents the incumbent transition conditional on the reserve described in Section 4. Let $\mathbb{E}[\cdot]$ be the expectation operator associated to (7.30) and $\pi^i(\theta, R)$ the static revenues (see (7.6) and (7.7)). Then, the value functions for the seller's problem are

$$V^i(\theta, n) = \max_R \pi^i(\theta, R) + \beta \mathbb{E}\left[V^i(\theta', n')\right]$$

and

$$V^i_\emptyset(n) = \max_R \pi_n(R) + \beta \mathbb{E}\left[V^i(\theta', n')\right].$$

Notice that these expressions differ from those in Theorem 7 only for the fact that the continuation value has an expectation over next-period's number of bidders.

Lemma 24. *In auction format $i = F, S$, the initial reserve price $R^i_\emptyset$ solves*

$$\beta \mathbb{E}\left[V_\emptyset(n') - V\left(R^i_\emptyset, n'\right)\right] = \psi\left(R^i_\emptyset\right), \tag{7.31}$$

while the tailing reserve price $\overline{R}^i(\theta, n)$ solves conditions (4.4) and (4.5), with the dynamic wedge replaced by

$$\mathbb{E}\left[V^i_\emptyset(n') - V^i\left(\overline{R}^i(\theta, n'), n'\right)\right]. \tag{7.32}$$

[62] The analysis of the base model in the paper corresponds to a degenerate distribution δ_n for some n.

[63] While the serial independence of n is only assumed for simplicity, the requirement that the seller cannot use the realization of n to infer whether the incumbent stayed or not is crucial. This can be microfounded assuming that n represents an exogenous physical constraint on the "bidding slots" that are open, and that there is always excess demand for slots: the number of bidders in the auction is equal to the number of slots available but the seller does not know whether the incumbent sits in one of the slots.

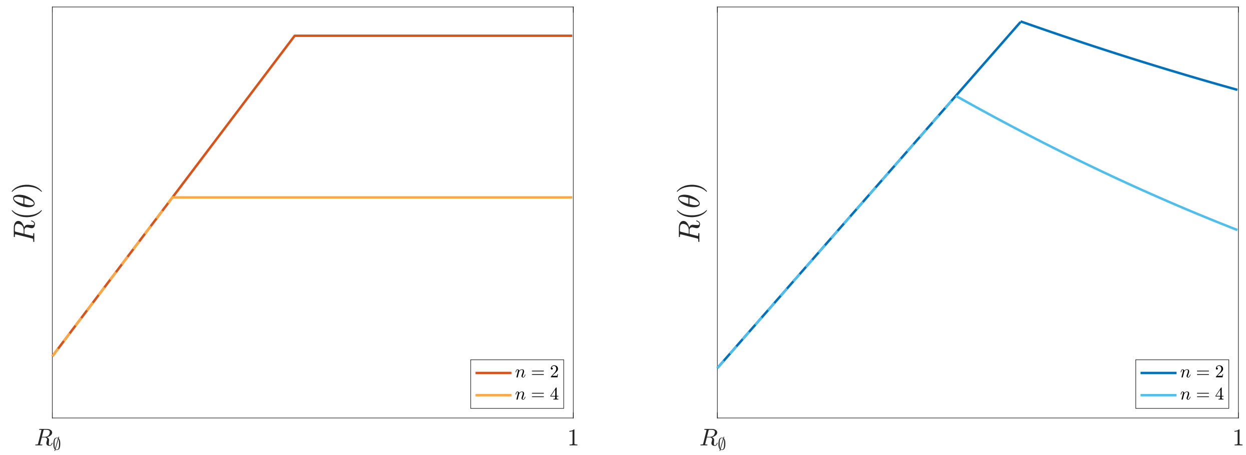

Figure 7.1: Optimal reserve prices in SPA (left) and FPA (right) as a function of the state θ, with $\beta = 0.5$, $\eta = 0.5$, $N = \mathcal{U}(\{2,4\})$.

