Making Carbon-Allowance Auctions Robust to Aftermarkets

MOSHE BABAIOFF, Microsoft Research, Israel
NICOLE IMMORLICA, Microsoft Research, USA
YINGKAI LI, Northwestern University, USA
BRENDAN LUCIER, Microsoft Research, USA

In many regions, pollution is controlled by requiring firms to purchase emission allowances to cover their carbon emissions. These allowances are commonly allocated by the government using uniform-price auctions, and firms can typically trade these allowances among themselves in an aftermarket that may not be fully under the auctioneer's control. While the uniform-price auction is approximately efficient in isolation, we show that speculation and resale in aftermarkets might result in a significant welfare loss. Motivated by this issue, we consider three approaches, each ensuring high equilibrium welfare in the combined market. The first approach is to adopt smooth auctions such as discriminatory auctions. This approach is robust to correlated valuations and to participants acquiring information about others' types. However, discriminatory auctions have several downsides, notably that of charging bidders different prices per allowance, resulting in fairness concerns that make the format unpopular. Two other approaches we suggest are either using posted-pricing mechanisms, or using uniform-price auctions with anonymous reserves. We show that when using balanced prices, both these approaches ensure high equilibrium welfare in the combined market. The latter also inherits many of the benefits from uniform-price auctions such as price discovery, and can be introduced with a minor modification to auctions currently in use.

Additional Key Words and Phrases: carbon markets, aftermarkets, price of anarchy, multi-unit auctions, posted prices

1 INTRODUCTION

Our planet is warming at alarming rates. According to the Annual 2020 Global Climate Report of the National Oceanic and Atmospheric Administration [NOAA, 2020],[1] every month that year except December was in the top four warmest on record for that month. The Swiss Re Institute, the research arm of the reinsurance company Swiss Re based in Zurich, Switzerland, estimated in an April 2021 report [Swiss Re Institute, 2021], on the impact of climate change that "The world stands to lose close to 10% of total economic value by mid-century if climate change stays on the currently-anticipated trajectory." Due to these dire circumstances, the United Nations made combating climate change and its impacts one of 17 sustainable development goals [UN Goals, 2019]. Greenhouse gas emissions caused by human activity are a primary contributor to climate change according to the United States Environmental Protection Agency [EPA, 2021]. Reducing human consumption of these gasses requires heroic efforts on the part of humanity.

Current efforts are largely focused on imposing limits on emissions. This makes the "right to emit" a scarce resource. Emission allowances are allocated via markets, with the goal of distributing them to the industries that can provide the highest value to society subject to the emissions they generate. Each year, each company must present an emission allowance to account for all its pollution, or face steep fines. While some allowances are reserved for regulated industries such as public utilities and aviation, a large fraction of these allowances in major markets, such as the EU and California, are distributed via auctions on a monthly or quarterly basis. The EU Emissions Trading System (EU ETS), for instance, allocated 57% of allowances via auction between 2013 and 2020.[2] Each of these auctions is typically run using a single-round sealed-bid uniform-price auction. As described on the European Energy Exchange website [EEX, 2021], which runs the EU ETS, the auction format is:

- Single round: Bids will be submitted during one given bidding window.
- Sealed bid: Bids will be submitted without seeing other participant's bids.
- Uniform price: All successful bidders will pay the same auction clearing price.

This format was adopted to reflect the priorities of the EU commission which requires, according to Article 10(4) of DIRECTIVE 2003/87/EC, that auctions are designed to ensure transparency, equitable informational and procedural access, and that "participants do not undermine the operation of the auctions."[3]

In regard to the last priority, one clear point of concern for these auctions is the existence of unregulated (or only partially regulated) secondary markets. These secondary markets are not fully controlled by the primary auctioning bodies and can take many forms – bilateral trade, brokered trade, and exchanges, to name a few. There are certainly potential benefits to secondary markets, such as providing simpler market access to smaller firms who do not feel confident participating in the primary auction. Nonetheless, the existence of these secondary markets is concerning because they can distort outcomes, creating significant inefficiency even when the underline auctions are (approximately) efficient in isolation. But these secondary markets are also unavoidable absent extreme regulation. The EU ETS goes so far as to even sanction an EEX platform for secondary spot, futures and options contracts for emission allowances.

What is the impact of these secondary markets on the primary market? According to [EU ETS, 2021a], the auction clearing prices closely track the mean of the best-ask and best-bid prices on the EEX spot secondary market. Furthermore, both prices rise steadily month-over-month. These facts, taken together, suggest there is room for agents to speculate by buying allowances in the

[1]The NAOA is an American scientific and regulatory agency within the United States Department of Commerce tasked with understanding changes in climate, among other related topics.
[2]See the article from EU ETS [2021b].
[3]See the legal document from [Europa Union Law, 2021].

auctions and reselling them at a later time in the secondary markets. While allowing participants to buy and resell has some potential advantages (such as providing insurance against future price fluctuations), it also opens the door for speculation that could lead to excessive volatility, price bubbles, or stockpiling of quantities large enough to influence prices.[4] As noted in prior work of Quemin and Pahle [2021], "regulators are currently ill-equipped to appraise the beneficial and detrimental facets of speculation, and proper warning systems are wanting." Their work provides a diagnostic toolkit to assess the degree and impact of speculation in these markets. In our work, we ask whether the design of the primary market itself can defend against detrimental speculative behavior. Namely, *are the welfare guarantees of certain auction formats (approximately) robust to reallocation by arbitrary secondary markets? Can small changes to currently-used auctions achieve such robustness?*

To address these questions we focus on an idealized model of this market. We model the market as multi-unit auction in which each item corresponds to an allowance for one unit of emission. Agents (buyers/bidders) have decreasing marginal value for units. Items (allowances) acquired in the auction can be resold in a secondary market (aftermarket). To distinguish a "secondary market" (or aftermarket) from a general mechanism, we impose some mild conditions on the form these markets can take. Specifically, we assume that these are *trade mechanisms*: mechanisms that are budget balanced[5] and do not force participation.

As agents anticipate the secondary market, the potential for resale can change behavior in the primary auction and may encourage socially-wasteful speculation. As a result, the final allocation and welfare might be very different than that of the auction in isolation. One subtlety is that behavior in the secondary market can depend on the information released after the primary auction, such as whether bids are publicly observed. We want results that are robust to this choice, so we allow an arbitrary revelation of signals correlated with the auction bids and outcomes before the secondary market begins. Finally, we assume that agents are fully aware of the secondary market (and what information they'll learn about the primary auction outcome) when participating in the primary auction. We call the resulting mechanism that combines the primary auction and the aftermarket trade mechanism the *combined market*. We seek conditions on the design of the primary market that guarantee high welfare in every equilibrium allocation in the combined market.

One naive approach would be to choose primary market designs, such as the uniform-price auctions run by the EU ETS, that are approximately efficient in isolation. Indeed, if participants of the primary market don't anticipate the secondary market, then any trade in the secondary market only Pareto improves the utilities of all agents (as the trade mechanism satisfies voluntary participation and budget balance), and so welfare can never be harmed. The problem is that the existence of the secondary market impacts the strategies of agents in the primary market by, for example, enabling speculation. As we show, even a uniform-price auction can suffer arbitrarily large loss of welfare due to the presence of the secondary market. For example, if an individual participant has sufficient market power in the secondary market (i.e., by holding many allowances), they can distort prices to increase their own revenue at the expense of efficiency. In addition to the direct effect on welfare, this behavior has implications on the auction as well: a speculator may rationally behave as though they have complementary preferences over allowances, since they can exert market power in the secondary market only if they win a large number of allowances at auction. Unfortunately, the welfare guarantees of uniform-price auctions break down in the presence of complementarities. Moreover, since bidders have a tendency to engage in demand reduction in

[4]Bloomberg [2021] points out the urgent need to protect carbon market from excessive speculation.

[5]Our results will hold not only for trade mechanisms that are strongly budget balanced (net payment of 0), but also for weakly budget balanced mechanisms (mechanisms that never lose money).

order to keep auction prices low, this type of aggressive resale strategy will not necessarily be resisted by the competing agents. In Section 3 we present an example illustrating all of these effects under sequentially rational equilibrium behavior.

We argue however that some market formats *do* have robust welfare guarantees in combined markets. In particular, we show that *smooth* auctions maintain their welfare guarantees in such environments. Smoothness is a technical condition introduced by Roughgarden [2009] in the context of proving worst-case guarantees on welfare properties of equilibria. As noted by Syrgkanis and Tardos [2013], smoothness can be thought of as a sort of approximate First Welfare Theorem, whereby any loss in efficiency due to a reduced allocation to one bidder can always be partially offset by the payment of another bidder.[6]

Unfortunately, uniform-price auctions are not smooth, and indeed we show by example that their welfare can be severely reduced in the presence of a secondary market. Thus, current emission allowances markets are not robust to speculation opportunities created by aftermarkets. Alternative multi-unit auctions are, however, smooth. For example, discriminatory price auctions, in which each buyer pays her marginal bid for each unit she wins, are smooth. The smooth discriminatory-price auction therefore guarantees high equilibrium welfare in the combined market. We also show that this guarantee continues to hold even when agent valuations can be arbitrarily correlated, and even if participants can choose to acquire costly information about others' types in advance of the auction (such as a potential speculator investigating market forecasts). Yet, discriminatory price auctions also have some downsides. First, bidding is rather challenging and highly depends on distributional knowledge by the bidders. Second, discriminatory auctions might be perceived as unfair since identical goods are sold for different prices, creating envy between buyers. Finally, winners typically realize in retrospect that they could lower their payment by lowering their bids, creating regret. These issues could discourage participation in the auction, and indeed such concerns have been cited as reasons why this auction format was not adopted [EU ETS, 2022].

To circumvent this concern, we show that one can achieve robustness to secondary market distortions in another way. Instead of running an auction, one could use *posted prices*: make a quantity of allowances available at a declared price and allow buyers to purchase (in an arbitrary order) while supplies last. This combines the fixed-price feature of a carbon tax with the quantity restriction of an auction. The price offered would naturally depend on aggregate demand, so this method loses the price-discovery aspect of an auction. But if the designer is aware of overall market conditions (in the form of the expected maximal welfare of the market), then a recent line of work has demonstrated that simply posting an appropriate price achieves a welfare guarantee comparable to that of a uniform-price auction when running in isolation. Are such posted-price mechanisms robust to secondary markets?

