

Cued to Queue: Information in Waiting-Line Auctions

Jack Hirsch

Harvard University

Eric Tang

Stanford University

February 26, 2025

Abstract

We study the effect of providing information to agents who queue before a scarce good is distributed at a fixed time. When agents have quasi-linear utility in time spent waiting, they choose entry times as they would bids in a descending auction. An information designer can influence their behavior by providing updates about the length of the queue. Many natural information policies release *sudden bad news*, which occurs when agents learn that the queue is longer than previously believed. We show that sudden bad news can cause assortative inefficiency by prompting a mass of agents to simultaneously attempt to join the queue. As a result, if the value distribution has an increasing (decreasing) hazard rate, information policies that release sudden bad news increase (decrease) total surplus, relative to releasing no information. When agents face entry costs to join the queue and the value distribution has a decreasing hazard rate, an information designer maximizes total surplus by announcing only when the queue is full.

Keywords: queues, dynamic information design, auctions

JEL classification: D44, D61, D82, D83

Hirsch: Department of Economics, Harvard University, Cambridge, MA 02138, jhirsch1@g.harvard.edu, Tang: Graduate School of Business, Stanford University, Stanford, CA 94305, ertang@stanford.edu. We thank Itai Ashlagi, Dirk Bergemann, Modibo Camara, Matthew Gentzkow, Yannai Gonczarowski, Oliver Hart, Ravi Jagadeesan, Zi Yang Kang, Scott Duke Kominers, Shengwu Li, Annie Liang, Eric Maskin, Paul Milgrom, Michael Ostrovsky, Andrzej Skrzypacz, Tomasz Strzalecki, Takuo Sugaya, Bob Wilson, Frank Yang, Weijie Zhong, and seminar participants for helpful discussions and suggestions. Tang thanks the NSF Graduate Research Fellowship for financial support. Any errors are our own.

1. INTRODUCTION

Queues are often used to distribute a scarce good at a fixed time. Agents choose when to join the queue, trading off their probability of obtaining the good with the time they waste in line. The result is a “waiting-line auction,” in which the form of payment is time instead of money.¹ For example, when a food pantry or soup kitchen has limited rations to distribute, individuals may queue well in advance of the opening time to receive a meal. Prior to the release of the iPhone X, some prospective customers queued for days outside of Apple stores to secure the device.² Homeless shelters allocating beds, theatres selling popular concert tickets, and gas stations selling rationed fuel face similar queues. Our paper studies how access to real-time information about the number of people in line changes the welfare properties of such queues.

As a running example, suppose that a food pantry plans to distribute 50 parcels of food at noon, and more than 50 individuals desire a parcel. These agents face a trade-off: the earlier they arrive, the higher their probability of obtaining the good, but the more time they waste in line. Agents balance these two forces to determine when they will join the queue. However, information may alter their decisions. If an agent learns that there are already 45 people in line an hour before the pantry opens, she may hurry to join the queue before supply is exhausted. Conversely, if she discovers that there are only five people in line, she may wait to join until it is nearly noon, since more supply remains. Indeed, agents may rely on such information to guide their decisions, using social media or word of mouth to learn about the current queue length.

The distributor of the good may also choose to provide information. For example, the food pantry could release a social media post updating individuals when the line has reached a certain length, or they could post the length of the line every hour. More generally, technology such as social media and live cameras may enable distributors to provide real-time updates about the length of the queue. The choice of whether or not to release information, and what information to release, changes individuals’ queueing behavior.

¹Holt and Sherman 1982 are the first to refer to these queues as “waiting-line auctions.”

²Business Insider (2017): “In London and elsewhere in the world, queues of die-hard Apple fans are forming, prepared to camp out for a night or more to buy the eagerly anticipated tenth anniversary Apple smartphone” (Price 2017). See also Business Insider (2017), quoting an Apple press release: “iPhone X will be available in more than 55 countries and territories, and in Apple Stores beginning Friday, November 3 at 8:00 a.m. local time. Stores will have iPhone X available for walk-in customers, *who are encouraged to arrive early*” [emphasis added] (Leswing 2017).

In our model, the distributor of the good does not change the allocation scheme: the good is always allocated to the agents who arrive earliest, with ties broken randomly.³ However, the distributor can influence allocation probabilities by committing to an information policy that controls the flow of public information that agents receive about the queue length. Returning to our example, if the food pantry commits to announce the length of the line every hour, individuals may modify their queueing behavior in anticipation of, and in response to, the incoming announcements. We study how information affects the equilibrium allocation of the good, how long individuals wait, and total welfare.

Theorem 1 proves that under a condition we call *sudden bad news*, providing real-time information may break *assortative efficiency*: agents with the highest values may no longer receive the good. In equilibrium, agents form common beliefs about the number of individuals who have already joined the queue. When information leads to a jump in beliefs about the length of the queue, we say that agents received *sudden bad news*. Many natural real-world information policies have a positive probability of releasing sudden bad news. These include information policies that (i) announce the length of the queue at a fixed time, (ii) announce when the queue reaches a certain length, or (iii) announce every time an agent joins the queue.

The intuition behind our main result is simple: if agents suddenly learn that the queue is longer than expected, a “rush” of agents may simultaneously attempt to join. If the number of agents rushing the queue exceeds the number of items remaining, then the objects are randomly allocated among those agents, breaking assortative efficiency. In contrast, if the information policy commits to never release information, the equilibrium is assortatively efficient. Thus, by controlling the flow of information about the length of the queue, the distributor can affect both assortative efficiency and the total waiting times of agents.

Theorem 1 shows that a wide range of information policies induce assortative inefficiency. While Gretscho et al. 2014 study a similar question in an auction setting, they only analyze the “information policy” in which agents learn each time someone purchases an item (for us, each time someone enters the queue). We instead employ a novel proof technique to analyze a wide range of information policies. Our proof uses the Revenue Equivalence Theorem to abstract away from potentially complex equilibria induced by information policies, which allows us to

³In many cases, revealing real-time information (e.g., making an online announcement) may be logistically simpler than altering the allocation procedure (e.g., running a lottery among all agents present). Distributors may also be motivated by fairness concerns or concerns about the presence of speculators.

simplify our analysis to sudden bad news. Our result also sharpens a general intuition for why real-time information causes inefficiency: “bad news” can drive many agents to suddenly rush the queue.

Furthermore, we derive welfare implications that differ from those of the auction setting in Gretschno et al. 2014. In an auction, assortative efficiency maximizes seller revenue. A queue, on the other hand, is a money-burning mechanism; the “payment” is time spent waiting in line, which only reduces total surplus. As a result, total surplus is equivalent to consumer surplus, and assortative efficiency may not maximize consumer surplus. Theorem 1 suggests that a distributor should consider the impact of releasing such information, and we provide welfare results to guide her decision. We use Theorem 1 and existing work on money-burning mechanisms (e.g., Hartline and Roughgarden 2008, Condorelli 2012) to partially characterize when information policies increase or decrease total surplus.

Concretely, Theorem 2 states that if the distribution of values for the good has an increasing hazard rate (as in, for example, the uniform distribution), then any information policy that can release sudden bad news *increases* total surplus, relative to no information being released. If the value distribution has a decreasing hazard rate (as in a Pareto distribution), then any such information policy *decreases* total surplus. Intuitively, sudden bad news can cause assortative inefficiency, but it can also reduce agents’ wait times. The hazard rate determines which effect dominates the welfare analysis. When the value distribution has a thick tail, it is more important to screen for high-value agents; when it has a thin tail, it is more important to reduce wait times.

Section 5 models a queue in which agents face an entry cost when attempting to join. For example, agents may incur a cost in time or money to travel to the distribution site. In such settings, an information designer may have further reason to release information: information can prevent agents from wastefully paying the entry cost without obtaining the good. Corollary 2 shows that when the hazard rate of the value distribution is decreasing, an information designer maximizes total welfare by only announcing when the queue is full. Indeed, when the hazard rate is strictly decreasing, this information policy yields a strictly higher surplus than either releasing no information or announcing every time an agent joins the queue. Thus to maximize surplus, an information designer must selectively release an “intermediate” amount of information.

The paper is organized as follows. Section 1.1 describes related literature. Section 2 introduces our model and formally demonstrates the equivalence between queues and strictly descending

multi-unit auctions. Section 3 shows that sudden bad news causes assortative inefficiency. Section 4 partially characterizes how information affects total surplus. Section 5 considers an extension with entry costs. Section 6 concludes. All omitted proofs are in the appendix.

1.1. Related Literature.

Multi-Unit Descending Auctions. Our model is similar to a multi-unit descending auction. Holt and Sherman 1982 is the first paper to demonstrate that some queues, which they call “waiting-line auctions,” resemble Dutch auctions. However, they do not consider how information affects queueing behavior. The rush of agents that enter the queue following sudden bad news in our model is similar in spirit to Bulow and Klemperer 1994. They show that if prices continuously fall, information about a sale can induce agents to purchase immediately. However, they maintain assortative efficiency by allowing prices to increase to clear the market, whereas we model the “price” or entry time as strictly descending, reflecting the real-life constraints of queues.