The proof of the Lemma is analogous to the proof of Theorem 7.Notice that equations (7.31) and (7.32) indicate that both the initial and the tailing reserve depend on the number of bidders in the next period n'. However, since the number of bidders n in the current auction does not affect n', the $R_\emptyset$ and the dynamic wedge are independent of n.[64] This implies that Lemmas 18 and 19 extend to an environment with a stochastic number of bidders and that, as in our main model, the seller never excludes an incumbent. Similarly, the tailing reserve (as well as the tracking threshold) depends on the current number of bidders n only through the static wedge, and is decreasing in n. Notice that when the seller observes a high number of bidders in the current auction she always sets a reserve price (weakly) lower than when she observes a low number of bidders (as shown in Figure 7). The reason is that the higher the number of bidders, the less sensitive current revenues are to the reserve price, which shrinks the static wedge.

Bidders Stay after No-Trade

Our base model assumes that bidders who do not win leave the repeated auction, even if the auction resulted in no-transaction, hence no bidder submitted a bid higher than the reserve price.

In this section, we analyze an extension of our model where losers of auction that resulted in no transaction participate in the next period, and maintain their valuation. This can be interpreted as bidders maintaining their match values until an advertisement is actually shown to the user. We show that most qualitative properties of the main model are preserved in this extension, which however features exclusion on path (and hence is closer to the static benchmark of Section 3). Formally, we consider the following participation dynamics:

1. When the object is sold, the winner participates in the subsequent auction with probability $1-\eta$, whereas losers leave.

2. When the object is *not* sold, all bidders stay and bid in the subsequent auction.

[64] If instead the distribution of n' depends on the current number of bidders n, the dynamic wedge is $\mathbb{E}_{n'|n}\left[\left(V\left(n'\right)-V\left(R,n'\right)\right)\right]$.

The dynamics after a period with transaction are the same as in the base model. If the object is not sold, however, bidders of the subsequent auction are symmetric and have valuations drawn from a distribution truncated at the previous reserve. Contrary to the base model, therefore, after a period with no transaction the repeated auction does not restart anew. The seller does not face a situation equivalent to the starting one, and the value of the reserve price that led to a no-trade state is relevant in the choice of the optimal reserve price.

The seller now solves a dynamic problem with state θ, which is either the value of the incumbent or, in case the object was not sold in the previous auction, the last reserve price. Formally, the state space is $[0,1] \times \{T, NT\}$.[65] Let V^{NT} be the value of the seller after a period with no trade, and V^T be her value after a period with trade. Then,

$$V^T(\theta) = \max_{R \le \theta} \pi(\theta, R) + \beta \left[\eta \left(F(R)^n V^{NT}(R) + \int_R^1 V^T(\theta') \, \mathrm{d}F(\theta')^n \right) + \right.$$
$$\left. + (1-\eta) \left(F(\theta)^{n-1} V^T(\theta) + \int_\theta^1 V^T(\theta') \, \mathrm{d}F(\theta')^{n-1} \right) \right],$$

and

$$V^{NT}(\theta) = \max_R \frac{1}{F(\theta)^n} \left[\int_R^\theta \psi(x|\theta) \, \mathrm{d}F(x)^n + \beta \left(F(R)^n V^{NT}(R) + \int_R^\theta V^T(x) \, \mathrm{d}F(x)^n \right) \right],$$

where $\psi(x|\theta) \equiv x - \frac{F(\theta) - F(x)}{f(x)}$ is the virtual value of the truncated distribution.

This model induces a dual reserve price policy. $R^T(\theta)$ — the counterpart of the policy $R(\theta)$ in the base model — denotes the reserve when there is an incumbent with value θ, while $R^{NT}(\theta)$ denotes the reserve when there are n symmetric bidders with value lower than θ (where θ represents the reserve price of the failed auction.)

The two main differences with respect to the base model are the following. First, there is no minimum reserve price $R_\emptyset$ because $R^{NT}(\theta)$ can be arbitrarily low.[66] Second, since there can be incumbents with arbitrarily low values, the seller can chose on path a reserve price higher than the incumbent, hence she can choose to exclude the incumbent.