In general, using an arbitrary posted-price mechanism is not enough to achieve the robustness we are looking for. We show by way of example that even if a posted-price mechanism achieves high welfare on its own, this welfare can be significantly decreased by speculation that occurs at equilibrium in the presence of a secondary market. This motivates us to focus on a particular form of posted-price mechanisms: those that use *balanced prices*, which are set proportional to the expected average welfare generated in the efficient allocation. It is known that balanced prices yield strong welfare guarantees for a variety of allocation problems, including multi-unit auctions [Dutting et al., 2020, Feldman et al., 2014], and we prove that the welfare guarantee from balanced prices continues to hold at any equilibrium even in the presence of a secondary market.[7] Specifically, for

[6]Informally, the First Welfare Theorem states that when prices clear the market, the allocation is socially efficient.
[7]Unlike with smooth auctions, this result does not extend to correlated value distributions or to settings where buyers can acquire additional information before the auction opens. For example, it is problematic if the buyers become more informed than the designer who set the prices.

our setting of selling identical allowances to buyers with decreasing marginal values, there is a per-item price that guarantees at least half of the expected optimal welfare.

To this point we have described two methods for allocating allowances that achieve robust welfare guarantees in the presence of secondary markets. One is to use a smooth mechanism, such as a discriminatory auction, and the other is to find a balanced price and sell allowances at that price while supplies last. We have already discussed potential drawbacks of the former solution, but the latter has its own set of practical challenges: for one thing, it is a dramatic change relative to the uniform-price auctions typically used; for another, if the government underestimates demand and sets its price too low, this could encourage a rush where buyers race to purchase allowances at the moment they become available, resulting in buyer frustration and a perception of unfairness. As it turns out, one can address these issues and obtain all of the benefits of balanced posted prices with a small tweak to a uniform-price auction – the introduction of appropriately-chosen reserves. Reserve prices are already common in many emission auctions, such as the one administered by the California Air Resources Board [CARB, 2022], in the form of price floors. We show an appropriate choice of reserve guards against welfare loss in combined markets: namely one can augment a uniform-price auction with a per-allowance reserve price set equal to the balanced per-allowance price one would use in a posted-price mechanism. We prove that the expected welfare at any equilibrium of this auction is at least half of the expected optimal welfare, and this guarantee persists in the presence of an arbitrary secondary market. We view this as a practical solution that can be implemented with minimal effort: as long as the government can estimate the expected average social value of a carbon allowance, they can mitigate the impact of speculation and other equilibrium effects of resale by employing an appropriately-determined auction reserve. We further show that the welfare guarantees degrade gracefully as one adjusts the prices, meaning that unavoidable misspecifications in the price determination will have a modest effect on the welfare guarantees.

1.1 Related Work

There is a rich literature exploring market and regulation-based techniques for reducing emissions and their effectiveness. Here we discuss a small sampling of this literature, referring the reader to many excellent overviews such as [Cotton, 2015] or [Cramton et al., 2017] for further details. Weitzman [1974] asks whether it is better to control emissions via imposing standards (quantity regulation) or charging taxes (price regulation), and notes that prices tend to fare better when the social cost of emissions is close to linear, whereas quantity regulation can be preferable in the face of uncertainty when marginal costs are variable. We note that taxes share many similarities with the posted-pricing mechanism we study. Cramton and Kerr [2002], in turn, propose selling emission allowances in an auction (they suggest an ascending auction). Their paper explicitly suggests these allowances be tradeable in aftermarkets to maximize liquidity. Goldner et al. [2020] study uniform-price auctions with price floors and ceilings, a common mechanism in practice, and prove welfare guarantees under certain conditions in the absence of aftermarkets. Our paper complements these by exploring the interplay of these aftermarkets and the primary auction and stating conditions under which welfare guarantees extend to the combined market.

One closely related line of theoretical work is price of anarchy for sequential auctions [Paes Leme et al., 2012, Syrgkanis and Tardos, 2012]. In Paes Leme et al. [2012], the authors illustrate that although price of anarchy of the sequential composition of first-price auction is small for unit-demand agents, the result breaks for agents with submodular valuations, and the price of anarchy can be unbounded in the latter case. In contrast, our results indicates that for submodular valuations, any trade mechanism[8] followed by simultaneous first-price auction will have constant price of anarchy for the combined market. The main difference that allows us to handle combinatorial auctions in sequential auction format is that all items are sold only in the first market, and the secondary market is only providing the platform for agents to retrade the items, rather than selling items sequentially, with each item sold once in one of the auctions.

Recently Eden et al. [2020] bound the price of anarchy when each agent is subject to an externality from the allocation of the other agents. The authors motive the externality by the resale model since the value of the an agent for winning any item depends on the utility gain of potentially reselling the item to other agents, and the latter depends on other agents' private assessment of the item. However, they assume that those resale behaviors are fixed and hence the externality among agents in the auction are exogenous. In contrast, we assume that the secondary market is outside the control of the seller, and the utilities of agents for winning any item in the auction are endogenously determined by the mechanism adopted in the auction and the corresponding equilibrium behavior of all agents in the combined market.

The challenges in analyzing the equilibrium in the combined market was acknowledged in Haile [2003] due to the fact that there exist endogenously induced common value components in the auction. In the simple single-item setting with winner posting prices as secondary markets,[9] Hafalir and Krishna [2008] characterized the equilibrium behavior of the agents in the combined market, and Hafalir and Krishna [2009] adopted the characterization to show that the expected welfare of the first-price auction with secondary markets may decrease by a multiplicative factor of ${}^{2e}/_{(2e-1)}$. In addition to the above discussions, there are many papers discussing various properties of the resale model in the economics literature, including but not limited to the observation of bid shading in the auction [Pagnozzi, 2007], and the revenue ranking of the simple auctions [Lebrun, 2010]. See the survey of Susin [2017] for more discussions on the equilibrium properties of the resale model. Finally, there are several recent papers focusing on designing optimal mechanisms when the seller has no control over the secondary market. Carroll and Segal [2019] show that second price auction with reserve prices is the robustly revenue optimal mechanisms with unknown resale opportunities. Dworczak [2020] considers the design of information released to the secondary markets and show that the information structure that induces truthful behaviors are cutoff rules. He also provides sufficient conditions for simple information structure to be optimal.

Our techniques leverage smoothness and balanced pricing. The smoothness framework is a powerful tool for analyzing the price of anarchy in auctions (see [Roughgarden et al., 2017] for a detailed discussion on the literature). This framework is first proposed in [Roughgarden, 2009] for complete information games and [Roughgarden, 2012, Syrgkanis, 2012] for incomplete information games. The idea of smoothness is further refined and generalized for simultaneous composition and sequential composition of smooth mechanisms in multi-item settings [e.g., Feldman et al., 2013, Syrgkanis and Tardos, 2013]. For multi-unit auctions in particular, this theory has been used to derive equilibrium welfare bounds for different auction formats in isolation [de Keijzer et al., 2013, Markakis and Telelis, 2015], and we extend this analysis to settings with aftermarkets. The balanced pricing framework is a general approach for designing posted-price mechanisms in a broad class of allocation problems [Dutting et al., 2020, Feldman et al., 2014, Kleinberg and Weinberg, 2012]. These constructions employ the theory of Prophet inequalities to bound the welfare obtained when buyers sequentially purchase their preferred bundles at the proposed prices, which are calculated using the distribution of buyer values. Similar to smoothness, we extend the existing analysis

[8]This means that the secondary market satisfies voluntary participation and weak budget balance. See Section 2 for a detailed definition of the assumption.

[9]In this model, the authors also assume that no information, especially the bids, are revealed in the secondary market to avoid the ratchet effect.

to show that the welfare guarantees attainable through balanced pricing extend to settings with aftermarkets.

2 PRELIMINARIES

2.1 The Basic Setting

For clarity we begin by describing a basic model focused on a multi-unit auction of identical allowances. Our results actually apply to a more general model with multiple allowance types and different information structures (including the ability to purchase information about aggregate demand); we describe these extensions in Section 2.2.

2.1.1 The Allocation Problem. We model a carbon allowances allocation problem in which a government agency holds a set M of m identical licenses (allowances) to be allocated among a set $\mathcal{N}$ of n buyers. A feasible allocation of the items is a profile $\mathbf{x} = (x_1, \ldots, x_n)$, where $x_i \in [m]$ is the number of licenses obtained by agent i and $\sum_i x_i \leq m$. We write $\mathcal{X}$ for the set of feasible allocations.

Buyer i has a private valuation function $v_i \colon [m] \to \mathbb{R}_{\geq 0}$ where $v_i(x_i)$ denotes buyer i's value for obtaining x_i licenses, normalized so that $v_i(0) = 0$. We emphasize that this is a consumption value, driven by the ability to emit pollution as afforded by the licenses. Valuations are assumed to have non-increasing marginal valuations: for each $j \geq 1$, $v_i(j) - v_i(j-1)$ is non-negative and weakly decreasing in j. We will sometimes refer to v_i as the type of agent i. We write $\Theta = \times_i \Theta_i$ for the set of valuation profiles. We assume that v_i is sampled independently from a known distribution F_i, and denote the prior product distribution over the valuations by $F = \times_i F_i$. The utility of agent i given allocation x_i and total payment p_i is $u_i(x_i, p_i) = v_i(x_i) - p_i$. Buyers are assumed to be risk-neutral and seek to maximize expected utility.

The *welfare* of an allocation $\mathbf{x} \in \mathcal{X}$ when the valuations are $\mathbf{v}$ is defined to be $\mathrm{Wel}(\mathbf{v}, \mathbf{x}) = \sum_i v_i(x_i)$. For any valuation profile $\mathbf{v}$, let $\mathrm{Wel}(\mathbf{v}, \mathcal{X}) = \sup_{x \in \mathcal{X}} \mathrm{Wel}(\mathbf{v}, x)$ be the optimal (highest) welfare given the valuation functions $\mathbf{v}$ and feasibility constraint $\mathcal{X}$. We say an allocation is *efficient* if it achieves the optimal welfare. Let $\mathrm{Wel}(F, \mathcal{X}) = \mathbf{E}_{\mathbf{v} \sim F}[\mathrm{Wel}(\mathbf{v}, \mathcal{X})]$ be the expected optimal welfare. When $\mathcal{X}$ is clear from the context, we omit it in the notation and use $\mathrm{Wel}(\mathbf{v}), \mathrm{Wel}(F)$ to denote the optimal welfare and expected optimal welfare, respectively.