The work most similar to ours is Gretschko et al. 2014, which studies a strictly-descending multi-unit Dutch auction. The price at which each object is sold is publicly announced, and therefore may induce a rush of agents. Gretschko et al. 2014 prove that any equilibrium is assortatively inefficient, and derive a symmetric equilibrium in a simple setting characterized by a set of initial value problems.

While the auction model of Gretschko et al. 2014 is isomorphic to our queue model, they only consider the full-disclosure information policy, which reveals each time an agent joins the queue. Our results apply to a much broader set of information policies, such as announcements at a fixed time or probabilistic announcements. The proof technique in Gretschko et al. 2014 does not extend to these information policies. Our work develops a new technique that uses the Revenue Equivalence Theorem, which also provides a novel intuition for the underlying cause of assortative inefficiency.⁴ Furthermore, because we analyze a queueing setting instead of an auction setting, information has different implications for welfare. We provide distinct results on how information affects total surplus, as well as associated policy recommendations for information designers.

Queueing Models. Many queueing models consider stochastic agent arrivals and a limited capacity to process arriving agents. The seminal early work is Naor 1969, and Hassin and Haviv

⁴In particular, we show that the phenomenon of “sudden bad news” leading to assortative inefficiency is a much more general occurrence than in the full-disclosure information policy.

2003 provide a survey. More recently, Leshno 2022 considers a setting with stochastic arrivals of agents who have heterogeneous preferences between two over-demanded items with waitlists, and demonstrates that information policies can increase total surplus. Related work on dynamic assignment includes Baccara et al. 2020 and Murra-Anton and Thakral 2024. Che and Tercieux 2024 consider a dynamic queue in which homogeneous agents arrive stochastically and decide whether to join the queue. In their model, the queue designer controls agents’ entry and exit from the queue, their queue priority, and the information they receive.

We model a fundamentally different type of queue. In our setting, all agents are present at the beginning of the game, all items are allocated at the same time, and each agent makes a single decision about when to join the queue. Our model describes a line that forms before a distributor (such as a food pantry or store) allocates goods at a fixed opening time. The dynamic queueing literature is more apt to model a setting such as a public housing waitlist, in which the waitlist’s length fluctuates as new housing becomes available and new agents arrive. In contrast, a queue forms in our setting not because agents arrive faster than items arrive, but instead because agents choose to join the queue earlier than the distribution time.⁵ Consequently, we employ analytic tools from auction theory rather than the stationary distributions used in the dynamic queueing literature. To our knowledge, our paper is the first to study the effects of information on queues that form due to a fixed distribution time and a fixed population of agents.

Money-Burning Mechanisms. A queue is an example of a “money-burning” or “ordeal” mechanism, in which agents are screened by a costly ordeal that does not transfer surplus (as opposed to screening via money). In such settings, there is a trade-off between wasteful costly screening and assortative efficiency. Dworzak 2023 compares screening with ordeal mechanisms to lump-sum transfers and Taylor et al. 2003 compares costly screening to lotteries, but neither work analyzes the role of information in queues. Hartline and Roughgarden 2008 and Condorelli 2012 construct consumer surplus-maximizing ordeal mechanisms for allocation problems. Theorem 2 combines Theorem 1 and the results in these papers to partially characterize how information policies change total surplus in our queueing model.

⁵In the work cited above, agents’ fundamental choices are whether to join the queue or not. In our model, agents’ fundamental choices are *when* to join the queue.

2. MODEL

2.1. Setup. Consider n risk-neutral agents indexed by $i \in \{1, \dots, n\}$. Agents have unit demand for one of $k < n$ identical copies of an item. Agents have private values v_i drawn independently and identically from distribution F , which we assume is atomless and has full support on $[0, \bar{v}]$.⁶ We fix model primitives n , k , and F unless stated otherwise, and we assume they are all common knowledge. Throughout the paper, $Y_k^{(n)} \sim F_k^{(n)}$ denotes the k^{th} order statistic of n draws from F .

Instead of paying to receive a copy of the good, agents queue in line. We assume agents have quasi-linear utility in time. We further assume that agents share a common disutility for each unit of time spent waiting in line, normalized to 1.⁷ Time is continuous and $t \in \mathbb{R}_+$ denotes the time remaining until the good is distributed, at time $t = 0$. If there are q agents in the queue, we say there are $k - q$ *items remaining*.

At any time, agents not in the queue may attempt to join the queue. An agent successfully joins the queue if there are still items remaining ($q < k$). If more agents attempt to join the queue than items remain, ties are broken uniformly at random. Each agent i obtains utility

$$u(v_i, t) = \begin{cases} v_i - t & \text{if } i \text{ successfully joins at } t \\ 0 & \text{otherwise.} \end{cases}$$

Our utility function implicitly assumes that (i) attempting to join the queue is costless; and (ii) once an agent joins the queue, she cannot leave.

The information designer commits to a commonly known and public information policy. The information policy maps from the game's history and time to an arbitrary message space \mathcal{M} and the time that message will be released. A message released at time t becomes instantaneously known to all n agents at time t . Natural information policies include announcing every time an agent joins the queue or announcing the length of the queue at a particular time. We let $\mathcal{Q}(p)$ denote the queueing game with information policy p .

⁶All results extend to distributions with unbounded support, so long the expected value of the distribution is finite. Even if the expectation of the distribution is infinite, all results still hold except for the strict welfare comparisons in Theorem 2 and Corollary 2. Further details are in Appendix B.

⁷We assume that agents have a homogeneous waiting cost, in which case all units of value can be normalized to value per time waited. If agents have heterogeneous waiting costs, one can normalize values to value per time waited. However, the implications for total welfare will then differ.

2.2. Game Procedure. Intuitively, our model is a descending clock game in which agents decide when to enter the queue. However, our model is complicated by the timing of interactions between messages and agents. We therefore formalize the model in a sequential format, with a finite number of “bidding” rounds. Throughout the paper, we interchangeably refer to the time at which an agent attempts to join the queue as his bid, often in the context of employing results from auction theory.

To summarize our formal model, at each stage, the announcer submits the next time she would like to make an announcement (conditional on no agent joining the queue before that time), and each agent submits the next time he would like to join the queue (conditional on no announcement arriving before that time). The next announcement or queue entry is processed, with the announcer winning ties. After an announcement, agents have the opportunity to instantaneously join the queue in response to the announcement. Similarly, after an agent joins the queue, the announcer has the opportunity to instantaneously make an announcement.

The formal game proceeds as follows. Initialize the clock to time $T = \bar{v}$.⁸ Initialize the number of available copies of the good to $K = k$. As the game progresses, we use T and K to denote the game clock and number of remaining items, respectively. We use she/her pronouns for the announcer and he/him pronouns for agents. The order of play is as follows:

Step 1: Submitting Bids. Each agent i still in the game submits a sealed bid $b_i \in [0, T]$. This is the time at which he will attempt to enter the queue conditional on receiving no new announcements before time b_i .

Step 2: Submitting Announcements. The announcer draws her next (possibly randomized) message and message time (m, t) with $t \in [0, T]$. If no agents join the queue before t , she will release message m at time t . Her draw does not depend on the sealed bids from Step 1.

Step 3: Resolving Bids and Announcements. We compare the highest bid $b := \max_i \{b_i\}$ to the announcement time t . First set the clock to $T = \max\{b, t\}$. There are two cases, depending on which party moves first; both begin the following nested game. Intuitively, in the nested game, the announcer and the agents alternate actions at time T , until either no further agents wish to join the queue at time T , or the announcer no longer wishes to send messages at time T . At that point, the nested game ends and the main game continues.

⁸A start time of \bar{v} is sufficiently large since any agent entering before would receive a strictly negative payoff.

- Case 1: The announcer’s message arrives first: $t \geq b$.⁹
 - (a) All agents observe announcement m . Each agent then has a binary choice: he may either attempt to join the queue at time T , or choose to wait. If no agent attempts to join, the game returns to Step 1.
 - (b) Agents attempting to join the queue are either successful or unsuccessful. Successful agents are allocated the good at cost T ; unsuccessful agents obtain payoff zero.
 - (c) If any agent joined the queue, then the announcer has the chance to make another announcement at time T . If she does, the game returns to (a) of Case 1; otherwise, the game returns to Step 1.
- Case 2: An agent joins first: $b > t$.

First, agents attempting to join the queue are either successful or unsuccessful, as in Case 1. Then the announcer has the chance to make another announcement, still at time T . If she does, the game proceeds from phase (a) of Case 1. Otherwise, the game returns to Step 2: the announcer draws a new announcement, but agents do not update their bids.¹⁰

Remark. The following three assumptions are equivalent: if the queue is full, then (1) the game formally ends, (2) the announcer reveals that the queue is full, or (3) all agents not currently in the queue receive payoff zero. The reason they are equivalent is that in all three cases, all agents not in the queue receive payoff 0, independent of their actions. For convenience, we assume the game ends when the queue is full. In Section 5, we model agents who pay an entry cost to join the queue. In that case, their behavior changes depending on whether or not the information designer announces when the queue is full.