Figure 7 displays the two reserve price policies in the SPA as a function of the state and shows that, after a period with no trade, the reserve price can be equal to 0. The reason is that, when the seller faces a pool of bidders with low values (that all bid below the reserve price in the previous auction), she prefers to reduce the reserve price enough to guarantee a transaction, to induce the low-value bidders to leave and secure the participation of a new set of bidders in subsequent auctions. In other words, the seller sacrifices current revenues in favor of future ones, even with relatively low values of β.

Additional Proofs

Proof of Lemma 14

Direct computation gives

$$\widetilde{\pi}^F(\theta) - \widetilde{\pi}^S(\theta) = \int_\theta^1 b(x, \theta, n) \, \mathrm{d}F(x)^{n-1} - \int_\theta^1 b(x, \theta, n-1) \, \mathrm{d}F(x)^{n-1} > 0$$

[65] State $\omega = \{\theta, T\}$ corresponds to information "the previous auction sold the object to type θ", while state $\omega = \{\theta, NT\}$ corresponds to information "the previous auction failed at reserve price θ".

[66] It is immediate to see that $R^{NT}(\theta)$ is below θ, as otherwise the dynamic auction would be stuck in a no-trade (zero revenue) state.

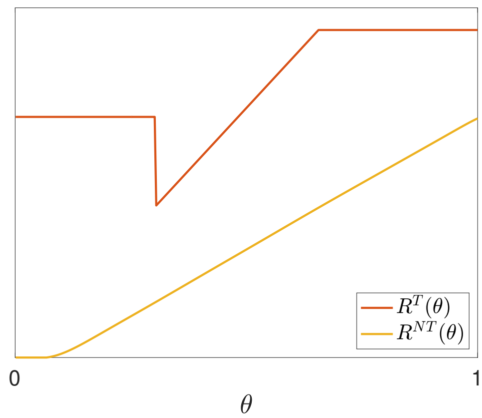

Figure 7.2: Reserve price policies in the SPA with $\beta = 0.2$, $n = 3$, $\eta = 0.5$.

by Fact 1. Moreover,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\theta}\widetilde{\pi}^F(\theta) - \frac{\mathrm{d}}{\mathrm{d}\theta}\widetilde{\pi}^S(\theta) &= -b(\theta,\theta,n) + b(\theta,\theta,n-1) + \int_\theta^1 \frac{\partial}{\partial R} b(x,\theta,n)\,\mathrm{d}F(x)^{n-1} \\ &\quad - \int_\theta^1 \frac{\partial}{\partial R} b(x,\theta,n-1)\,\mathrm{d}F(x)^{n-1} \\ &= \int_\theta^1 \underbrace{\frac{\partial}{\partial R} b(x,\theta,n) - \frac{\partial}{\partial R} b(x,\theta,n-1)}_{\leq 0} \mathrm{d}F(x)^{n-1} \leq 0,\end{aligned}$$

where the central inequality is Fact 2.

Proof of Lemma 15

If $\theta > \theta'$, then $\pi^i(\theta, R) > \pi^i(\theta', R)$ and $F(\theta_{t+1}|\theta, R) \trianglerighteq^{\text{FOSD}} F(\theta_{t+1}|\theta', R)$ for all R.[67] The result then follows from Theorem 9.7 in SLP.

Proof of Lemma 16

Recall that $R^i_\emptyset$ solves

$$\max_R \pi_n(R) + \beta\left[F(R)^n V^i_\emptyset + \int_R^1 V^i(\theta')\,\mathrm{d}F(\theta')^n\right].$$

The FOC of this problem is

$$-\mathrm{d}F(R)^n \psi(R) + \beta \mathrm{d}F(R)^n \left(V^i_\emptyset - V^i(R)\right) = 0,$$

which is rearranged to the desired expression.