2.1.2 Mechanisms. In our model, agents can acquire licenses by participating in an auction then trading licenses among themselves in a secondary market. We will formally describe both the auction and the secondary market as *mechanisms*. A mechanism $\mathcal{M}$ defines a set of actions for each agent, and a (possibly random) mapping from profile of actions to a feasible allocation and payment from each agent. Formally, a mechanism $\mathcal{M} = (x^{\mathcal{M}}, p^{\mathcal{M}}) : A \to \Delta(\mathcal{X} \times \mathbb{R}^n)$ is defined by an *allocation rule* $x^{\mathcal{M}} : A \to \Delta(\mathcal{X})$ and a *payment rule* $p^{\mathcal{M}} : A \to \mathbb{R}^n$, where $A = \times_i A_i$ and A_i is the action space of agent i in the mechanism. Thus, for action profile $\mathbf{a} = (a_1, a_2, \ldots, a_n) \in (A_1, A_2, \ldots, A_n) = A$ the *outcome* of the mechanism is the (randomized) allocation $x^{\mathcal{M}}(\mathbf{a})$, and each agent i is charged (in expectation) a *payment* of $p_i^{\mathcal{M}}(\mathbf{a}) \geq 0$. For example, in an auction each agent's action is typically a bid, and the allocation and payment rules describe how to use the submitted bids to determine how many licenses are won by each buyer and at what prices. The utility of agent i with valuation v_i when participating in the mechanism $\mathcal{M}$ in which agents take actions $\mathbf{a} \in A$ is $u_i(\mathcal{M}(\mathbf{a})) = v_i(x^{\mathcal{M}}(\mathbf{a})) - p_i^{\mathcal{M}}(\mathbf{a})$.

A mechanism $\mathcal{M}$ with valuation distribution F defines a game. A *strategy* $\sigma_i : v_i \to \Delta(a_i)$ for agent i is a mapping from her valuation v_i to a distribution over her actions. With slight abuse of notation denote by $\sigma_{-i}(v_{-i})$ the profile of actions taken by agents other than i when each $j \neq i$ has valuation v_j. A strategy σ_i is a *best response* for agent i given strategies of the others σ_{-i} if for any strategy σ_i' it holds that $\mathbf{E}[u_i(\mathcal{M}((\sigma_i(v_i), \sigma_{-i}(v_{-i}))))] \geq \mathbf{E}\big[u_i(\mathcal{M}((\sigma_i'(v_i), \sigma_{-i}(v_{-i}))))\big]$ for every valuation v_i, where the expectation is over the valuations of the other agents as well as any randomness in the mechanism and strategies. A profile of strategies $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a *Bayesian Nash equilibrium* (BNE) for mechanism $\mathcal{M}$ with distribution F, if for every agent i, strategy σ_i is a best response for agent i given strategies of the others σ_{-i}.

By slightly overloading the notation, we also denote $\mathrm{Wel}(\mathcal{M}, \sigma, F)$ as the expected welfare obtained in mechanism $\mathcal{M}$ using equilibrium strategy profile σ. Let the *price of anarchy of mechanism $\mathcal{M}$ within the family of distributions $\mathcal{F}$* be

$$\mathrm{PoA}(\mathcal{M}, \mathcal{F}) = \sup_{F \in \mathcal{F}} \frac{\mathrm{Wel}(F)}{\inf_{\sigma \in \mathrm{BNE}(F, \mathcal{M})} \{\mathrm{Wel}(\mathcal{M}, \sigma, F)\}}$$

where $\mathrm{BNE}(F, \mathcal{M})$ is the set of Bayesian Nash equilibria given distributions F and mechanism $\mathcal{M}$.

Auctions. We can describe multi-unit auctions as mechanisms. For example, in most common multi-unit auctions, an action of bidder i is a *bid* of the form $(a_{i1}, \ldots, a_{im}) \in A_i = \mathbb{R}^m_{\geq 0}$ representing her m marginal values, with $a_{i1} \geq \cdots \geq a_{im} \geq 0$. The agents simultaneously declare these bids to the auctioneer. The $n \times m$ received marginal bids are then sorted from largest to smallest, and the m identical licenses are greedily allocated to the bidders of the m highest marginal bids (breaking ties arbitrarily). The two most common auction formats, uniform and discriminatory, then differ in how payments are calculated:

- *discriminatory auction*: Each agent pays her winning bids. That is, if agent i wins x_i licenses, then she pays her highest x_i marginal bids: $p_i = \sum_{j=1}^{x_i} a_{ij}$.
- *uniform-price auction*: A common price p per unit is chosen, and each agent pays p for each license won. This price p is taken to lie between the (m)-th highest marginal bid and the $(m+1)$-st highest (within the $n \times m$ reported marginals). In other words, p lies between the highest losing bid and the lowest winning bid. In this paper we focus on the case where p is the highest losing bid.

Auctions also sometimes impose further restrictions on the bids that can be submitted. An auction with *standard bidding* imposes no restrictions aside from non-increasing marginal values. This is in contrast to *uniform bidding* which requires that each agent submit a bid in which all non-zero marginal values are equal to each other. In this paper we will focus on standard bidding, which is the format used by the EU ETS.

We will also be interested in a simple auction format known as a posted-price mechanism. In such a mechanism, the auctioneer selects in advance a price p per license. The buyers are then approached sequentially in a fixed order. Each buyer can choose to buy as many licenses as desired, up to the amount remaining, at a price of p per license. That is, if the buyers purchase in the order $1, 2, \ldots, n$, then each bidder i can choose any non-negative $x_i \leq m - \sum_{j<i} x_j$ and pays $p \cdot x_i$. Once all licenses are sold the mechanism ends.

Secondary Markets. Informally, a secondary market allows users to trade items that they obtained from the auction. For example, one might imagine that agents could offer to sell their licenses at a certain price, and other agents might choose to purchase at the suggested price (or not). Similar to the auction, we will model the secondary market as a mechanism. The starting point of the secondary market is the allocation picked by the auction, which is publicly revealed. The secondary market is therefore parameterized by an allocation $x \in \mathcal{X}$, which we think of as the auction outcome. We will tend to use $\mathcal{M}^2$ to refer to secondary market mechanisms, and in a slight abuse of notation we will write $\mathcal{M}^2(\mathbf{a};\mathbf{x}) = (x^{\mathcal{M}^2}(\mathbf{a};\mathbf{x}), p^{\mathcal{M}^2}(\mathbf{a};\mathbf{x}))$ for the allocation and payment rules of a secondary market $\mathcal{M}^2$ as a function of the initial allocation $\mathbf{x} \in \mathcal{X}$.[10]

To capture our intuitive notion of a secondary market, we will introduce two mechanism properties that we will assume in secondary markets we consider. First, we assume voluntary participation, which informally means that each agent can choose not to participate. More formally, voluntary participation requires that each agent has an "opt-out" action that guarantees their utility is not reduced by the secondary market.

Definition 1. *A secondary market $\mathcal{M}^2$ satisfies* voluntary participation *if for each agent i, all valuations v_i, and all feasible allocations $\mathbf{x}$, there exists an action a_i^* such that, for any action profile $\mathbf{a}_{-i}$ of the other agents, $v_i(x_i) \leq u_i(\mathcal{M}^2((a_i^*, \mathbf{a}_{-i}); \mathbf{x}))$.*

We argue that this condition is quite mild. For example, if the secondary market is one in which license holders can suggest take-it-or-leave-it prices, and trade happens if another user agrees to trade at that price, then a license holder might decide not to make an offer ("not participate"), and other agents can decide to decline any offer made (again, "not participate"). Each agent can therefore ensure that her utility in the secondary market is the same as the utility obtained from the initial allocation $\mathbf{x}$.

We also assume that our secondary market satisfies weak budget balance, which means that it does not run a deficit with respect to payments.

Definition 2. *A secondary market $\mathcal{M}^2$ satisfies* weak budget balance *if $\sum_i p_i^{\mathcal{M}^2}(\mathbf{a};\mathbf{x}) \geq 0$ for any action profile $\mathbf{a}$ and feasible allocation $\mathbf{x} \in \mathcal{X}$.*

A mechanism that satisfies both voluntary participation and weak budget balance is called a *voluntary-non-subsidized-trade mechanism*, or a *trade mechanism* for short.

2.1.3 A Combined Game. We are finally ready to formally model our setting of an auction followed by a secondary market. We model this scenario as a two-round game $\mathcal{G}$ that consists of two mechanisms, $\mathcal{M}^1$ and $\mathcal{M}^2$, run sequentially. We refer to $\mathcal{M}^1$ as the *auction* and $\mathcal{M}^2$ as the *secondary market.*

In the first round of the game, the agents participate in the auction $\mathcal{M}^1$. We denote the action space of $\mathcal{M}^1$ by $A^1 = \times_i A_i^1$. The agents simultaneously choose actions $a^1 \in A^1$, resulting in outcome $x^{\mathcal{M}^1}(a^1)$ and payments $(p_i^{\mathcal{M}^1}(a^1))_i$. Each agent observes the outcome of the auction and her own payment.

The second round then starts and the agents participate in the secondary market $\mathcal{M}^2$. The allocation $x^{\mathcal{M}^1}(a^1)$ from the auction is used as the initial allocation in the secondary market. We denote the action space of $\mathcal{M}^2$ by $A^2 = \times_i A_i^2$. Note that the action spaces of the two mechanisms can be different, but they share a common set of feasible allocations.

To summarize, the two-round game $\mathcal{G}(\mathcal{M}^1, \mathcal{M}^2)$ proceeds as follows:

(1) Each agent i picks an action $a_i^1 \in A_i^1$, all agents choose actions simultaneously. Mechanism $\mathcal{M}^1$ runs on actions a^1.
(2) Each agent i observes $x^{\mathcal{M}^1}(a^1)$ and $p_i^{\mathcal{M}^1}(a^1)$.
(3) Each agent i picks an action $a_i^2 \in A_i^2$, all agents choose actions simultaneously. Mechanism $\mathcal{M}^2$ runs on actions a^2, starting from allocation $x^{\mathcal{M}^1}(a^1)$.

[10] In some secondary market formats it is more natural to think of actions being chosen sequentially rather than simultaneously. E.g., in the example above, a seller first chooses a price then buyers choose whether to purchase. One can model this by having a buyer's "action" be a mapping from all possible observations (e.g., prices) to a realized action (e.g., whether to buy).

(4) The total payoff to agent i in the combined market is $u_i(\mathcal{M}^2(a^2; x^{\mathcal{M}^1}(a^1))) - p_i^{\mathcal{M}^1}(a^1)$.[11]

Note that an instance of the two-round game $\mathcal{G}(\mathcal{M}^1, \mathcal{M}^2)$ naturally corresponds to a combined mechanism $\mathcal{M}^C$ which we denote by $\mathcal{M}^C = \mathcal{G}(\mathcal{M}^1, \mathcal{M}^2)$, in which an action has two components: (1) an action $a_i^1 \in A_i^1$ for the auction mechanisms, and (2) a mapping for each agent i from the tuple of (allocation, payment) from the auction into an action for the secondary market. The notions of BNE and PoA then extend to such a mechanism $\mathcal{M}^C$ as before.