2.3. Discussion of Game Procedure. Our baseline model assumes that attempting to join the queue is costless. In Section 5 we loosen this assumption to accommodate potential entry costs for joining the queue, such as time or money spent to travel to the queue. We also assume that when an agent attempts to join the queue, he immediately learns if the queue is full. This assumption is plausible when agents know how many items are available for distribution when they join a physical queue. In many cases, an agent joining a queue can observe how many

⁹In our model, the announcer wins ties. Intuitively, having the announcer win ties allows agents to instantaneously react to the absence of an announcement, but prevents them from preempting one.

¹⁰If the model instead gave agents the chance to update their bids, they would become aware that an agent had joined the queue, even if there were no announcement.

others are currently in line. In such a case, he can determine whether that queue length exceeds the number of items available.

Our model assumes that ties between agents are broken uniformly at random: if two items remain and three agents attempt to join the queue simultaneously, each has a two-thirds chance of success, then the game ends. This assumption is motivated by the reality of physical queues. In most real-world queues, if several agents arrive to a queue at the same time, their order in line is essentially random. Stated differently, agents who arrive simultaneously can no longer be screened by their willingness to wait for the good. Our uniform tiebreaking assumption could also model a short, stochastic delay between when agents attempt to join the queue and when they actually arrive at the queue.

We have also assumed agents cannot leave the queue and rejoin at a later time. We believe that in these settings in practice, while some individuals choose to leave and rejoin the queue, most individuals who arrive at the queue (before it is full) stay at the queue until the distribution time. In our model without entry costs, allowing agents to exit and rejoin the queue would also allow them to costlessly observe the queue length, rendering information policies meaningless. In a model with entry costs, agent behavior can become quite complex if they can leave and rejoin the queue. For example, an agent who joins the queue and observes that there are few other agents currently in the queue may find it optimal to leave the queue and rejoin at a later time. For these reasons, we simplify agents' behavior by assuming they cannot exit and rejoin the queue.

2.4. Information Policies. We describe information policies and our equilibrium concept with verbal descriptions in this section, and we provide mathematical formulations in the appendix.

The *message history* captures the content of past announcements and the time at which they were sent, and is visible to both the announcer and the agents. The *queue history* captures the history of past arrivals to the queue, and is visible only to the announcer. An information policy can depend on both histories.

Definition 1 (Information Policy). An *information policy* specifies the time and content of publicly announced messages as a function of the history and current time.

We restrict our attention to public information policies. Information policies may be randomized or deterministic, and must send a finite number of messages. The baseline information

policy which we use for comparison is the “trivial” information policy that never makes an announcement.

Example 1 (Trivial). The announcer never makes an announcement.

We also include two natural examples of substantive information policies. Note that information policies can also be followed probabilistically, so that announcements may or may not arrive.

Example 2 (Fixed Time). The announcer reveals the current queue length at time τ .

Example 3 (Full Revelation). The announcer reveals the updated queue length each time an agent joins the queue. As a result, agents always know the length of the queue.

2.5. Equilibrium. We study Perfect Bayesian Equilibria, defined by agents’ strategies and beliefs. Each agent follows a strategy that maps his value, the current message history, and the current time to a bid. In equilibrium, all agents derive beliefs $\mu \in \Delta(\{1, \dots, k\})$ about the the number of items remaining, determined by the information policy and model primitives, the message history, and the current time. Because information policies are public, all agents who have not yet joined the queue share a common belief μ at any time. Agents have beliefs about the *positive* quantities of items remaining because the game ends when no items remain.¹¹

We index agents’ beliefs in a way that allows us to describe limiting beliefs. We let μ_{t+} denote agents’ beliefs about the number of remaining items at time t , but *before* the announcer has had the opportunity to send any message at time t .¹² We let μ_{t-} denote agents’ beliefs about the number of remaining items at time t immediately *after* the announcer has had the chance to send a message.¹³

Our main analysis studies the equilibrium properties of queueing games defined by various information policies. We primarily focus on *assortative efficiency* and total surplus.

¹¹We also define “sudden bad news” with respect to agents’ beliefs about the positive quantities of items remaining.

¹²In our sequential game procedure, μ_{t+} can be defined as follows. Consider the latest round that occurs strictly before time t . Then μ_{t+} is an agent’s belief in this round about the number of items currently remaining, adjusted downward by their belief about the number of agents who will submit bids strictly greater than t , conditional on no further messages being announced before time t . Then μ_{t+} describes agents’ beliefs about the items remaining at time t but before knowing whether a message at time t has arrived.

¹³When necessary, we continue in this fashion to index further rounds of the nested game occurring at the same time-step. We use t^{--} to index beliefs after the announcer has had two chances to send messages at time t . That is, t^{--} indexes beliefs immediately after the second instance that step (a) is reached at time t , if it is reached. We then use t^{---} for the third instance that step (a) is reached at time t , and so forth.

Definition 2 (Assortatively Efficient). An equilibrium is *assortatively efficient* if the k agents with highest values are always allocated the good.

Since a queue is a money-burning mechanism, time spent waiting in line is lost surplus. As a result, an assortatively efficient equilibrium does not necessarily maximize total surplus: the screening costs necessary to ensure assortative efficiency may outweigh the surplus gains of allocating the items to the agents with the highest values. In Section 4 we analyze how information policies affect total surplus.

3. RESULTS

3.1. Benchmark: No Information. We begin by analyzing Example 1, the benchmark case in which there can never be an announcement. Since there are never announcements, the setting is isomorphic to a multi-unit sealed-bid discriminatory price auction. Because agents do not learn new information, an agent's strategy is equivalent to choosing, at the beginning of the game, the time at which he will enter the queue. This is the same as choosing a bid to maximize expected profit in a discriminatory price auction. As in a discriminatory price auction, agents face a trade-off between their probability of winning and their surplus conditional on winning. The only difference between the queue model and the auction model is the form of payment: money or time.

Reny 1999 proves that a multi-unit sealed-bid discriminatory price auction has a symmetric, assortatively efficient equilibrium characterized by the bidding function¹⁴

$$(1) \quad \beta(v) = \mathbb{E}[Y_k^{(n-1)} | Y_k^{(n-1)} < v].$$

In any equilibrium of this auction, agents cannot pool.¹⁵ Hence any equilibrium is assortatively efficient, and by the Revenue Equivalence Theorem, payments must be defined by (1). The equilibrium is therefore unique.

Corollary 1. *If information policy p is trivial, then the unique equilibrium of $\mathcal{Q}(p)$ is assortatively efficient.*

¹⁴As in our setting, Reny 1999 considers agents with unit demand and in which values are private and drawn independently and identically. See Krishna 2009.

¹⁵If agents pool, agents in the pooling region could profitably deviate by increasing their bid by some $\varepsilon > 0$.

In our benchmark, releasing no information results in assortative efficiency, as the agents with highest values join the queue earliest. Theorem 1 shows that a wide range of information policies instead result in an assortatively inefficient allocation.

3.2. Main Results. To analyze how information policies affect the ultimate allocation, we define a notion of “sudden bad news” and show that many natural information policies release sudden bad news with positive probability. Theorem 1 then shows that if an equilibrium has a positive probability of releasing sudden bad news, it is assortatively inefficient.

Definition 3. We say *sudden bad news* is revealed at time t if (i) μ_{t-} is strictly FOSD dominated by μ_{t+} and (ii) μ_{t+} is weakly FOSD dominated by $\mu_{t'}$ for all $t' > t$.

Condition (i) holds when agents learn that there are strictly fewer objects remaining than they had previously believed. Intuitively, this leads agents to bid more aggressively, but the current time t imposes a cap on how aggressively agents can bid. Condition (ii) ensures that μ_{t-} is the “most pessimistic belief” agents have had in the game thus far.¹⁶ Many intuitive information policies have a positive probability of releasing sudden bad news in any equilibrium:

- Example 2 (fixed time). The announcement has a positive probability of revealing that one item remains, which FOSD dominates any previous belief. Even if the announcement occurs probabilistically or after a delay, there is still a positive probability of sudden bad news.
- Example 3 (full revelation). Every announcement is sudden bad news. Agents always know the exact queue length, so learning that a new agent entered is sudden bad news. Delays or probabilistic announcements still result in sudden bad news.¹⁷
- An information policy that reveals when exactly k' items remain. Before the announcement, agents know that there must be strictly more than k' items remaining, so the announcement reveals sudden bad news.

Theorem 1. *Fix information policy p . Any equilibrium of $\mathcal{Q}(p)$ with a positive probability of revealing sudden bad news is assortatively inefficient.*

¹⁶Note that sudden bad news can also occur in the absence of an announcement. For example, suppose the designer commits to send one message at time τ if there are fewer than five agents in the queue. If no message arrives at time τ , agents learn there are at least five agents in the queue, and their beliefs update discontinuously.

¹⁷If the announcements occur probabilistically or with a delay, there is still a positive probability of announcing that only one item remains, which is sudden bad news.