[67] If the incumbent leaves, the dynamics depends only on the reserve; if he stays then, let y be the max of $n-1$ standard bidders. If $y < \theta'$ then $\theta_{t+1}|R,\theta = \max\{R,\theta\} > \max\{R,\theta'\} = \theta_{t+1}|R,\theta'$. If $\theta' < y < \theta$ then $\theta_{t+1}|R,\theta = \max\{R,\theta\} > \max\{R,y\} = \theta_{t+1}|R,\theta'$. If $y > \theta$ then $\theta_{t+1}|R,\theta = \max\{R,y\} = \max\{R,y\} = \theta_{t+1}|R,\theta'$.

Proof of Lemma 17

Differentiating $V^{+,i}(\theta)$, we get that $R^i(\theta)$ solves

$$\frac{\mathrm{d}}{\mathrm{d}R}\pi^{+,i}(\theta,R)+\beta\eta\mathrm{d}F(R)^n\left(V_\emptyset^i-V^i(R)\right)=0. \tag{7.33}$$

Substituting the static marginal revenues (7.8)-(7.9) and rearranging yields conditions (4.4)-(4.5). Notice the dynamic wedge $\beta\eta\mathrm{d}F(R)^n\left(V_\emptyset^i-V^i(R)\right)$ is independent of θ in both auction formats. Since $\frac{\mathrm{d}}{\mathrm{d}R}\pi^{+,S}(\theta,R)$ is also independent of θ (see (7.8)), the reserve price that solves (7.33) is independent of θ. In the FPA, sub-modularity of $\pi^{+,F}$ implies sub-modularity of the dynamic objective since

$$\frac{\partial^2}{\partial R\partial\theta}\eta\left(F(R)^nV_\emptyset^i+\int_R^1V^i(y)\,\mathrm{d}F(y)^n\right)+(1-\eta)\left(F(\theta)^{n-1}V^i(\theta)+\int_\theta^1V^i(y)\,\mathrm{d}F(y)^{n-1}\right)=0.$$

Therefore, as in the static case, $R^F(\theta)$ is decreasing in the tailing region.

Proof of Lemma 20

Throughout the proof, all derivatives are evaluated at $\eta=1$. Substituting the reserve prices and values at $\eta=1$ into (7.19),

$$\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)\left(1-\beta F\left(r^M\right)^n\right)=\beta\int_{r^M}^1\frac{\partial}{\partial\eta}\left(V^S(y)-V^F(y)\right)\mathrm{d}F(y)^n. \tag{7.34}$$

Applying the envelope theorem to (7.14) (which is the relevant specification of the value function when all incumbents are tailed),

$$\begin{aligned}\frac{\partial}{\partial\eta}V^i(\theta)&=\frac{\partial}{\partial\eta}\pi^{+,i}\left(\theta,r^M\right)+\beta\frac{\pi_n\left(r^M\right)}{1-\beta}\left(\underbrace{F\left(r^M\right)^n+\int_{r^M}^1\mathrm{d}F(\theta')^n-F(\theta)^{n-1}-\int_\theta^1\mathrm{d}F(\theta')^{n-1}}_{=0}\right)\\&\quad+\beta\left(F\left(r^M\right)^n\frac{\partial}{\partial\eta}V_\emptyset^i+\int_{r^M}^1\frac{\partial}{\partial\eta}V^i(x)\,\mathrm{d}F(x)^n\right).\end{aligned} \tag{7.35}$$

Therefore,

$$\begin{aligned}\frac{\partial}{\partial\eta}\left(V^S(y)-V^F(y)\right)&=\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\\&\quad+\beta F\left(r^M\right)^n\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)+\beta\int_{r^M}^1\frac{\partial}{\partial\eta}\left(V^S(x)-V^F(x)\right)\mathrm{d}F(x)^n.\end{aligned}$$