Since our combined market is an extensive-form game that proceeds in multiple rounds, one might wonder about natural refinements of the BNE solution concept such as perfect Bayesian equilibrium or the stronger refinement of sequential equilibrium. We formally define these equilibrium notions in Appendix B. Roughly speaking, such refinements require (a) subgame perfection, where behavior in the secondary market is always rational given any auction outcome, and (b) that agents accurately update their beliefs after the auction outcome and behave in accordance with those beliefs in the secondary market. As it turns out, our positive results about welfare hold at *any* BNE, whether or not they satisfy these requirements. So in particular our welfare bounds also hold for any sequential equilibria as well. Moreover, each example we use to illustrate a negative result will not only be a BNE, but rather also be a sequential equilibrium.

2.2 Extensions

All of the positive results in our paper actually apply in a more general framework than the basic setting described above. So far we have focused on the simpler setting for ease of notation and to more directly connect to the application of allocating emission licenses, but we will now describe two generalizations of the model. All of our positive results in the remainder of the paper will actually be proven for this generalized setting.[12]

Combinatorial Allocation and Multiple License Types. In the basic model, the licenses to be allocated are all identical, so each buyer is concerned only with the number of licenses she obtains. More generally, we can consider a *combinatorial auction* scenario where there is a set of (possibly different) goods to allocate and each buyer has a value for each possible combination of goods. We then interpret an allocation x_i to buyer i as a subset of the available goods, $x_i \subseteq M$. An allocation profile $\mathbf{x} = (x_1, \ldots, x_n)$ is then feasible if no item is double-allocated, meaning that $x_i \cap x_j = \emptyset$ for all $i \neq j$. The basic model is the special case where all of the items are identical and agents have non-increasing marginal values for the items. This generalization also captures scenarios where licenses of different types being sold alongside each other, such as licenses that apply to different calendar years or that permit different forms of emissions. The natural generalization of "non-increasing marginal values" is then that agent valuations are submodular, meaning that $v_i(S) + v_i(T) \geq v_i(S \cap T) + v_i(S \cup T)$ for all sets of items $S, T \subseteq M$.

Post-Auction Information Revealed. In the basic model, after the auction but before the secondary market, each agent observes the outcome of the auction and her own auction payment. More generally, each agent might also observe some additional information about the auction before the secondary market begins. For example, it may be that all agents' payments are revealed, or it might be that all bids are made public. Formally, we can think of each buyer i as observing a private signal $s_i \in S_i$ after the auction that can be correlated with A^1, $x^{\mathcal{M}^1}(a^1)$, and $(p_i^{\mathcal{M}^1}(a^1))_i$. In fact, we could also allow these signals to be correlated with the valuations of the agents, which

[11]Note that this expression includes the payments from both the auction and the secondary market, as the secondary market transfers are included in the utility term.
[12]In particular, our welfare bounds in Sections 4, 5, and 6 all apply in this generalized setting, assuming that agent valuations are submodular (which is a natural generalization of the assumption that agents have non-increasing marginal values).

allows the agents to receive additional information that even the auction has no direct access to. We can write Γ for the (possibly randomized) mapping from $\mathbf{v}$, a^1, $x^{\mathcal{M}^1}(a^1)$, and $(p_i^{\mathcal{M}^1}(a^1))_i$ to the profile of signals $(s_1, \ldots, s_n)$ that the agents receive after the auction. Under this generalization we would include Γ in the description of the combined mechanism, so that $\mathcal{M}^C = \mathcal{G}(\mathcal{M}^1, \Gamma, \mathcal{M}^2)$. In the basic setting the agents receive no signals, so we can think of Γ as being the empty mapping that always returns a null signal.

3 WELFARE LOSS AND SECONDARY MARKETS

Uniform-price auctions are approximately efficient in isolation: as long as bidders avoid dominated overbidding strategies, the price of anarchy is at most a constant [de Keijzer et al., 2013]. In this section we present an example showing that for such an auction, the presence of a secondary market can induce a sequential equilibrium with low expected welfare, even when bidders avoid weakly dominated strategies. Intuitively, this means that behaviour in the combined market does not hinge on non-credible threats regarding behaviour in the secondary market or on playing dominated strategies in the auction.

Example 1. *There are $m > 3$ units to be allocated and 3 agents named A, B and C. Agent A has marginal value 2 for the first unit, value uniformly sample from $[1, 1.5]$ for the second unit, and 0 for any subsequent units. Agent B has the following distribution over valuations. She always has marginal value 2 for the first unit acquired, then a value $z_B > 0$ for each subsequent unit acquired. Here z_B is a random variable drawn from a distribution with CDF $F(z) = 1 - \frac{1}{1+(2m-1)z}$ for $z \in [0, 1)$ and $F(z) = z - \frac{1}{2m}$ for $z \in [1, 1 + \frac{1}{2m}]$. Note that given buyer value with distribution F, the unique revenue maximizing price is $p = 1$ with expected revenue of $p \cdot \Pr[z_B \geq p] = 1/(2m)$ per-unit. Agent C has value 0 for any number of units; we refer to agent C as a* speculator.

The primary auction will be a uniform-price auction with standard bidding. The auction is resolved by allocating units in order of bids (breaking ties in any arbitrary but consistent manner) until all units have been allocated. The price p per unit is set equal to the highest unallocated bid. Each agent then pays p per unit received.

In the secondary market, the speculator C can put some or all of the licenses that she has acquired in the auction up for sale, at a take-it-or-leave-it price of her choice. Agent A has the first opportunity to purchase any (or all) of the licenses made available by C. Then agent B has the option to purchase any licenses that are still available.[13]

We now describe a particular choice of bidding strategies in the primary auction and agent behavior in the secondary market. We will then prove that these form an equilibrium in the combined market (and actually, a sequential equilibrium that avoids weakly dominated strategies). In the auction, agents A and B each bid 2 for exactly a single unit and 0 for the rest of the units. The speculator C bids 1 for $m-2$ units and 0 for the rest of the units. Then, in the secondary market, the speculator offers all the units she has acquired for the price of 1. Agent A buys one unit, and agent B buys all $m-3$ remaining units if $z_B = 1$, and nothing otherwise. Note that under this behaviour the price in the auction is 0. Agents A and B have utility of 2 in every realization of values. Agent C has expected utility of $1 + (m-3)/(2m)$, as she makes 1 from selling to A, and $(m-3)/(2m)$ in expectation from selling to B.

Consider the social welfare obtained in the combined market. The total expected welfare obtained under this behavior is at most 6, whereas the optimal expected welfare is at least $4 + (m-2)\mathbb{E}[z]$. It

[13] A more general secondary market can allow all agents to offer acquired licenses for sale, not just agent C. The result presented extends to such more general version as well, but for simplicity we focus on a secondary market where only agent C can sell.

is easy to compute that $\mathbb{E}[z] = \Theta((\log m)/m)$, and hence the optimal expected welfare is $\Theta(\log m)$. Thus, if this behavior occurs at equilibrium, this implies that the price of anarchy for this combined market is $\Theta(\log m)$, growing unboundedly large with m.

We now claim that this behavior does indeed form a sequential equilibrium. Moreover, the behavior of the agents in the auction is not "weakly dominated," in the sense that there is no "unambiguously better" alternative bidding strategy for any agent given the way that agents will behave in the secondary market given the auction outcome. The formal definition of sequential equilibrium and weakly dominated strategies, as well as the proof of Proposition 1, is provided in Appendix B.

Proposition 1. *The above behaviour in the combined market forms a sequential equilibrium. Moreover, no agent is using a weakly dominated strategy in the auction.*

Discussion. Let's interpret this example. One thing to notice is that, at equilibrium, the speculator is placing a very high bid for a very large number of licenses; much higher than the revenue she obtains in the secondary market. Of course, the speculator can afford these licenses because of the low price, but what is the point of placing such a high bid? One might expect that this sort of behavior is a form of "bullying" and it may be surprising at first glance that it is not excluded by removing dominated strategies (similar to overbidding in a second-price auction). But we note that, at equilibrium, the speculator behaves as a monopolist in the secondary market. She generates a modest amount of revenue from each license sold to agent B (namely, $1/(2m)$ each), plus a large amount of revenue (revenue of 1) by selling a license to agent A. However, this sale to agent A can only occur if agent A obtains fewer than 2 licenses at auction. This creates an extra incentive for the speculator C to obtain many licenses, to prevent agent A from obtaining a second one. This can cause the licenses to appear complementary to speculator C, depending on the bidding behavior of agent A: obtaining one fewer licenses could dramatically reduce C's utility if that one license is won by agent A instead. This rationalizes the overbidding that occurs at equilibrium in the primary auction, where the speculator makes an effective bid much higher than her obtained revenue in the secondary market. It is for this reason that the overbidding behavior of the speculator is not weakly dominated.

Another implication of speculator C's monopolistic behavior is that she posts a high price that distorts the allocation to agent B. Although agent B has high expected welfare for the goods that agent C holds, the speculator C maximizes revenue by setting the probability of trade very low and significantly reducing welfare.

Finally, we note that bidders A and B are systematically under-bidding in this equilibrium. This is driven by demand reduction effects, where the bidders are strictly incentivized to underbid in order to keep prices low. Importantly, this behavior is not driven by indifference, and is undominated.

4 PRICE OF ANARCHY VIA SMOOTH FRAMEWORK

In the previous section we saw that the expected welfare of a uniform-price auction may decrease catastrophically when there is a secondary market, even in "natural" equilibria that are sequentially rational and avoid weakly dominated strategies. Can the welfare loss due to aftermarkets be bounded for other auction formats? Unfortunately, explicitly characterizing the associated welfare loss is a laborious task: it requires one to construct and analyze the equilibria in the combined market, which can depend on the agents' distributions in subtle ways.

In this section, we circumvent the challenges of explicitly characterizing all equilibrium strategies by showing that while adding a secondary market might harm welfare, the worst-case welfare guarantees of several classical mechanisms (including discriminatory auctions) will not decrease in the combined market, as long as the auction mechanism satisfies certain smoothness properties. In other words, while the equilibrium welfare may decrease in particular market instances, worst-case guarantees are retained for smooth mechanisms. The following definition captures the notion of smoothness we require.