Intuitively, sudden bad news reveals that the good is scarcer than previously believed, driving agents with high values to bid more aggressively (join the queue earlier). However, since the clock is already at τ , these agents cannot bid higher than τ , and some are forced to pool by bidding τ . To formalize the intuition, we consider the game immediately after sudden bad news is revealed, and show that conditional payments in the paused queueing game are not equal to conditional payments in the assortatively efficient mechanism. Then by the Revenue Equivalence Theorem, the allocation of the queueing game cannot be assortatively efficient.

To prove Theorem 1, we first develop a series of lemmas. Fix information policy p . By the Revelation Principle, any equilibrium of $\mathcal{Q}(p)$ corresponds to an incentive-compatible direct mechanism. We will use the Revenue Equivalence Theorem to compare the equilibrium payments of $\mathcal{Q}(p)$ to the equilibrium payments of an assortatively efficient mechanism. We first formally demonstrate that the Revenue Equivalence Theorem holds when there is supply and demand uncertainty.

Lemma 1. *[Revenue Equivalence Theorem] Fix agents' beliefs about the number of remaining items and agents. Then any two mechanisms with the same allocation rule have the same expected payments for all types.*

We use Lemma 1 to characterize a symmetric equilibrium bidding function for a discriminatory price auction.¹⁸ Once the game begins, agents have beliefs about the remaining supply and demand, pinned down by their beliefs about the number of objects remaining. Since our main proof technique considers intermediate stages of the game, it is useful to establish various properties when there is demand and supply uncertainty. In this case, the equilibrium bidding function is given by the weighted average of the bidding functions in Equation (1). However, the weighted average is not just over the probability the state occurs, but also over the probability the agent wins in that state.

Proposition 1. *Consider a discriminatory auction with unit-demand and uncertainty about supply (number of items available) and demand (number of bidders). There exists an assortatively efficient equilibrium in which agents bid the weighted average of the equilibrium bidding functions for discriminatory price auctions with known supply and demand, given by Equation (5).*

¹⁸One consequence of Lemma 1 is that if an equilibrium of queueing game $\mathcal{Q}(p)$ has the same allocation rule as an equilibrium of queueing game $\mathcal{Q}(p')$, then the two equilibria have the same expected payments for all types.

The proof and exact bidding function (Equation (5)) are in the appendix. We use Proposition 1 to characterize transfers in the queueing game with more complex information policies.

Definition 4. Suppose agents have beliefs μ and values drawn independently and identically from $F|_{[0,\hat{v}]}$. Let $t_{AE}(v|\mu, \hat{v})$ denote the expected payment of type v conditional on winning in the assortatively efficient mechanism. Let $t_Q(v|\mu, \hat{v}, p)$ denote the same for an equilibrium of the queueing game with information policy p .

We suppress p in the notation for t_Q when clear.

Lemma 2. *The expected transfer of the highest remaining type in an assortatively efficient mechanism, $v \mapsto t_{AE}(v|\mu, v)$, is continuous and strictly increasing.*

Lemma 2 says that if we fix beliefs μ and values are drawn independently and identically from some $F|_{[0,v]}$, then the expected transfer in an assortatively efficient mechanism for an individual with the highest possible type v is continuous in v . Note that changing v also changes the support over which values are drawn. The next lemma describes how expected payments in the assortatively efficient mechanism change as beliefs change.

Lemma 3. *Suppose that belief μ strictly FOSD dominates μ' . Then $t_{AE}(v|\mu', v) > t_{AE}(v|\mu, v)$ for all v .*

Because μ FOSD dominates μ' , there are fewer items available under μ' . As a result, an agent with the highest type must pay more conditional on winning (and always wins) in equilibrium. In the equilibrium of a discriminatory price auction, such an agent would bid more aggressively.

Formally, from Proposition 1 and Equation (5), an agent's conditional payment is the average of their conditional payment in each state weighted by both the probability that the state occurs and the probability that the agent wins in that state. Since the highest type always wins in an assortatively efficient mechanism, the payment of the highest type is greater under μ' than under μ .

We use t_{AE} to pin down the time at which each type must join the queue, given their beliefs as the game progresses.

Lemma 4. *Suppose $\mathcal{Q}(p)$ has an assortatively efficient equilibrium. In this equilibrium, type v attempts to enter the queue at some $\tilde{t}' \in \{t^+, t^-, t^{--}, \dots\}$ if and only if $t_{AE}(v|\mu_{\tilde{t}}, v) = t$ for some $\tilde{t} \in \{t^+, t^-, t^{--}, \dots\}$.*

In an assortatively efficient equilibrium, an agent's payment conditional on winning is determined by his beliefs, his value, and the highest type remaining in the game. This in turn determines the time at which he attempts to enter the queue, because that time is also his payment conditional on winning.

We use Lemma 1, Lemma 2, Lemma 3, and Lemma 4 to show that if information policy p has a positive probability of sudden bad news, then transfers in $\mathcal{Q}(p)$ cannot be equal to transfers in an assortatively efficient mechanism.

Proof of Theorem 1. Fix an equilibrium of $\mathcal{Q}(p)$ with a positive probability of revealing sudden bad news. For the sake of contradiction, assume that the equilibrium is assortatively efficient. Suppose sudden bad news is revealed at time t .

Consider two paused versions of $\mathcal{Q}(p)$: (i) the game immediately before sudden bad news is revealed, when agents have beliefs μ_{t+} and (ii) the game immediately after, when agents have beliefs μ_{t-} . Let \hat{v} denote the highest type that could still be in the game before the announcement (determined by agents' equilibrium bidding strategies), so that agents have types drawn from $F|_{[0, \hat{v}]}$. Since the equilibrium is assortatively efficient, only one type can join in any round of the nested game. Thus after the announcement the remaining agents have types in $[0, \hat{v})$ or $[0, \hat{v}]$.

By assumption, the equilibria of both games must still be assortatively efficient. By Lemma 1, each type's expected transfer conditional on winning in the queueing game must be equal to their expected transfer conditional on winning in the assortatively efficient mechanism, given their respective beliefs. In notation, $t_Q(v|\mu, \hat{v}) = t_{AE}(v|\mu, \hat{v})$ for any v , μ and \hat{v} .

We claim $t_{AE}(\hat{v}|\mu_{t+}, \hat{v}) = t$. Assume not. Since the clock is already at time t , type \hat{v} cannot pay more than t , so $t_{AE}(\hat{v}|\mu_{t+}, \hat{v}) \leq t$. By Lemma 2, we may choose $v' > \hat{v}$ sufficiently close to \hat{v} such that $t_{AE}(v'|\mu_{t+}, v') < t$. By the definition of sudden bad news and Lemma 3, for all $t' > t$ we know $t_{AE}(v'|\mu_{t'}, v') < t < t'$.¹⁹ By Lemma 4, type v' attempts to enter the queue at time t' if and only if $t_{AE}(v'|\mu_{t'}, v') = t'$. But then type v' cannot have joined the queue before time t and is still in the game. This contracts the maximality of \hat{v} . We conclude that $t_{AE}(\hat{v}|\mu_{t+}, \hat{v}) = t$.

¹⁹Intuitively, our theorem condition states that these are the most pessimistic beliefs agents have had thus far in the game, hence all prior expected payments (conditional on winning) were lower than the current expected payment (conditional on winning).

By Lemma 3, $t_{AE}(\hat{v}|\mu_{t-}, \hat{v}) > t_{AE}(\hat{v}|\mu_{t+}, \hat{v}) = t$. By the continuity of t_{AE} there exists some type $v < \hat{v}$ with $t_{AE}(v|\mu_{t-}, \hat{v}) > t$. That means type v is still in the game but satisfies

$$t_Q(v|\mu_{t-}, \hat{v}) = t_{AE}(v|\mu_{t-}, \hat{v}) > t.$$

A type- v agent has positive probability of winning the item, and if she does, she must pay $t_Q(v|\mu_{t-}, \hat{v}) > t$ in expectation. Since the game clock has already reached t , her payment cannot exceed t , a contradiction. ■

We have presented many examples of information policies that are assortatively inefficient, but we also believe it is possible to have (non-trivial) information policies that maintain assortative efficiency. As suggested by Theorem 1, the driving force behind assortative inefficiency is the instantaneous release of bad news, which causes a mass of agents to rush to join. If an information policy always releases bad news continuously, we conjecture that there is never a rush. Example 4 considers an information policy that (1) provides bad news continuously and (2) can only announce that the queue is empty (which is the best possible news).

Example 4 (Continuous Bad News). Fix $\lambda > 0$. While the queue is empty and no messages have yet been sent, the announcer sends a message at a $Poi(\lambda)$ rate.

Consider the game with information policy defined by Example 4. As time passes and a message does not arrive, agents continuously update their beliefs about the queue length, with increasingly pessimistic views about the probability that no one has yet joined. If a message arrives, it signifies good news and should relax bidding functions.

4. WELFARE

We next analyze how information affects total welfare. Descriptively, this can help predict the welfare impacts of decentralized information arriving from social media and the internet. More importantly, our results may help an information designer increase total surplus with a well-designed information policy. Suppose the distributor of a good wishes to allocate a good using a queue instead of, for example, a lottery. This may be due to logistical concerns (a lottery may be more difficult to organize), concerns about allocation to very-high-need individuals (such as in the distribution of food aid), or concerns about the presence of speculators. In such settings, when an organization wishes to maintain a queue but is willing to release real-time information about that queue, our results can guide what information should be released.