Integrating,

$$\int_{r^M}^1\frac{\partial}{\partial\eta}\left(V^S(y)-V^F(y)\right)\mathrm{d}F(y)^n=\int_{r^M}^1\left[\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\right]\mathrm{d}F(y)^n+$$

$$+\beta\left(1-F\left(r^M\right)^n\right)F\left(r^M\right)^n\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)+\beta\left(1-F\left(r^M\right)^n\right)\int_{r^M}^1\frac{\partial}{\partial\eta}\left(V^S(x)-V^F(x)\right)\mathrm{d}F(x)^n$$

$$\Leftrightarrow\quad\int_{r^M}^1\frac{\partial}{\partial\eta}\left(V^S(y)-V^F(y)\right)\mathrm{d}F(y)^n=$$

$$=\frac{\int_{r^M}^1\left[\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\right]\mathrm{d}F(y)^n+\beta\left(1-F\left(r^M\right)^n\right)F\left(r^M\right)^n\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)}{1-\beta\left(1-F\left(r^M\right)^n\right)}.$$

Substituting in (7.34) yields

$$\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)\left(1-\beta F\left(r^M\right)^n\right)=$$

$$=\beta\frac{\int_{r^M}^1\left[\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\right]\mathrm{d}F(y)^n+\beta\left(1-F\left(r^M\right)^n\right)F\left(r^M\right)^n\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)}{1-\beta\left(1-F\left(r^M\right)^n\right)}$$

$$\Leftrightarrow\quad\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)\left(1-\beta F\left(r^M\right)^n-\frac{\beta^2\left(1-F\left(r^M\right)^n\right)F\left(r^M\right)^n}{1-\beta\left(1-F\left(r^M\right)^n\right)}\right)=$$

$$=\beta\frac{\int_{r^M}^1\left[\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\right]\mathrm{d}F(y)^n}{1-\beta\left(1-F\left(r^M\right)^n\right)}$$

$$\Leftrightarrow\quad\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)(1-\beta)=\beta\int_{r^M}^1\left[\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\right]\mathrm{d}F(y)^n.$$

Hence,

$$\frac{\partial}{\partial\eta}\left(V_\emptyset^S-V_\emptyset^F\right)\propto\int_{r^M}^1\left[\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)\right]\mathrm{d}F(y)^n.\tag{7.36}$$

Notice that

$$\frac{\partial}{\partial\eta}\pi^{+,i}\left(y,r^M\right)=\pi_n\left(r^M\right)-\pi_{n-1,y}^i\left(r^M\right)$$

so that

$$\begin{aligned}\frac{\partial}{\partial\eta}\pi^{+,S}\left(y,r^M\right)-\frac{\partial}{\partial\eta}\pi^{+,F}\left(y,r^M\right)&=\pi_{n-1,y}^F\left(r^M\right)-\pi_{n-1,y}^S\left(r^M\right)\\&=\int_y^1\left(b\left(x,r^M,n\right)-b(x,y,n-1)\right)\mathrm{d}F(x)^{n-1}.\end{aligned}\tag{7.37}$$

Using (7.5),

$$\int_y^1 b(x,y,n-1)\,\mathrm{d}F(x)^{n-1}=\int_y^1\psi(x)\,\mathrm{d}F(x)^{n-1}.$$

Substituting the last two expressions in (7.36) yields (7.23).

Stationary Distribution of Winners' Values

Starting from expressions (7.27), (7.28) and (7.29), we solve for the stationary density of θ as follows. We first obtain the stationary distribution treating $G_\emptyset$ as a constant and then obtain $G_\emptyset$ using continuous pasting of the stationary CDF.

Step 1. Distributions fixing $G_\emptyset$. Given $G_\emptyset$, $\mathrm{d}G^A(\theta)$ and $\mathrm{d}G^B(\theta)$ represent two well-behaved differential equations that can be solved (forward and backwards, respectively) in the relevant domains starting from the initial conditions $G(R_\emptyset) = 0$ and $G(1) = 1 - G_\emptyset$, respectively. Let $\hat{G}^A_{G_\emptyset} : \mathcal{A} \to [0,1]$ and $\hat{G}^B_{G_\emptyset} : \mathcal{B} \to [0,1]$ denote these solutions. Notice that, for all $\theta \in \mathcal{A}$, $\hat{G}^A_{G_\emptyset}(\theta)$ is continuous and increasing in $G_\emptyset$ and, for all $\theta \in \mathcal{B}$, $\hat{G}^B_{G_\emptyset}(\theta)$ is continuous and decreasing in $G_\emptyset$.