Definition 3 (Syrgkanis and Tardos, 2013). *Auction $\mathcal{M}$ with action space A is (λ, μ)-*smooth *for $\lambda > 0$ and $\mu \geq 1$, if for any valuation profile $\mathbf{v}$, there exists action distributions $\{D_i(\mathbf{v})\}_{i\in[n]}$ such that for any action profile $a \in A$,*

$$\sum_{i\in[n]} \mathbf{E}_{a'_i \sim D_i(\mathbf{v})}\left[u_i(\mathcal{M}(a'_i, a_{-i}))\right] \geq \lambda \cdot \mathrm{Wel}(\mathbf{v}) - \mu \cdot \mathrm{Rev}(a; \mathcal{M})$$

It is known that a smooth auction in isolation achieves approximately optimal welfare at any equilibrium.

Proposition 2 (Roughgarden, 2012, Syrgkanis and Tardos, 2013). *Let $\mathcal{F}^{\Pi}$ be the family of all possible product type distributions. If a mechanism $\mathcal{M}$ is (λ, μ)-smooth for $\lambda > 0$ and $\mu \geq 1$, then the price of anarchy of $\mathcal{M}$ within the family of distributions $\mathcal{F}^{\Pi}$ is at most $\frac{\mu}{\lambda}$, i.e.,* $\mathrm{PoA}(\mathcal{M}, \mathcal{F}^{\Pi}) \leq \frac{\mu}{\lambda}$.

We now show the main result of this section: if a smooth auction is followed by a secondary market that satisfies voluntary participation and weak budget balance, then the combined market is smooth as well, and hence the price of anarchy is bounded for product type distributions.

Theorem 1. *Let $\mathcal{F}^{\Pi}$ be the family of all possible product type distributions. For any signaling protocol Γ and any trade mechanism $\mathcal{M}^2$ in the secondary market, if an auction mechanism $\mathcal{M}$ is (λ, μ)-smooth for $\lambda \in (0, 1]$ and $\mu \geq 1$, the combined mechanism $\mathcal{M}^C = \mathcal{G}(\mathcal{M}, \Gamma, \mathcal{M}^2)$ is (λ, μ)-smooth. Thus, the price of anarchy of $\mathcal{M}^C$ within the family of distributions $\mathcal{F}^{\Pi}$ for the combined market is at most $\frac{\mu}{\lambda}$, i.e.,* $\mathrm{PoA}(\mathcal{M}, \mathcal{F}^{\Pi}) \leq \frac{\mu}{\lambda}$.

The proof of Theorem 1 is given in Appendix A. Here we comment on the implications of the theorem. First, note that the welfare bound does not depend on the details of the information revelation structure Γ or the trade mechanism $\mathcal{M}^2$ adopted in the secondary market, so the bounds hold for any choice of each. Moreover, our reduction framework does not require refinements on the equilibrium such as sequential equilibrium in the combined market to show that the price of anarchy is small — the result holds for *any* Bayes-Nash equilibrium.

We also note that Theorem 1 extends directly to settings with multiple secondary markets executed sequentially, with any information released between each market. This is because combining a smooth auction with a trade mechanism results with a new smooth mechanism (with the same parameters), which we can now view as a smooth auction to be combined with the next trade mechanism.

Theorem 1 establishes a robust welfare guarantee for a combined market as long as the initial auction is (λ, μ)-smooth. For multi-unit allocation problems (such as in our basic setting), it is known that the discriminatory price auction is $(1 - 1/e, 1)$-smooth [de Keijzer et al., 2013]. We therefore obtain the following corollary.

Corollary 1. *Consider any multi-unit auction setting with non-increasing marginal values where the agents' valuations are distributed independently. Let $\mathcal{M}^C$ be a combined mechanism that runs the discriminatory price auction followed by an arbitrary signaling protocol Γ and trade mechanism $\mathcal{M}^2$. Then at any BNE of $\mathcal{M}^C$ the expected welfare is at least $(1 - 1/e)$ times the expected optimal welfare.*

As we know from Section 3, a similar welfare bound does not hold for uniform-price auctions. This is because, unlike the discriminatory price auction, the uniform-price auction is not (λ, μ)-smooth for any positive constants λ and μ.[14]

[14] The uniform-price auction does satisfy a relaxed version of smoothness: it is weakly $(1 - 1/e, 1)$-smooth, which implies a constant welfare bound in isolation as long as agents avoid weakly dominated "overbidding" strategies [de Keijzer et al.,

We can also apply Theorem 1 to other smooth auctions for more general allocation problems, such as submodular combinatorial auctions. This can capture, for example, a scenario where different license types are being auctioned off simultaneously. See Appendix A.3 for further details on the implied welfare bounds.

In the remainder of this section we will extend Theorem 1 in two ways. In Section 4.1 we show that if the auction $\mathcal{M}$ satisfies a stronger notion of smoothness known as semi-smoothness, the price of anarchy of the combined market is bounded even for correlated type distributions. It turns out that the discriminatory auction is $(1-1/e, 1)$-semi-smooth, so the welfare bound from Corollary 1 applies even if agent valuations are correlated. Second, in Section 4.2 we show that our welfare bound continues to hold even if agents are allowed to purchase signals correlated with the value realizations of agents in advance of the auction. For example, this captures settings in which a speculator could invest in market research before participating in the auction.

4.1 Extension: correlated valuations

As Proposition 2 is proven only for independent value distributions, Theorem 1 likewise applies only to product distributions. As it turns out, we can extend Theorem 1 to derive similar results for correlated distributions based on semi-smoothness [Lucier and Paes Leme, 2011].

Definition 4 (Lucier and Paes Leme, 2011, Roughgarden et al., 2017)**.** *Auction $\mathcal{M}$ with action space A is (λ, μ)-*semi-smooth *for $\lambda > 0$ and $\mu \geq 1$, if for any valuation profile $\mathbf{v}$, there exists action distributions $\{D_i(v_i)\}_{i\in[n]}$ such that for any action profile $a \in A$,*

$$\sum_{i\in[n]} \mathbf{E}_{a'_i \sim D_i(v_i)}\left[u_i(\mathcal{M}(a'_i, a_{-i}); v_i)\right] \geq \lambda \mathrm{Wel}(\mathbf{v}) - \mu \mathrm{Rev}(a; \mathcal{M})$$

The main difference between the definition of semi-smooth and smooth is that for each agent i, the deviating action distribution $D_i(v_i)$ in semi-smooth only depends on her private valuation v_i, not the entire valuation profile $\mathbf{v}$.

Proposition 3 (Lucier and Paes Leme, 2011)**.** *If a mechanism $\mathcal{M}$ is (λ, μ)-semi-smooth for $\lambda > 0$ and $\mu \geq 1$, then the price of anarchy of $\mathcal{M}$ within the family of all distributions $\mathcal{F}$ is at most $\frac{\mu}{\lambda}$, i.e.,* $\mathrm{PoA}(\mathcal{M}, \mathcal{F}) \leq \frac{\mu}{\lambda}$.

Similarly to Theorem 1, we next show that combining a (λ, μ)-semi-smooth auction with and any signaling protocol and any trade mechanism happening aftermarkets, the resulting mechanism in the combined market has small price of anarchy for arbitrary distributions.

Theorem 2. *Let $\mathcal{F}$ be the family of all possible type distributions. For any signaling protocol Γ and any trade mechanism $\mathcal{M}^2$ in the secondary market, if a mechanism $\mathcal{M}$ is (λ, μ)-semi-smooth for $\lambda \in (0, 1]$ and $\mu \geq 1$, the combined mechanism $\mathcal{M}^C = \mathcal{G}(\mathcal{M}^1, \Gamma, \mathcal{M}^2)$ is (λ, μ)-semi-smooth. Thus, the price of anarchy of $\mathcal{M}$ within the family of distributions $\mathcal{F}$ for the combined market is at most $\frac{\mu}{\lambda}$, i.e.,* $\mathrm{PoA}(\mathcal{M}, \mathcal{F}) \leq \frac{\mu}{\lambda}$.

The proof of Theorem 2 is essentially identical to Theorem 1 (up to replacing $D_i(\mathbf{v})$ by $D_i(v_i)$) and hence omitted here. We can now use results regarding semi-smooth auction from the literature to prove that the price of anarchy of the corresponding combined markets is bounded.

Proposition 4 (de Keijzer et al., 2013)**.** *For multi-unit auctions with non-increasing marginal values, the discriminatory auction is $(1 - {}^1/_e, 1)$-semi-smooth.*[15]

2013, Syrgkanis and Tardos, 2013]. In contrast, our example in Section 3 shows that such a refinement is not sufficient to guarantee good welfare guarantees when this auction is part of a combined market.

[15]In de Keijzer et al. [2013], the authors only explicitly state that the discriminatory auction is smooth. However, their construction directly implies that the discriminatory auction is semi-smooth.

4.2 Extension: Acquiring Additional Information

We consider the extension where agents can acquire costly information about other agents' private types before the auction starts. This captures the application where speculators gather information on the demands in the carbon market, and use the acquired information to improve their utilities through buying items in the auction and reselling them more expensively in the secondary market. In general this could have a negative impact on the equilibrium welfare of the combined market. In this section, we show that if the designer uses smooth auctions, then the welfare guarantees we obtained in Theorems 1 and 2 hold even when agents can acquire costly information.

Information is captured by a signal from the types of the others to some signal space. Specifically, for each agent i, let Ψ_i be the set of feasible signal structures for agent i. Any signal structure $\psi_i \in \Psi_i$ is a mapping from the opponents' valuations $\mathbf{v}_{-i}$ to a distribution over the signal space S_i. Note that the information acquisition is potentially costly, i.e., there is a non-negative cost $c_i(\psi_i; v_i)$ for any $\psi_i \in \Psi_i$ and any type v_i of i. Let $\bar{\psi}$ be the signal that acquires no information with zero cost. We assume that $\bar{\psi} \in \Psi_i$ for any agent i. In our model, both the cost function c_i and the set Ψ_i of any agent i are common knowledge among all agents.

In the following theorem, we extend Theorem 1 and show that the price of anarchy for smooth auctions in the combined market is bounded when agents can acquire information on the competitors.

Proposition 5. *Let $\mathcal{F}^{\Pi}$ be the family of all possible product type distributions. For any set of signals Ψ and any cost function c, if mechanism $\mathcal{M}$ is (λ, μ)-smooth for $\lambda \in (0, 1]$ and $\mu \geq 1$, then the price of anarchy of $\mathcal{M}$ within the family of distributions $\mathcal{F}^{\Pi}$ for the combined market with information acquisition is at most $\frac{\mu}{\lambda}$, i.e.,* $\mathrm{PoA}(\mathcal{M}, \mathcal{F}^{\Pi}) \leq \frac{\mu}{\lambda}$.