Theorem 1 allows us to characterize some cases in which providing information increases or decreases total surplus. Intuitively, when information breaks assortative efficiency, it reduces the value generated by the allocation, but it also reduces total waiting times. The hazard rate of the value distribution determines which effect dominates. Theorem 2.9 of Hartline and Roughgarden 2008 allows us to characterize when assortative efficiency maximizes or minimizes surplus.

Lemma 5. *(Hartline and Roughgarden 2008.) Consider the class of IC and IR money-burning mechanisms that allocate all items. If the hazard rate of the value distribution is weakly increasing (decreasing), then any mechanism with an assortatively inefficient allocation has a weakly higher (lower) total surplus than the assortatively efficient mechanism. If the hazard rate is strictly monotone, the inequality is strict.*

To our knowledge, no prior work has connected this result to information design for queueing. Lemma 5, together with Theorem 1, yields Theorem 2.

Theorem 2. *Suppose the hazard rate of the value distribution is weakly increasing (decreasing). Fix information policy p and an equilibrium of $\mathcal{Q}(p)$. If in equilibrium there is a positive probability of sudden bad news, then total surplus is weakly higher (lower) than under a trivial information policy. If the hazard rate is strictly monotone, this inequality is strict.*

Proof. By the Revelation Principle, any equilibrium of the queueing game is equivalent to a direct-revelation mechanism that is incentive compatible and individually rational. By Corollary 1, providing no information results in assortative efficiency. By Theorem 1, any equilibrium of $\mathcal{Q}(p)$ that may reveal sudden bad news is assortatively inefficient. If the hazard rate of the value distribution is increasing, then by Lemma 5, assortative efficiency minimizes total surplus, hence any equilibrium of $\mathcal{Q}(p)$ yields higher total surplus. The results for a decreasing hazard rate and strict monotonicity follow identically from Lemma 5. ■

Theorem 2 provides a policy recommendation for an information designer who knows the value distribution of her customers. If the hazard rate is increasing, she can increase consumer surplus by providing information, and if the hazard rate is decreasing, she maximizes consumer surplus by providing no information. Our results also suggest that an information designer may find it worthwhile to learn the distribution of values of her customers to guide her decisions.

As discussed in subsection 2.3, Theorem 2 depends on several modeling assumptions about the queueing environment, including the assumption that agents can costlessly join the queue.

In Section 5, we develop a model that uses the same framework, but incorporates entry costs. We derive different implications for how an information designer should optimize total surplus. In particular, maintaining assortative efficiency may require releasing an *intermediate* level of information: disclosing when the queue is full, but not revealing sudden bad news.

5. ENTRY COSTS

In many real-world situations, agents face an upfront cost to join a queue. When a food bank distributes aid or a store releases an over-demanded product, agents physically travel to join the queue, which can be costly in money or time. In the absence of information, agents risk incurring the entry cost without securing the good, which may reduce consumer surplus. Furthermore, some low-value agents may choose not to queue at all, which results in wasted surplus when unclaimed items remain. An information designer may share information about the queue’s length to avoid wasted entry costs or unallocated items, but as we showed in Section 4, information can have subtle effects on welfare.

In this section, we show that if agents face entry costs, and the value distribution has a strictly decreasing hazard rate, then the information designer maximizes surplus by revealing only some information. In particular, Corollary 2 states that one optimal information policy is to reveal when the queue is full, and nothing else. The resulting equilibrium has strictly higher welfare than revealing no information, and also has strictly higher welfare than any policy which reveals sudden bad news (such as full revelation, revealing the queue length at a fixed time, or announcing when the queue reaches a certain length).

To model entry costs, suppose that all agents face a homogeneous entry cost $c \in (0, \bar{v})$ to join the queue. If an agent attempts to join the queue, they first pay cost c , then the game proceeds as in the main model.²⁰ Without entry costs, announcing when the queue is full does not affect payoffs or the equilibrium outcome: if the queue is full, agents obtain zero utility regardless of when or whether they attempt to join the queue. With entry costs, if the queue is full, agents lose utility if they attempt to join. As a result, agents change their behavior if the announcer discloses when the queue is full.

²⁰For tractability, we assume that the cost is paid instantly. If the cost is paid in time, there may be a delay between when an agent receives information and when they arrive at the queue. If the cost is paid in money, there is no delay, although consequences for total surplus may differ: what we describe as total surplus is now consumer surplus.

Modeling entry costs also changes the assortatively efficient allocation. In the entry cost model, an equilibrium is assortatively efficient if the k agents with highest values above c , or all such agents if there are fewer than k , receive the good. When there are entry costs, releasing no information is no longer assortatively efficient: some agents who could have profitably entered the queue choose not to, due to uncertainty about whether they will obtain the item.

Proposition 2. *Suppose agents face an entry cost $c \in (0, \bar{v})$ and a trivial information policy. Then in equilibrium, the good is allocated to the (up to) k agents with highest values above some $v^R > c$. In particular, the allocation is not assortatively efficient.*

To achieve assortative efficiency, we consider an information policy that discloses when the queue is full (and nothing else). By announcing when the queue is full, agents with values close to c still have an incentive to join before they receive an announcement.

Proposition 3. *Suppose information policy p discloses when the queue becomes full and makes no other announcements. Then the unique equilibrium of $\mathcal{Q}(p)$ is assortatively efficient.*

If the announcer reveals “too much” information, in the sense formalized in Definition 3, we again lose assortative efficiency due to a mass of agents potentially rushing the queue. Thus to maintain assortative efficiency, the announcer must release some information (disclosing when the queue is full) without disclosing sudden bad news. As described in Section 3, a wide variety of natural information policies can therefore break assortative efficiency.

Proposition 4. *Consider the entry-cost queue game with an information policy p that always reveals when the queue is full. It may also make other announcements. If in equilibrium $\mathcal{Q}(p)$ has a positive probability of revealing sudden bad news, then the equilibrium is assortatively inefficient.*

If the announcer always discloses when the queue is full, we can reinterpret the game as a standard queueing game in which there are no entry costs and agents’ values are reduced by c . Sudden bad news is still defined relative to beliefs over the *positive* number of items remaining. Theorem 1 then implies that there is no assortatively efficient equilibrium. To connect assortative efficiency and welfare in the entry cost setting, we establish a lemma that is nearly identical to the results in Hartline and Roughgarden 2008.

Lemma 6. *Suppose agents' values are distributed with support $[\underline{v}, \bar{v}]$ for $\underline{v} \leq 0$ and $\bar{v} \geq 0$. Consider the class of IC and IR money-burning mechanisms. If the hazard rate of the value distribution is weakly (strictly) decreasing, then an assortatively efficient mechanism yields a weakly (strictly) higher total surplus than any mechanism in this class that is not assortatively efficient.*

The proof is nearly identical to that of Theorem 2.9 of Hartline and Roughgarden 2008. The only differences are that (i) we allow for agents with negative values and (ii) we maximize surplus over all possible mechanisms, not only over mechanisms that always allocate the good.²¹

Combining Proposition 2, Proposition 3, and Proposition 4 with Lemma 6 yields a condition under which the information designer prefers to release an intermediate amount of information.

Corollary 2. *Suppose the hazard rate of the value distribution is weakly decreasing. Then the information policy that discloses when the queue becomes full, and makes no other announcements, yields a weakly higher total surplus than any other information policy. If the hazard rate is strictly decreasing, this information policy yields a strictly higher total surplus than the trivial information policy, or any information policy that reveals sudden bad news.*

Returning to our leading example of the food pantry, suppose that the hazard rate of the value distribution is indeed strictly decreasing, and agents face an entry cost to join the queue. Suppose that the food pantry wishes to maximize total surplus by altering the information it releases about the queue length. The food pantry can do so by announcing only when the queue is full; furthermore, this policy should be *strictly* preferred to a wide range of other natural policies. It is strictly suboptimal for the food pantry to reveal no information. Yet it is also strictly suboptimal to fully reveal the length of the queue, to disclose the queue length at a fixed time, or to announce when the queue reaches a fixed length (other than k). These results show that information has a nuanced effect on total surplus, and may require organizations to carefully consider what information they release. Our results can guide organizations in making their decision.

²¹Indeed, when $\underline{v} < 0$, IR mechanisms do not always allocate the good. Negative values model agents for whom the entry cost exceeds their value.

6. CONCLUSION

We study how access to information affects welfare when queues form in advance of the distribution of a good. Under a condition we call “sudden bad news,” releasing information leads to an assortatively inefficient allocation. We partially characterize when providing real-time information increases or decreases total surplus, and we show that the effect on total surplus depends on the hazard rate of the value distribution. When agents face an entry cost to join the queue, an information designer may maximize surplus by revealing when the queue is full and not revealing further information.