Because the last addendum of the numerator in (7.29) depends on $G\left(R(\theta)^{-1}\right)$, solving for $\mathrm{d}G^C(\theta)$ requires the stationary CDF in $\mathcal{B}$ (because $\theta \in \mathcal{C} \Rightarrow R(\theta)^{-1} \in \mathcal{B}$). Hence, we use $\hat{G}^B_{G_\emptyset}$ and the inverse reserve price policy to solve (7.29) in $\mathcal{C}$, proceeding backwards from the initial condition $G\left(\overline{\theta}\right) = \hat{G}^B_{G_\emptyset}\left(\overline{\theta}\right)$. Let $\hat{G}^C_{G_\emptyset} : \mathcal{C} \to [0,1]$ denote this solution and notice that it is continuous and decreasing in $G_\emptyset$.

Step 2. Computing $G_\emptyset$. The procedure in Step 1 takes reserve prices as given. The auction format i then determines the reserve prices and the conditions used to pin down $G_\emptyset$. We distinguish three cases.

Case 1. Full Tracking. If $R^i(1) = \overline{\theta}^i = 1$, then regions $\mathcal{B}$ and $\mathcal{C}$ are empty and (7.27) gives the stationary density for all possible incumbents. In this case, $G^i_\emptyset$ is a solution to $\hat{G}^A_x(1) = 1 - x$, which exists and is unique since the LHS is increasing and satisfies $\hat{G}^A_0(1) = 0$.

Case 2. SPA with Tailing. Recall from Theorem 7 that $R^S(1) = \overline{\theta}^S$, so region $\mathcal{C}$ is always empty in the SPA. In this case, $G^S_\emptyset$ is a solution to $\hat{G}^A_x\left(\overline{\theta}^S\right) = \hat{G}^B_x\left(\overline{\theta}^S\right)$, which exists and is unique since the LHS is increasing and satisfies $\hat{G}^A_0\left(\overline{\theta}^S\right) = 0$ and the RHS is decreasing and satisfies $\hat{G}^B_1\left(\overline{\theta}^S\right) = 0$.

Case 3. FPA with Tailing. In this case $G^F_\emptyset$ is a solution to $\hat{G}^A_x\left(R^F(1)\right) = \hat{G}^C_x\left(R^F(1)\right)$, which exists and is unique since the LHS is increasing and satisfies $\hat{G}^A_0\left(R^F(1)\right) = 0$ and the RHS is decreasing and satisfies $\hat{G}^C_1\left(R^F(1)\right) = 0$.

Plugging $G_\emptyset$ in the three cases in the corresponding differential equations yields the desired stationary distributions.

We now show that $G_\emptyset$ in the SPA is non-monotonic in η in an example with $\theta \sim \mathcal{U}$, $\beta = 0$ and $n = 2$. In this case, $\frac{\mathrm{d}R_\emptyset}{\mathrm{d}\eta} = 0$, $R(\theta) = \theta$ in the tracking region, and $R(\theta) = \overline{\theta}^S = \frac{1+\eta}{4\eta}$ in the tailing region. Therefore, the reserve price effect (5.4) simplifies to

$$\eta \frac{\mathrm{d}\overline{\theta}^S}{\mathrm{d}\eta} 2\overline{\theta}^S \left(1 - G_\emptyset - G(\overline{\theta}^S)\right) = -\frac{1+\eta}{8\eta^2}\left(1 - G_\emptyset - G\left(\frac{1+\eta}{4\eta}\right)\right) < 0.$$