The proof of Proposition 5 is provided in Appendix A. Similarly, one can extend Theorem 2 for all family of all possible type distributions when the mechanism $\mathcal{M}$ is (λ, μ)-semi-smooth. As the proof is similar to the proof of Proposition 5, we omit it.

5 WELFARE GUARANTEES UNDER POSTED PRICING

In the previous section we showed that the welfare guarantees derived from smooth auctions, such as discriminatory price auctions, are robust to the presence of a secondary market. Such smooth auctions have the advantage of being agnostic to the prior distributions from which agent valuations are drawn. A downside is that bidding in such auctions can be quite complex: constructing an optimal bidding strategy requires sophisticated reasoning and the ability to predict market conditions. If the designer (i.e., government) has a sense of the market conditions, then a tempting alternative to running a smooth auction is to sell a fixed quantity of licenses at a pre-specified unit price, while supplies last. Such posted-price mechanisms have the advantage of being very simple to participate in, since each potential buyer can simply purchase her utility-maximizing bundle of the remaining licenses at the given prices.

A recent literature on static posted pricing and Prophet Inequalities has illustrated that such pricing methods can provide strong welfare guarantees in a variety of allocation problems, even when the order of buyer arrival is adversarial. See [Lucier, 2017] for a recent survey. For example, consider the special case where there is a single indivisible good to be sold and each buyer's value is drawn independently. It is known that if there is no secondary market, then setting a single take-it-or-leave-it price and (while the item is still available) letting buyers in sequence each choose whether to purchase, guarantees half of the expected maximum. Does this guarantee still hold in the presence of a secondary market?

As it turns out, the answer depends on how the fixed price is selected. The canonical solution to the single-item problem, based on the classic Prophet inequality, chooses price $p^* = \sup\{p :$

$\Pr[\max_i v_i > p] \geq 1/2\}$. That is, the median of the distribution over maximum values. This choice of price guarantees half of the expected optimal welfare with no secondary market [Samuel-Cahn, 1984], but the following example shows that this is no longer the case in the presence of a secondary market.[16]

Example 2. *There is a single item to be sold and two potential buyers. The first buyer has value v_1 drawn uniformly from $[0, 1]$. The second buyer has value v_2 equal to 0 with probability $1 - \epsilon < 1$, and with the remaining probability ϵ the value v_2 is set equal to ϵ^{-1} times a random variable drawn from the equal-revenue distribution capped at H. That is, with probability $\epsilon > 0$, $v_2 = \epsilon^{-1} z$ where z is drawn from a distribution with CDF $F(z) = \frac{z-1}{z}$ for $z \in [1, H]$ and $F(H) = 1$, and with the remaining probability $v_2 = 0$.*

The efficient allocation gives the good to buyer 2 whenever $v_2 > 0$, leading to an expected welfare of at least $\epsilon \times \epsilon^{-1} \times \mathbf{E}[z] = \Theta(\log H)$.

Note that the median price is $p^* = \sup\{p : \Pr[\max_i v_i > p] \geq 1/2\} = \frac{1}{2(1-\epsilon)}$. Suppose we offer this price to each buyer in sequence, starting with buyer 1. If there is no secondary market, then the first buyer will purchase only if $v_1 \geq p^*$, which occurs with probability less than $1/2$. The item is therefore available for purchase for the second buyer with probability at least $1/2$, leading to an expected welfare of $\Omega(\log(H))$. This mechanism therefore obtains a constant fraction of the optimal welfare when running in isolation.

Now suppose that the posted-price mechanism is followed by a secondary market in which the winning buyer (if any) can post a take-it-or-leave-it price offer to the losing buyer. In this case the first buyer would always prefer to purchase the item at price $p^* = \frac{1}{2(1-\epsilon)}$, and then offer to resell it to the second buyer in the secondary market at price $p' = \epsilon^{-1} H$. Note that this choice of p' is revenue-maximizing, assuming that no extra information about buyer valuations is revealed between the auction and the secondary market, and obtains expected revenue $1 \geq v_1$. This is therefore a sequential equilibrium. The expected welfare at this equilibrium is $O(1)$, since $v_1 = O(1)$ and $\Pr[v_2 \geq \epsilon^{-1} H] = \epsilon H^{-1}$. Taking H sufficiently large leads to an arbitrarily large welfare gap.

This example shows that the pricing strategy based the median of maximum values can lead to significant incentives for a low-value buyer to purchase with the intention to resell. Then, due to monopolist distortions, significant welfare is subsequently lost in the secondary market.

Another approach to setting static posted prices is based on so-called “balanced prices" [Dutting et al., 2020, Kleinberg and Weinberg, 2012]. In the single-item example described above, this corresponds to setting a price equal to $\frac{1}{2}\mathbf{E}[\max_i v_i]$. This approach likewise guarantees half of the optimal welfare for the single-item prophet inequality problem [Kleinberg and Weinberg, 2012]. We will show that, unlike Example 2, this guarantee continues to hold even in the presence of a secondary market. This is true not only for single-item auctions, but for the broader broad class of multi-unit auctions, and even to combinatorial allocation problems.

To formalize this claim, we make use of a general definition of balanced prices due to Dutting et al. [2020]. We will state this definition in a general combinatorial auction setting in which the set of items M are not necessarily identical. A pricing function p assigns to each set of items $x \subseteq M$ a price $p(x) \geq 0$. For example, this function might assign a price to each individual item and set $p(x)$ to be the sum of item prices, or more generally $p(x)$ might assign arbitrary prices to each bundle.[17]

[16]The guarantee of half of the expected maximum welfare with no secondary market requires some care in the case that a value is precisely equal to p. But this occurs with probability 0 in our example, so our negative result holds regardless of how this is handled.

[17]This definition naturally extends to fractional or randomized allocations, but in order to keep notation simple, in this section we will restrict attention to deterministic allocations. Theorem 3 can be extended to fractional or randomized allocations with the appropriate notational adjustments.

First some notation. Write $\mathrm{OPT}(\mathbf{v}, S)$ for the welfare-optimal allocation of the items in $S \subseteq M$ when the valuations are $\mathbf{v}$, and denote $\mathrm{OPT}(\mathbf{v}) = \mathrm{OPT}(\mathbf{v}, M)$. In a slight abuse of notation, write $M \backslash \mathbf{x}$ for $M \backslash (\cup_i x_i)$, the set of items that are unallocated in feasible solution $\mathbf{x}$. Then, in particular, $\mathrm{OPT}(\mathbf{v}, M \backslash \mathbf{x})$ is the welfare-optimal allocation of the items that are unallocated in $\mathbf{x}$.

The following is a definition from Dutting et al. [2020] specialized to our setting.

Definition 5. *For a given valuation realization $\mathbf{v}$, a pricing function p is (α, β)-balanced if, for any pair of allocations $\mathbf{x}$, $\mathbf{x}'$ that are disjoint and jointly feasible (i.e., they allocate disjoint sets of items), we have*

- $\sum_{i \in N} p(x_i) \geq \frac{1}{\alpha}(\mathbf{v}(\mathrm{OPT}(\mathbf{v})) - \mathbf{v}(\mathrm{OPT}(\mathbf{v}, M \backslash \mathbf{x})))$,
- $\sum_{i \in N} p(x_i') \leq \beta \cdot \mathbf{v}(\mathrm{OPT}(\mathbf{v}, M \backslash \mathbf{x}))$

That is, prices p are balanced if the price paid for any allocation $\mathbf{x}$ is at least the loss in optimal welfare due to losing the items in $\mathbf{x}$ (up to a factor of α). Secondly, the total price of any allocation $\mathbf{x}'$ that remains feasible after $\mathbf{x}$ is removed, is at most β times the optimal welfare achievable using the items not allocated in $\mathbf{x}$. Note that $\mathbf{x}'$ need not be the welfare-optimal allocation of the items in $M \backslash \mathbf{x}$.

Note that this definition of balanced prices is with respect to a particular realization $\mathbf{v}$ of agent types. We will use this to construct a static pricing rule (i.e., prices that are independent of realizations) by taking an expectation over types. Formally, we say that a posted-price mechanism uses (α, β)-balanced prices if it (a) defines an (α, β)-balanced price function $p^{\mathbf{v}}$ for all valuation profiles $\mathbf{v}$, then (b) sets its actual static price function p according to $p(x) = \frac{\alpha}{1+\alpha\beta} E_{\mathbf{v}}[p^{\mathbf{v}}(x)]$. That is, prices are set by taking expectations of the type-specific prices over the buyer types and multiplying by the constant $\frac{\alpha}{1+\alpha\beta}$. For example, in the single-item prophet inequality setting, a $(1,1)$-balanced price for a given realization of values $\mathbf{v}$ is the price $p^{\mathbf{v}} = \max_i v_i$. The appropriate choice of posted price for the prophet inequality is then $\frac{1}{2}\mathbf{E}[\max_i v_i]$, the expectation of the balanced prices times $\frac{1}{2}$.

It is known that in the absence of a secondary market, using balanced prices leads to a strong welfare guarantee.

Proposition 6 (Dutting et al., 2020)**.** *Fix the valuation distributions $F = \times_i F_i$. If a posted-price mechanism $\mathcal{M}$ uses (α, β)-balanced prices for $\alpha, \beta \geq 1$, then the expected welfare obtained by $\mathcal{M}$ is at least $\frac{1}{1+\alpha\beta}$ times the expected optimal welfare.*

As we now show, this result extends to any equilibrium in the combined market setting with an arbitrary trade mechanism being used as a secondary market. This theorem applies to multi-unit and combinatorial auction allocation problems.

Theorem 3. *Fix the valuation distributions $F = \times_i F_i$. For any signaling protocol Γ and any trade mechanism $\mathcal{M}^2$ in the secondary market, if a posted-price mechanism $\mathcal{M}^1$ uses (α, β)-balanced prices for $\alpha, \beta \geq 1$, then the expected welfare obtained by the combined mechanism $\mathcal{M}^C$ at any Bayesian-Nash equilibrium is at least $\frac{1}{1+\alpha\beta}$ times the expected optimal welfare.*

Our argument is an adaptation of a proof method due to Dutting et al. [2020]. The details of the proof is given in Appendix C.

5.1 Application: Balanced Prices for Carbon Markets

We can apply Theorem 3 to any allocation problem for which balanced prices exist. For example, it is known that one can design $(1,1)$-balanced item prices for submodular combinatorial auctions [Feldman et al., 2014]. A multi-unit auction with weakly decreasing marginal values is a special case of a submodular combinatorial auction in which all items are identical. We can therefore conclude from Theorem 3 the existence of item prices that guarantee half of the expected optimal welfare in our model of a carbon license market with an arbitrary aftermarket.[18]

One thing to note is that in the standard construction for submodular combinatorial auctions, there is no guarantee that all items will be assigned the same price even if the items are identical. We next argue that in fact there always exist $(1, 1)$-balanced prices that are identical across items, so in fact there is only a single price that can be interpreted as a per-unit price offered to all prospective buyers.