One direction for future work is to determine the surplus-maximizing information policy for a given distribution of values. If the hazard rate is increasing, Theorem 2 shows that releasing sudden bad news increases surplus relative to releasing no information, but even in simple examples it is difficult to determine the optimal policy. There are many ways to extend our model as well. When there is uncertainty about the initial supply of items, agents have multi-dimensional beliefs and our “sudden bad news” criterion no longer necessarily induces a rush. Similarly, if agents do not learn the length of the queue when they arrive, they may wait without ultimately receiving the good; the model is similar to an all-pay auction (instead of a discriminatory price auction). There are also settings with multiple queues: for example, homeless shelters in the same city may each have a queue, or multiple stores may distribute the same scarce good. One could also model a setting in which those earlier in line receive a higher-quality good or face a shorter processing time. Finally, future work could expand the set of permissible information policies and actions by modeling costly information acquisition, private messages to a subset of agents, or agents who can leave and re-enter the queue.

APPENDIX A. DEFINITIONS

The *message history* $h^M = \{(m_i, t_i)\}_{i=1}^j$ captures the log of past messages, where (m_i, t_i) signifies that the i^{th} message m_i was announced at time t_i . We denote the space of all possible message histories \mathcal{H}^M . The message history is visible to both the announcer and the players.

The *queue history* h^Q captures the history of past arrivals to the queue, and is visible only to the announcer. We denote a generic queue history as $h^Q = \{t_1, \dots, t_m\} \in \mathcal{H}^Q$, where t_i denotes the time that the i^{th} agent in the queue joined the queue.

Whereas the message history is public knowledge, the queue history is not. We therefore combine h^M and h^Q into one history $h \in \mathcal{H}$ only visible to the announcer that records the order of messages and queue entrances.

Definition 5 (Information Policy). An *information policy* specifies the time and content of messages as a function of the history and current time. Formally, an *information policy* is a function

$$p : \mathbb{R}_+ \times \mathcal{H} \rightarrow \Delta(\mathcal{M} \times \mathbb{R}_+),$$

where $p(t, h)$ is the distribution of possible messages sent when the current time is t and the history is h . The announcer cannot send messages earlier than the current time; any tuple (m', t') in the support of $p(t, h)$ has $t' \leq t$.²² We also restrict information policies to send a finite number of messages over the course of the game. Most explicit information policies we consider are deterministic, in which case we write p as a function mapping to the space $\mathcal{M} \times \mathbb{R}_+$.

Definition 6 (Strategy). A strategy s_i for player i is a map from the current message history, time, and player value to their bid. Formally,

$$\begin{aligned} s_i : \mathcal{H}^M \times \mathbb{R}_+ \times \mathbb{R}_+ &\rightarrow \mathbb{R}_+ \\ (h^M, t, v) &\mapsto b \end{aligned}$$

Each strategy implicitly assumes a fixed information policy p , so that $s_i(h^M, t, v|p) = b$.

Definition 7 (Equilibrium). We consider Perfect Bayesian Equilibrium (PBE). If strategy profile S^* forms a PBE of $\mathcal{Q}(p)$, then

- (i) For each player i , strategy s_i^* is a best response to s_{-i}^* , the common belief μ , and current message history h^M .
- (ii) The on-path common belief μ is correctly derived via Bayes' Rule. In particular, let X_t be a random variable denoting the number of items remaining at time t . Let $\mu_t(j)$ denote the probability that j items remain, given beliefs μ_t . Then $\mu_t(j)$ is given by

$$(2) \quad \mu_t(j) := Pr(X_t = j | h^M, t) = \frac{Pr((h^M, t), X_t = j)}{Pr(h^M, t)},$$

²²The message is only announced if it conforms with the time restrictions given in the game procedure: namely, $t' < t$ in the main stage of the game, or $t' \leq t$ in the nested game.

When $Pr(h^M, t) = 0$, we replace the probability mass expressions with their respective probability densities.

APPENDIX B. UNBOUNDED SUPPORT

Our results also extend to distributions F with support on $[0, \infty)$. Let $Y \sim F$. Then our results hold exactly as written so long as $\mathbb{E}[Y] < \infty$. If $\mathbb{E}[Y] = \infty$, then the “strict” welfare results in Theorem 2 and Corollary 2 do not hold, because the expected total surplus from these equilibria are infinite. Note that extending our results to unbounded F is a meaningful extension: the only distributions with monotonically decreasing hazard rate have unbounded support.

All results in Section 3 are valid as written. The only change to the game procedure is that the clock begins at time $T = \infty$, allowing agents to submit arbitrarily high bids in the first round, but the game procedure still has a finite number of rounds. No further regularity conditions are needed, because all expectations used in the proofs of Section 3 are conditional expectations of bounded random variables, and are therefore finite. In addition, the applications of the Envelope Theorem are still valid with unbounded support.

In Section 4 and Section 5, the finite expectation ensures that the strict inequalities hold. For general distributions with unbounded support, if $\mathbb{E}[Y] = \infty$, then the expected total surplus is infinite as well. For example, consider the Pareto distribution with $\alpha \leq 1$. Lemma 5 shows that randomly allocating all items weakly minimizes surplus. However, that policy would yield expected surplus $\frac{k}{n} \times \mathbb{E}[Y] = \infty$. Therefore an assortatively efficient mechanism cannot yield strictly higher surplus.

For any distribution with finite expectation, all mechanisms have finite expected surplus. Any mechanism has a smaller expected surplus than giving *every* agent a free copy of the good, which would yield expected total surplus of $n\mathbb{E}[Y] < \infty$. Thus so long as $\mathbb{E}[Y]$ is finite, our strict welfare comparisons are meaningful. Our regularity condition therefore states that all results hold so long as no potential mechanism yields infinite surplus.

Some examples of distributions with strictly decreasing hazard rate and finite expectation include the Pareto distribution with scale parameter $\alpha > 1$ and the Weibull distribution with scale parameter $\lambda > 0$ and shape parameter $\beta \in (0, 1)$.

APPENDIX C. EXAMPLES OF INFORMATION POLICIES

Definition 8 (Trivial). Information policy p is trivial if

$$\sup \{p(t, h)\} \subset \Delta(\mathcal{M} \times \{0\}).$$

This means that messages can only arrive when $t = 0$. In this case, each agent chooses a single bid b to maximize her expected profit, with no chance to revise her bid in the future.

One function to obtain the information policy p in Example 2 is given by

$$p(t, h) = \begin{cases} (|h^Q|, \tau) & \text{if } t > \tau \\ (0, 0) & \text{otherwise.} \end{cases}$$

To see why, suppose the clock has not yet reached τ and there are $|h^Q|$ agents in the queue. If no agent joins the queue before τ , the information policy will announce that the queue has length $|h^Q|$ at time τ . If an agent does join before τ , the policy updates the message accordingly. Also note that while h^Q is not an explicit input of p , it can be recovered from h .

One function to obtain the information policy p in Example 3 is to announce whenever someone enters the queue. Then

$$p(t, h) = \begin{cases} (|h^Q|, t) & \text{if } h_{-1}^M \neq |h^Q| \\ (0, 0) & \text{otherwise,} \end{cases}$$

where h_{-1}^M denotes the most recent message sent by the announcer.

Each time an agent enters the queue, a new message is drawn from p , and it will trigger a message to announce the current length of the queue. Afterward, p reverts to subsequent announcement time of $t = 0$, which corresponds to no announcement. This default state remains until another agent joins the queue. The policy ensures all agents know the current length of the queue at all times.

Example 5 (Fixed Time and State). Say p announces at time τ if and only if the queue contains exactly k agents. Then

$$p(t, h) = \begin{cases} (|h^Q|, t) & \text{if } |h^Q| = k \text{ and } t \geq \tau \\ (0, 0) & \text{otherwise.} \end{cases}$$

The key intuition here is that the absence of an announcement can also be informative, and may change how agents behave after the time at which an announcement could have come.

Proof of Lemma 1. Consider a mechanism with uncertainty about the number of potential agents and number of items to be allocated. Let $\mathcal{N} = [\underline{n}, \bar{n}]$ denote the set of potential agents and let $\mathcal{K} = [\underline{k}, \bar{k}]$ represent the set of potential number of goods available, with $\underline{k} > 0$ and $\bar{k} < \underline{n}$ so that demand always exceeds supply. We allow for arbitrary correlation, so the probability that there are n agents and k items is given by $(n, k) \sim H \in \Delta(\mathcal{N} \times \mathcal{K})$ with probability mass function $h(n, k)$.²³ Agents have values drawn independently and identically from atomless F . All participating agents share common beliefs about the number of competitors and items for sale.

By assumption, both mechanisms have the same probability $\tilde{G}(v; n, k)$ of allocating an item to an agent of type v when there are n agents and k items. The overall probability that an agent reporting type z wins an item is therefore

$$G(z) := \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} h(n, k) \tilde{G}(z; n, k).$$

The rest of the proof is a straightforward application of the Revenue Equivalence Theorem; the expected payment of type v is then given by

$$p(v_i) = p(0) + \int_0^{v_i} v G'(v) dv$$

which depends only on the allocation rule G and does not depend on the mechanism format. ■

Proof of Proposition 1. As in Lemma 1, let $\mathcal{N} = [\underline{n}, \bar{n}]$ denote the set of potential agents, let $\mathcal{K} = [\underline{k}, \bar{k}]$ represent the set of potential number of goods available and assume $(n, k) \sim H \in \Delta(\mathcal{N} \times \mathcal{K})$ with probability mass function $h(n, k)$. Agents have values drawn independently and identically from atomless F , and let $Y_k^{(n)} \sim F_k^{(n)}$ represent the random variable that is the k^{th} highest of n draws from F . All participating agents share common beliefs about the number of competitors and items for sale.