Moreover, (7.28) simplifies to

$$\mathrm{d}G^B(\theta) = \frac{\left[G_\emptyset + \eta(1 - G_\emptyset)\right] 2\theta + (1-\eta) G(\theta)}{1 - (1-\eta)\theta}$$

that yields

$$\hat{G}^B_{G_\emptyset}(\theta)=\frac{\eta\theta^2-G_\emptyset\left(1-(1-\eta)\theta^2\right)}{1-(1-\eta)\theta},$$

and (7.27) simplifies to

$$\mathrm{d}G^A(\theta)=\frac{[G_\emptyset+\eta G(\theta)]\,2\theta+(1-\eta)\,G(\theta)}{1-(1-\eta)\theta}. \tag{7.38}$$

Consider $\eta=\frac{1}{2}$, so that $\overline{\theta}^S=\frac{3}{4}$. Then (7.38) yields

$$\hat{G}^A_{G_\emptyset}(\theta)=G_\emptyset\frac{\exp(1-2\theta)\left(636+e^3\left(4\Gamma(5,4-2\theta)-\Gamma(6,4-2\theta)\right)\right)}{16(2-\theta)^5},$$

where $\Gamma(x,y)$ denotes the incomplete gamma function, and

$$\hat{G}^A_{G_\emptyset}\left(\frac{3}{4}\right)=G_\emptyset\frac{6\left(\frac{6784}{\sqrt{e}}-3821\right)}{3125},$$

while

$$\hat{G}^B_{G_\emptyset}\left(\frac{3}{4}\right)=\frac{1}{20}(9-23G_\emptyset).$$

Imposing smooth pasting $\hat{G}^A_{G_\emptyset}\left(\frac{3}{4}\right)=\hat{G}^B_{G_\emptyset}\left(\frac{3}{4}\right)$ (see case 2 above) yields $G_\emptyset\approx 0.263$. Since $G_\emptyset$ is (i) 0 when $\eta=0$ and (ii) $\left(r^M\right)^2=0.25$ when $\eta=1$ (the static benchmark), this proves non monotonicity. Analogous computations show that $G_\emptyset>0.25$ when $\eta>0.4$. This confirms the pattern displayed in the left panel of Figure 5.2.

Proof of Lemma 23

Consider $x_{t+1}-x_t$,

$$\begin{aligned}x_{t+1}-x_t&=y_{t+1}p_{t+1}-y_tp_t+(1-p_{t+1})z_{t+1}-(1-p_t)z_t\\&=p_t(y_{t+1}-y_t)+(y_t-z_t)(p_{t+1}-p_t)+(p_{t+1}-p_t)(y_{t+1}-y_t)\\&\quad+(1-p_t)(z_{t+1}-z_t)+(p_t-p_{t+1})(z_{t+1}-z_t).\end{aligned}$$

By assumption *ii*) all the sequences converge (and p_t converges to 1). This implies that the last three addenda are dominated by $\max\{p_t(y_{t+1}-y_t),(y_t-z_t)(p_{t+1}-p_t)\}$ for large t.[68] Therefore, for large t and up to first order,

$$x_{t+1}-x_t=p_t(y_{t+1}-y_t)+(y_t-z_t)(p_{t+1}-p_t).$$

Using the monotonicity assumption, $y_{t+1}-y_t<0$ and $p_{t+1}-p_t>0$. Using that the limit of $\{y_t\}$ is larger than the limit of $\{z_t\}$ we get that eventually $y_t-z_t>0$. Combining the inequalities we obtain that for t large $x_{t+1}-x_t<0$.

[68] Formally, $\max\{(p_{t+1}-p_t)(y_{t+1}-y_t),(1-p_t)(z_{t+1}-z_t),(p_t-p_{t+1})(z_{t+1}-z_t)\}=o(\max\{p_t(y_{t+1}-y_t),(y_t-z_t)(p_{t+1}-p_t)\})$, or equivalently $\lim_{t\to\infty}\frac{(p_{t+1}-p_t)(y_{t+1}-y_t)}{\max\{p_t(y_{t+1}-y_t),(y_t-z_t)(p_{t+1}-p_t)\}}=\lim_{t\to\infty}\frac{(1-p_t)(z_{t+1}-z_t)}{\max\{p_t(y_{t+1}-y_t),(y_t-z_t)(p_{t+1}-p_t)\}}=0$.