Our construction is as follows. For each valuation profile $\mathbf{v}$, calculate the optimal allocation $\mathbf{x}^*(\mathbf{v})$ by greedily allocating in order of highest marginal value. Define $w^{\mathbf{v}} \triangleq \frac{1}{m}\sum_i v_i(x_i^*(\mathbf{v})) \geq 0$ to be the average per-unit welfare of this optimal allocation. We use the average per-unit welfare $w^{\mathbf{v}}$ as a price for each unit of the item, so in particular all units have the same price. That is, the price to acquire k units is $p^{\mathbf{v}}(k) = w^{\mathbf{v}} \times k$. We claim that this choice of unit prices is $(1, 1)$-balanced.

Claim 1. *For a multi-unit allocation problem, for any profile* $\mathbf{v}$ *of valuations with non-increasing marginal values, the price function* $p^{\mathbf{v}}(x) = w^{\mathbf{v}} \cdot |x|$ *where* $w^{\mathbf{v}} = \frac{1}{m}\sum_i v_i(x_i^*(\mathbf{v}))$ *is* $(1, 1)$*-balanced.*

Proof. Choose any pair of allocations $(\mathbf{x}, \mathbf{x}')$ as in the definition of balancedness. Suppose $\mathbf{x}$ allocates $k \leq m$ items in total, meaning that $\sum_i x_i = k$ Since allocation $\mathbf{x}'$ is feasible after the items allocated under $\mathbf{x}$ are removed, we have that $\mathbf{x}'$ allocates at most $m - k$ items.

For the first condition of balancedness, note that $\sum_i p^{\mathbf{v}}(x_i)$ is k times the per-unit welfare of the optimal allocation. On the other hand, $\mathbf{v}(\mathrm{OPT}(\mathbf{v})) - \mathbf{v}(\mathrm{OPT}(\mathbf{v}, M\backslash\mathbf{x}))$ is the loss of welfare due to removing any k items from the auction, which is precisely the sum of the k smallest marginal values in the optimal allocation. The average of these k smallest marginal values is at most the overall average, and hence $\sum_i p^{\mathbf{v}}(x_i) \geq \mathbf{v}(\mathrm{OPT}(\mathbf{v})) - \mathbf{v}(\mathrm{OPT}(\mathbf{v}, M\backslash\mathbf{x}))$.

For the second condition, note that $\sum_i p^{\mathbf{v}}(x_i')$ is at most $m - k$ times the per-unit welfare of the optimal allocation. On the other hand, $\mathbf{v}(\mathrm{OPT}(\mathbf{v}, M\backslash\mathbf{x}))$ is the value of the optimal allocation of $m - k$ elements, which is the sum of the largest $m - k$ marginals in the overall optimal allocation. The average of these $m - k$ largest marginals is at least the overall average, and hence $\sum_i p^{\mathbf{v}}(x_i') \leq \mathbf{v}(\mathrm{OPT}(\mathbf{v}, M\backslash\mathbf{x}))$.

We conclude that both conditions of balancedness are satisfied with $(\alpha, \beta) = (1, 1)$. $\square$

Given this choice of $(1, 1)$-balanced prices, Theorem 3 then implies that setting a static price per unit that equals to half of the expected per-unit welfare, $\frac{1}{2m}\mathbf{E}[\sum_i v_i(x_i(\mathbf{v}))]$, guarantees half of the optimal expected welfare in any equilibrium, under any secondary market implementable as a trade mechanism and for any set of information revealed between the posted-price mechanism and the secondary market.

Example 3. *Recall the lower bound market example in Section 3, which illustrated a low-welfare equilibrium in a uniform auction with an aftermarket. In that example, the welfare-maximizing allocation always allocates two goods to bidder A and the remaining* $m - 2$ *goods to bidder B, for a total expected welfare of* $\Theta(\log m)$*. In this example, Theorem 3 together with the* $(1, 1)$*-balanced prices we present in Claim 1 would set a per-unit price of* $E[v(\mathrm{OPT}(v))]/(2m) = \Theta((\log m)/m)$ *and allow bidders to purchase as many units as desired at this price, while supplies last. Theorem 3 implies that, at this price, any purchasing behavior at equilibrium will achieve at least half of the expected optimal welfare. Intuitively, the price is high enough to discourage the speculator agent C from buying more than two items (since at most two items can be sold to agent A, and any items sold to agent B generate an expected revenue of at most* $1/m$ *each, which is less than the auction per-unit reserve) but also low enough that agent B will nearly always decide to purchase most of the available items.*

[18]The approximation ratio is tight even for the single-item setting [Kleinberg and Weinberg, 2012].

6 BEST OF BOTH WORLDS: UNIFORM-PRICE AUCTION WITH RESERVES

We realize that moving from a uniform auction to a posted price mechanism is a radical change to the carbon allowances market. We thus consider the problem of finding a minimal change to the current auction which will still take care of the problem of significant efficiency loss that can happen as a result of combing the auction with a secondary market (as illustrated in the example presented in Section 3). As we now show, a sufficient change is to add to the uniform-price auction a per-unit reserve price based on balanced prices. Any marginal bids strictly less than this reserve price are ignored, but bids exactly equal to the reserve are allowed. Formally, in the following theorem we show that the welfare bound from Theorem 3 continues to hold when we add an appropriately chosen per-unit reserve price to the uniform-price auction. The proof of Theorem 4 is given in Appendix C.

Theorem 4. *For a multi-unit allocation problem with non-increasing marginal values sampled independently from $F = \times_i F_i$, consider a uniform-price auction subject to a per-unit reserve price of $p = \frac{1}{2m} E_{\mathbf{v}\sim F}[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))]$ Then for any trade mechanism, and at any Bayesian Nash equilibrium of the combined market, the expected welfare is at least half of the expected optimal welfare.*

Estimation Errors and Calculating Reserve Prices. The reserve price in Theorem 4 depends on the expected optimal welfare attainable by allocating licenses. As it turns out, the welfare guarantee in Theorem 4 is robust to mistakes in the prices. This robustness was noted by Feldman et al. [2014] for posted-price mechanisms in submodular combinatorial auctions, and the argument extends to our setting without change. Specifically, for any $\epsilon > 0$, if instead of setting reserve $p = \frac{1}{2m} E[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))]$ we set a reserve price p' such that $|p' - p| < \epsilon$, then the expected welfare of the resulting combined market is at least half of the expected optimal welfare less $m\epsilon$, where recall that m is the total number of licenses for sale. To see why, recall the proof of Theorem 4 and note that lowering a license's price by ϵ reduces the revenue obtained by at most ϵ, and increasing a license's price by ϵ reduces buyer surplus from purchasing that item by at most ϵ. Collecting up these additive losses leads to a total loss of at most $m\epsilon$. A formal statement and proof are deferred to Appendix D.

One implication of this robustness result is that if one sets reserve prices based on a noisy estimate of $E[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))]$, such as obtained by samples or historical data, then the welfare guarantee one obtains will degrade in proportion to the estimation error.

Corollary 2. *For a multi-unit allocation problem with non-increasing marginal values sampled independently from $F = \times_i F_i$, suppose that the seller has access to an estimate ψ to the optimal welfare such that with probability at least $1 - \delta$, $\psi \in [(1-\epsilon)\mathbf{v}(\mathrm{OPT}(\mathbf{v})), (1+\epsilon)\mathbf{v}(\mathrm{OPT}(\mathbf{v}))]$ for constants $\delta, \epsilon \in [0, 1]$. Consider the uniform-price auction subject to a per-unit reserve price of $p = \frac{\psi}{2m}$. Then with probability at least $1 - \delta$, for any trade mechanism, and at any Bayesian Nash equilibrium of the combined market, the expected welfare is at least $\frac{1-\epsilon}{2}$ fraction of the expected optimal welfare.*

REFERENCES

Bloomberg (2021). EU urged to protect carbon market from excessive speculation. https://www.bloomberg.com/news/articles/2021-12-15/eu-urged-to-protect-carbon-market-from-excessive-speculation.

CARB (2022). Cap-and-Trade Program. https://ww2.arb.ca.gov/resources/documents/faq-cap-and-trade-program.

Carroll, G. and Segal, I. (2019). Robustly optimal auctions with unknown resale opportunities. *The Review of Economic Studies*, 86(4):1527–1555.

Cotton, D. (2015). Emissions Trading Design – A Critical Overview, edited by stefan e. weishaar. published by edward elgar, uk, 2014, pp. 249, isbn: 978 1 78195 221 4, aud$114.00 (hardcover). *Australian Journal of Agricultural and Resource Economics*, 59(1):156–158.

Cramton, P. and Kerr, S. (2002). Tradeable carbon permit auctions: How and why to auction not grandfather. *Energy Policy*, 30(4):333–345.

Cramton, P., MacKay, D. J., Ockenfels, A., and Stoft, S. (2017). *Global carbon pricing: the path to climate cooperation*. The MIT Press.

de Keijzer, B., Markakis, E., Schäfer, G., and Telelis, O. (2013). Inefficiency of standard multi-unit auctions. In *European Symposium on Algorithms*, pages 385–396. Springer.

Dutting, P., Feldman, M., Kesselheim, T., and Lucier, B. (2020). Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582.

Dworczak, P. (2020). Mechanism design with aftermarkets: Cutoff mechanisms. *Econometrica*, 88(6):2629–2661.

Eden, A., Feldman, M., Talgam-Cohen, I., and Zviran, O. (2020). Price of anarchy of simple auctions with interdependent values. *arXiv preprint arXiv:2011.00498*.

EEX (2021). EU ETS Auctions. https://www.eex.com/en/markets/environmental-markets/eu-ets-auctions.

EPA (2021). Causes of climate change. https://www.epa.gov/climatechange-science/causes-climate-change.

EU ETS (2021a). Auctions by the common auction platform. https://ec.europa.eu/clima/system/files/2021-10/policy_ets_auctioning_cap_report_202109_en.pdf.

EU ETS (2021b). European Commission Auctioning. https://ec.europa.eu/clima/eu-action/eu-emissions-trading-system-eu-ets/auctioning_en.

EU ETS (2022). Auctioning regulation. https://www.emissions-euets.com/auctioning-regulation.

Europa Union Law (2021). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02003L0087-20180408.