First consider a unit-demand uniform-price auction. For any n and k with $n > k$, if an agent faces n other bidders and there are k items, it is an equilibrium to bid truthfully.²⁴ Consequently, even when there is uncertainty over agents and items, it is still optimal to bid truthfully. An

²³In the middle of a queueing game, the number of items remaining uniquely determines the number of competitors remaining. We use this lemma on “paused” moments in the middle of the game, and thus need to allow for correlation.

²⁴See e.g. Krishna 2009, chapter 13.

agent with value v wins if $v > Y_k^{(n-1)}$, which occurs with probability $F_k^{(n-1)}(v)$. The overall probability she wins is therefore

$$G(v) := \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} h(n, k) F_k^{(n-1)}(v),$$

and the expected payment is

$$(3) \quad p^{UPA}(v) = \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} h(n, k) F_k^{(n-1)}(v) \mathbb{E}[Y_k^{(n-1)} | Y_k^{(n-1)} < v].$$

Next consider a discriminatory price auction. Let β be a symmetric, increasing equilibrium. For an agent with value v , the expected payment is

$$(4) \quad p^{DPA}(v) = G(v) \beta(v).$$

By the Revenue Equivalence Theorem (Lemma 1), the payments in Equations (3) and (4) are equal.

$$(5) \quad \begin{aligned} G(v) \beta(v) &= \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} h(n, k) F_k^{(n-1)}(v) \mathbb{E}[Y_k^{(n-1)} | Y_k^{(n-1)} < v] \\ \beta(v) &= \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} \frac{h(n, k) F_k^{(n-1)}(v)}{G(v)} \times \mathbb{E}[Y_k^{(n-1)} | Y_k^{(n-1)} < v]. \end{aligned}$$

We conclude that the bidding function in a discriminatory price auction with supply and demand uncertainty is the average of the bidding functions with known supply and demand, weighted by the probability of the state times the probability the agent wins in the state. \blacksquare

Proof of Lemma 2. Let $\mathcal{N} = [\underline{n}, \bar{n}]$ denote the set of potential agents and let $\mathcal{K} = [\underline{k}, \bar{k}]$ represent the set of potential number of goods available, both with support on the positive integers and $\bar{k} < \underline{n}$. We allow for arbitrary correlation, so the common beliefs about the probability that there are n agents and k items is given by $(n, k) \sim H \in \Delta(\mathcal{N} \times \mathcal{K})$ with probability mass function $h(n, k)$. Any belief μ determines one such distribution h . Agents have values drawn independently and identically from $\bar{F} = F|_{[0, v]}$, denoted by random variable \bar{Y} . By Lemma 1, for any assortatively efficient mechanism a type- v agent has expected transfer given by Equation (3):

$$t_{AE}(v | \mu, v) = \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} h(n, k) \bar{F}_k^{(n-1)}(v) \mathbb{E}[\bar{Y}_k^{(n-1)} | \bar{Y}_k^{(n-1)} < v].$$

Since v is the highest type, they win with probability 1, and the above expression simplifies to

$$t_{AE}(v|\mu, v) = \sum_{n=\underline{n}}^{\bar{n}} \sum_{k=\underline{k}}^{\bar{k}} h(n, k) \mathbb{E}[\bar{Y}_k^{(n-1)}].$$

It is thus sufficient to show $\mathbb{E}[\bar{Y}_k^{(n-1)}]$ is continuous and strictly increasing in v . Since \bar{F} is a truncated version of F , as the upper support increases, so does the expectation of each draw. For continuity, note that

$$\begin{aligned} \mathbb{E}[\bar{Y}_k^{(n-1)}] &= \int_0^v x \frac{\partial}{\partial x} [\bar{F}_k^{(n-1)}(x)] dx \\ &= v \bar{F}_k^{(n-1)}(v) - \int_0^v \bar{F}_k^{(n-1)}(x) dx \\ &= v - \int_0^v \left[\sum_{j=0}^{k-1} \binom{n-1}{j} \left(\frac{F(x)}{F(v)} \right)^{n-1-j} \left(1 - \left(\frac{F(x)}{F(v)} \right) \right)^j \right] dx. \end{aligned}$$

The second line integrates by parts, and the third line uses the fact that $\bar{F}_k^{(n-1)}(v) = 1$ and expands the term under the integral. Since F is atomless, it is continuous, hence the integrand is continuous, as is the whole expression.²⁵ ■

Proof of Lemma 3. If k items remain, there must be $N - K + k$ remaining agents. Because v is the highest remaining type, agents have values drawn independently and identically from $F|_{[0,v]}$. Consider a uniform-price auction: the goods are allocated to the k bidders with the highest bids at the price of the $(k+1)^{\text{st}}$ -highest bid. It is a weakly dominant strategy to bid truthfully, in which case the auction is assortatively efficient. An agent with value v has an expected payment conditional on winning of $\mathbb{E}[Y_k^{(N-K+k-1)} | Y_k^{(N-K+k-1)} < v] = \mathbb{E}[Y_k^{(N-K+k-1)}]$, where Y is a random variable drawn from $F|_{[0,v]}$. Averaging the expected payment conditional on winning across all states,

$$\begin{aligned} t_{AE}(v|\mu, v) &= \frac{\sum_{k=1}^K \mu(k) \cdot \Pr(Y_k^{(N-K+k-1)} < v) \cdot \mathbb{E}[Y_k^{(N-K+k-1)} | Y_k^{(N-K+k-1)} < v]}{\sum_{k=1}^K \mu(k) \cdot \Pr(Y_k^{(N-K+k-1)} < v)} \\ &= \sum_{k=1}^K \mu(k) \cdot \mathbb{E}[Y_k^{(N-K+k-1)}]. \end{aligned}$$

²⁵As $v \rightarrow 0$, the support of \bar{F} collapses to 0, and the payment is also 0.

Observe that $\mathbb{E}[Y_k^{(N-K+k-1)}]$ is decreasing in k .²⁶ Thus

$$\begin{aligned} t_{AE}(v|\mu, v) &= \mathbb{E}_{k \sim \mu}[\mathbb{E}[Y_k^{(N-K+k-1)}]] \\ &< \mathbb{E}_{k \sim \mu'}[\mathbb{E}[Y_k^{(N-K+k-1)}]] \\ &= t_{AE}(v|\mu', v), \end{aligned}$$

where the inequality follows from strict FOSD dominance. ■

Proof of Lemma 4. By the Revelation Principle, the assortatively efficient equilibrium of $\mathcal{Q}(p)$ can be implemented by a direct revelation mechanism with assortatively efficient allocation. Hence by Lemma 1, when type v is in the game with beliefs μ and highest type \hat{v} , her expected transfer (conditional on winning) is $t_{AE}(v|\mu, \hat{v})$.

Forward Implication: Suppose type v attempts to join the queue at some $\tilde{t} \in \{t^+, t^-, t^{--}, \dots\}$. By assortative efficiency, v must be the highest type still remaining in the game (otherwise, all types higher than v remaining in the game would also attempt to join at time t , yielding assortative inefficiency). Thus type v 's expected payment conditional on winning is $t_{AE}(v|\mu_{\tilde{t}}, v)$. But if he attempts to enter at time \tilde{t} , his payment conditional on winning is t , hence $t = t_{AE}(v|\mu_{\tilde{t}}, v)$.

Reverse Implication: Suppose $t_{AE}(v|\mu_{\tilde{t}}, v) = t$ for some $\tilde{t} \in \{t^+, t^-, t^{--}, \dots\}$. Then type v must be the highest type remaining in the game: otherwise, the highest type remaining in the game \hat{v} would have expected payment (conditional on winning) of $t_{AE}(\hat{v}|\mu_{\tilde{t}}, \hat{v}) > t$, which is a contradiction because type \hat{v} can no longer pay more than t in the queueing game. Therefore, by Lemma 1, type v has expected payment (conditional on winning) in the queueing game of $t_{AE}(v|\mu_{\tilde{t}}, v) = t$. If type v waits until any time $t - \delta < t$ before attempting to join the queue, their expected payment conditional on winning will be strictly lower than t , contradicting the prior sentence. We conclude that type v must attempt to join the queue at time t . ■

Proof of Proposition 2. Because there are no announcements, the game is equivalent to a discriminatory price auction in which each agent may choose to pay cost c to enter the auction and submit a sealed bid (his entry time). By Topkis's Theorem, there exists some threshold type v^R such that agents enter if and only if their value is above v^R . By continuity, a type- v^R agent must be indifferent between entering or not entering, and therefore obtains zero expected utility

²⁶Compare the expression for k and $k + 1$. For any realization of $N - K + k - 1$ value draws, consider adding one more value draw. If the new draw exceeds $Y_k^{(N-K+k-1)}$, then $Y_{k+1}^{(N-K+(k+1)-1)} = Y_k^{(N-K+k-1)}$. Otherwise, $Y_{k+1}^{(N-K+(k+1)-1)} < Y_k^{(N-K+k-1)}$. Hence the expected value is decreasing in k .

if he enters. If $v^R \leq c$, then he obtains a strictly negative expected utility: he always pays c to enter, and with some probability less than 1 he receives an item with value at most c . Hence $v^R > c$. ■

Proof of Proposition 3. One equilibrium is the following. Agents with values less than c do not participate. All other agents write down a bid $\beta(v)$ given by their bidding function in a sealed-bid discriminatory price auction with reserve price c . In the auction equilibrium, the bidding function is given by

$$\beta(v) = \mathbb{E}[\max(c, Y_k^{(n-1)}) | \max(c, Y_k^{(n-1)}) < v].$$

Participating agents join the queue at time $\beta(v) - c$, if the announcer has not yet announced that the queue is full. Otherwise, they do not attempt to join the queue.

Agents' expected allocations and payoffs are equivalent to their expected allocations and payoffs in the equilibrium of the sealed-bid discriminatory price auction with reserve price of c . Hence agents have no profitable deviations. Furthermore, the equilibrium bidding functions are monotonically increasing for agents with values above c , hence the equilibrium is assortatively efficient.

We next show that this equilibrium is unique. In any equilibrium of the game, multiple types cannot pool by bidding the same value, because any such type could profitably deviate to increase their bid by some $\varepsilon > 0$. Furthermore, agents cannot play mixed strategies in equilibrium. Suppose some agent mixed between bids b and $b' > b$. By Topkis's theorem, no other type can bid in the interval $[b, b']$. But then the agent should strictly prefer bidding b to bidding b' . We conclude that agents' strategies are pure and fully separating. By Topkis's theorem, agents' bidding strategies are therefore strictly increasing in their values, and the allocation is assortatively efficient. By the Revenue Equivalence Theorem, any equilibrium bidding strategy must therefore correspond to β . ■

Proof of Proposition 4. Suppose strategy profile S results in an assortatively efficient equilibrium. We claim there would also be an assortatively efficient equilibrium of the following standard queueing game, which contradicts Theorem 1. Agents in the entry cost game have values drawn

according to F . Define a standard queueing game with value distribution G given by

$$G(v) = \begin{cases} 0 & \text{if } v < 0 \\ F(v + c) & \text{if } v \geq 0. \end{cases}$$

Intuitively, we have shifted down the values in the entry cost game by c to form a standard game in which a discrete mass of agents have value 0, and the rest have their old value reduced by c . Call this the “shifted standard game.”

Let S' be a strategy profile of the shifted standard game defined by an agent with value v using the strategy of an agent with value $v + c$ in strategy profile S of the entry-cost game, and only entering if the queue is not yet full. We claim that S' and the associated beliefs are an assortatively efficient equilibrium of the shifted standard game.

First, because they were assortatively efficient in the entry cost game, the resulting allocation is assortatively efficient in the shifted standard game. Because the values have been shifted down by c , all agent payoffs are equivalent in the two games as well. Furthermore, this standard queueing game also must release sudden bad news with positive probability, as entry behavior and announcements are identical.

We next show that S' is an equilibrium of the shifted standard queueing game. We claim that if an agent has a profitable deviation in the standard game with strategy profiles S , they have a profitable deviation in the entry cost game. Consider enacting this same profitable deviation in the entry-cost game, with the only difference that the agent attempts to join the queue only if the announcer has not yet announced that the queue is full. Because the conjectured equilibria were assortatively efficient, if the deviating agent wins the good in the standard game when deviating, they also win the good in the entry-cost game when deviating. Hence their deviation payoffs are equal in both the entry-cost game and standard game. Because their conjectured equilibrium payoffs were equal in both games as well, this deviation is profitable in the entry-cost game. Since by assumption there is no profitable deviation in the entry-cost game, we conclude there is no profitable deviation from S in the standard game we have constructed.

This implies that our strategy profile S is an assortatively efficient equilibrium of the standard queueing game with this information policy. This information policy also releases sudden bad news with some probability. If the distribution G did not have an atom at 0, this would immediately contradict Theorem 1, which holds for any atomless value distribution.

G does in fact have an atom at zero. Despite this, note that in our proofs the distribution being atomless is necessary to show only the (i) Revenue Equivalence Theorem and (ii) continuity of t_{AE} . Both of these results hold when there is an atom at 0, because the payoff of agents with value 0 must be 0 in this mechanism. Hence Theorem 1 holds for any value distribution that is atomless, except for a possible atom at zero. Thus our resulting assortatively efficient equilibrium contradicts Theorem 1. ■

Proof of Lemma 6. Because the mechanism is IR and makes no positive transfers, all types $v < 0$ have utility zero and are never allocated the good. Let $x(\theta)$ denote the ex-interim probability that type θ is allocated the good in equilibrium. By the Envelope Theorem,

$$u(v) = \int_{\underline{v}}^v x(\theta) d\theta = \int_0^v x(\theta) d\theta.$$

Again by the Envelope Theorem, the expected total surplus from the mechanism is

$$\begin{aligned} \mathbb{E}[U(x)] &= \int_{\underline{v}}^{\bar{v}} u(\theta) f(\theta) d\theta \\ &= \int_0^{\bar{v}} u(\theta) f(\theta) d\theta \\ &= u(\theta) F(\theta) \Big|_0^{\bar{v}} - \int_0^{\bar{v}} u'(\theta) F(\theta) d\theta \\ &= u(\bar{v}) - \int_0^{\bar{v}} x(\theta) F(\theta) d\theta \\ &= \int_0^{\bar{v}} x(\theta) d\theta - \int_0^{\bar{v}} x(\theta) F(\theta) d\theta \\ &= \int_0^{\bar{v}} \frac{1 - F(\theta)}{f(\theta)} x(\theta) f(\theta) d\theta. \end{aligned}$$

The final expression is maximized pointwise by allocating the good to the agents with the highest inverse hazard rate, $\frac{1-F(\theta)}{f(\theta)}$. If the inverse hazard rate is weakly (strictly) increasing, then the assortatively efficient allocation rule has weakly (strictly) higher expected total surplus than any other allocation rule. ■

REFERENCES

- Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv (2020). “Optimal Dynamic Matching”. In: *Theoretical Economics* 15.3, pp. 1221–1278.
- Bulow, Jeremy and Paul Klemperer (1994). “Rational Frenzies and Crashes”. In: *Journal of Political Economy* 102.1.

- Che, Yeon-Koo and Olivier Tercieux (2024). “Optimal Queue Design”. In: *Journal of Political Economy*, *Forthcoming*.
- Condorelli, Daniele (2012). “What Money Can’t Buy: Efficient Mechanism Design with Costly Signals”. In: *Games and Economic Behavior* 75.2, pp. 613–624.
- Dworczak, Piotr (2023). “Equity-Efficiency Trade-off in Quasi-Linear Environments”. In: *Working paper*.
- Gretschko, Vitali, Alexander Rasch, and Achim Wambach (2014). “On The Strictly Descending Multi-Unit Auction”. In: *Journal of Mathematical Economics* 50, pp. 79–85.
- Hartline, Jason D. and Tim Roughgarden (2008). “Optimal Mechanism Design and Money Burning”. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, pp. 75–84.
- Hassin, Refael and Moshe Haviv (2003). *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Boston, MA: Springer US. ISBN: 978-1-4615-0359-0. DOI: [10.1007/978-1-4615-0359-0_2](https://doi.org/10.1007/978-1-4615-0359-0_2). URL: https://doi.org/10.1007/978-1-4615-0359-0_2.
- Holt, Charles A. and Roger Sherman (1982). “Waiting-Line Auctions”. In: *Journal of Political Economy* 90.2, pp. 280–294.
- Krishna, Vijay (2009). *Auction Theory*. 2nd ed. 30 Corporate Drive, Suite 400, Burlington, MA 01803, USA: Academic Press.
- Leshno, Jacob (2022). “Dynamic Matching in Overloaded Waiting Lists”. In: *American Economic Review* 112.12, pp. 3876–3910.
- Leswing, Kif (2017). “Apple says to line up early if you want an iPhone X”. In: *Business Insider*.
- Murra-Anton, Zeky and Neil Thakral (2024). “The Public Housing Allocation Problem”. In: *Working paper*.
- Naor, Pinhas (1969). “The Regulation of Queue Size by Levying Tolls”. In: *Econometrica* 37.1, pp. 15–24.
- Price, Rob (2017). “People have been queueing outside the Apple Store for days to buy the iPhone X”. In: *Business Insider*.
- Reny, Philip J. (1999). “On The Existence of Pure and Mixed Strategy Nash Equilibria in Continuous Games”. In: *Econometrica* 67.5, pp. 1029–1056.
- Taylor, Grant A., Kevin K.K. Tsui, and Lijing Zhu (2003). “Lottery or waiting-line auction?” In: *Journal of Public Economics* 87.5, pp. 1313–1334. ISSN: 0047-2727.