Feldman, M., Fu, H., Gravin, N., and Lucier, B. (2013). Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM Symposium on Theory of Computing*, pages 201–210.

Feldman, M., Gravin, N., and Lucier, B. (2014). Combinatorial auctions via posted prices. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 123–135. SIAM.

Goldner, K., Immorlica, N., and Lucier, B. (2020). Reducing inefficiency in carbon auctions with imperfect competition. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*.

Hafalir, I. and Krishna, V. (2008). Asymmetric auctions with resale. *American Economic Review*, 98(1):87–112.

Hafalir, I. and Krishna, V. (2009). Revenue and efficiency effects of resale in first-price auctions. *Journal of Mathematical Economics*, 45(9-10):589–602.

Haile, P. A. (2003). Auctions with private uncertainty and resale opportunities. *Journal of Economic theory*, 108(1):72–110.

Kleinberg, R. and Weinberg, S. M. (2012). Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136.

Lebrun, B. (2010). Revenue ranking of first-price auctions with resale. *Journal of Economic Theory*, 145(5):2037–2043.

Lucier, B. (2017). An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47.

Lucier, B. and Paes Leme, R. (2011). GSP auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 71–80.

Markakis, E. and Telelis, O. (2015). Uniform price auctions: Equilibria and efficiency. *Theory of Computing Systems*, 57(3):549–575.

NOAA (2020). NOAA Annual 2020 Global Climate Report. https://www.ncdc.noaa.gov/sotc/global/202013.

Paes Leme, R., Syrgkanis, V., and Tardos, É. (2012). Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM Symposium on Discrete Algorithms*, pages 869–886. SIAM.

Pagnozzi, M. (2007). Bidding to lose? Auctions with resale. *The RAND Journal of Economics*, 38(4):1090–1112.

Quemin, S. and Pahle, M. (2021). Financials threaten to undermine the functioning of emissions markets. *Available at SSRN 3985079*.

Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM Symposium on Theory of Computing*, pages 513–522.

Roughgarden, T. (2012). The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 862–879.

Roughgarden, T., Syrgkanis, V., and Tardos, E. (2017). The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101.

Samuel-Cahn, E. (1984). Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216.

Susin, I. (2017). Auctions with resale: a survey. *HSE Economic Journal*, 21(2):333–350.

Swiss Re Institute (2021). Swiss Re Institute April Report. https://www.swissre.com/dam/jcr:e73ee7c3-7f83-4c17-a2b8-8ef23a8d3312/swiss-re-institute-expertise-publication-economics-of-climate-change.pdf.

Syrgkanis, V. (2012). Bayesian games and the smoothness framework. *arXiv preprint arXiv:1203.5155*.

Syrgkanis, V. and Tardos, E. (2012). Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 929–944.

Syrgkanis, V. and Tardos, E. (2013). Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM Symposium on Theory of Computing*, pages 211–220.

UN Goals (2019). Goal 13: Take urgent action to combat climate change and its impacts. https://www.un.org/sustainabledevelopment/climate-change.

Weitzman, M. L. (1974). Prices vs. quantities. *The review of economic studies*, 41(4):477–491.

A MISSING PROOF FOR SMOOTH AUCTIONS

A.1 Proof of Theorem 1

First we finish the task of proving Theorem 1.

Proof of Theorem 1. Let $\mathcal{M}^1 = \mathcal{M}$ for clarity of notation. Let A^1, A^2 be the action spaces of mechanisms $\mathcal{M}^1, \mathcal{M}^2$ respectively, and let A^C be the action space of the combined market. Note that $a^C \in A^C$ is an extensive form action, which is choosing action $a^1 \in A^1$ in the first market, and then choosing action $a^2 \in A^2$ based on the allocation, payment, and the signal realized in the first market. For each agent i, by voluntary participation, there exists action $\hat{a}_i^2 \in A_i^2$ such that her payoff is at least her value of the initial allocation in the secondary market. Since mechanism $\mathcal{M}^1$ is (λ, μ)-smooth, for any valuation profile $\mathbf{v}$, there exists action distributions $\{D_i^1(\mathbf{v})\}_{i\in[n]}$ such that for any action profile $a^1 \in A^1$,

$$\sum_{i\in[n]} \mathbf{E}_{\hat{a}_i^1\sim D_i^1(\mathbf{v})}\left[u_i(\mathcal{M}^1(\hat{a}_i^1, a_{-i}^1))\right] \geq \lambda \mathrm{Wel}(\mathbf{v}) - \mu \mathrm{Rev}(a^1; \mathcal{M}^1).$$

For any valuation profile $\mathbf{v}$, let $\{D_i^C(\mathbf{v})\}_{i\in[n]}$ be the distributions over actions $\hat{a}_i^C$ for each agent i where $\hat{a}_i^C$ chooses action $\hat{a}_i^1$ according to distribution $D_i^1(\mathbf{v})$, and then always chooses action $\hat{a}_i^2$ regardless of the signal realization s_i from Γ, or the allocation and payment in the first market. For any $a^C \in A^C$, we have

$$\begin{aligned}\sum_{i\in[n]} \mathbf{E}_{\hat{a}_i^C\sim D_i^C(\mathbf{v})}\left[u_i(\mathcal{M}^C(\hat{a}_i^C, a_{-i}^C))\right] &\geq \sum_{i\in[n]} \mathbf{E}_{\hat{a}_i^1\sim D_i^1(\mathbf{v})}\left[u_i(\mathcal{M}^1(\hat{a}_i^1, a_{-i}^1))\right] \\ &\geq \lambda \mathrm{Wel}(\mathbf{v}) - \mu \mathrm{Rev}(a^1; \mathcal{M}^1) \geq \lambda \mathrm{Wel}(\mathbf{v}) - \mu \mathrm{Rev}(a^C; \mathcal{M}^C)\end{aligned}$$

where the first inequality holds by the definition of $\hat{a}_i^2$, the second as $\mathcal{M}^1$ is (λ, μ)-smooth, and the last inequality holds since mechanism $\mathcal{M}^2$ is weakly budget balanced. Thus, the combined mechanism $\mathcal{M}^C = \mathcal{G}(\mathcal{M}^1, \Gamma, \mathcal{M}^2)$ is (λ, μ)-smooth. Finally, by Proposition 2, we have that the price of anarchy in the combined market is at most $\frac{\mu}{\lambda}$. □

A.2 Proof of Proposition 5

Next we formally prove Proposition 5 for the setting in which agents can acquire costly signals about the valuations of other agents prior to participating in a combined mechanism.

Proof of Proposition 5. Since mechanism $\mathcal{M}$ is (λ, μ)-smooth, by definition, for any value type θ, there exists action distributions $\{D_i(\theta)\}_{i\in[n]}$ such that for any action profile $a \in A$,

$$\sum_{i\in[n]} \mathbf{E}_{a_i'\sim D_i(\theta)}\left[u_i(\mathcal{M}(a_i', a_{-i}); \theta_i)\right] \geq \lambda \mathrm{Wel}(\theta) - \mu \mathrm{Rev}(a; \mathcal{M}).$$

Suppose in equilibrium, for any agent i the information acquisition strategy is $\bar{\sigma}_i(\theta)$ and the bidding strategy is $\hat{\sigma}_i(\theta, s_i)$. Note that since the information acquisition decisions are not revealed to the opponents, the distribution over bids of any agent i is not affected by the information acquisition decisions taken by agent $j \neq i$. Let G_i be the distribution over actions in the auction for agent i under equilibrium strategies $\bar{\sigma}_i$ and $\hat{\sigma}_i$, when her type is distributed according to F_i.

Now consider the following deviating strategy for agent i. Agent i will not acquire any information by adopting signal structure $\bar{\psi}$. Then in the auction, agent i simulates the behavior of the other agents by first sampling $\hat{\theta}_j$ according distribution F_j for any $j \neq i$, and then follow the action distribution $D_i(\theta_i, \hat{\theta}_{-i})$. The expected utility of all agents given this deviating strategy is

$$\begin{aligned}
&\sum_{i\in[n]} \mathbf{E}_{\theta_i\sim F}\left[\mathbf{E}_{a'_i\sim D_i(\theta_i,\hat{\theta}_{-i});a_{-i}\sim G_{-i}}\left[u_i(\mathcal{M}(a'_i, a_{-i}); \theta_i)\right]\right] \\
&= \sum_{i\in[n]} \mathbf{E}_{\theta\sim F}\left[\mathbf{E}_{a'_i\sim D_i(\theta);a_{-i}\sim G_{-i}}\left[u_i(\mathcal{M}(a'_i, a_{-i}); \theta_i)\right]\right] \\
&\geq \mathbf{E}_{\theta\sim F}[\lambda \cdot \mathrm{Wel}(\theta) - \mu \cdot \mathbf{E}_{a\sim G}[\mathrm{Rev}(a; \mathcal{M})]] \\
&= \lambda \cdot \mathrm{Wel}(F) - \mu \cdot \mathbf{E}_{a\sim G}[\mathrm{Rev}(a; \mathcal{M})]
\end{aligned}$$

where the first equality holds by renaming the random variables $\hat{\theta}_{-i}$ as θ_{-i}. The inequality holds by applying the definition of the smoothness and taking expectation over the actions according to distribution G. Note that in every equilibrium $\bar{\sigma}(\theta) = (\bar{\sigma}_1(\theta_1), \ldots, \bar{\sigma}_n(\theta_n))$ and $\hat{\sigma}(\theta, s) = (\hat{\sigma}_1(\theta_1, s_1), \ldots, \hat{\sigma}_n(\theta_n, s_n))$, the utility of any agent is at least her utility given the above deviating strategy. Thus,

$$\begin{aligned}
&\sum_{i\in[n]} \mathbf{E}_{\theta\sim F}\left[\mathbf{E}_{\psi\sim\bar{\sigma}(\theta);s\sim\psi(\theta)}\left[\mathbf{E}_{a\sim\hat{\sigma}(\theta,s)}\left[u_i(\mathcal{M}(a); \theta_i)\right] - c_i(\psi_i, \theta_i)\right]\right] \\
&\geq \lambda \cdot \mathrm{Wel}(F) - \mu \cdot \mathbf{E}_{a\sim G}[\mathrm{Rev}(a; \mathcal{M})]\,. \qquad (1)
\end{aligned}$$

By rearranging the terms and noting that the sum of expected utility is the difference between equilibrium welfare and the expected revenue, we also have

$$\begin{aligned}
&\sum_{i\in[n]} \mathbf{E}_{\theta\sim F}\left[\mathbf{E}_{\psi\sim\bar{\sigma}(\theta);s\sim\psi(\theta)}\left[\mathbf{E}_{a\sim\hat{\sigma}(\theta,s)}\left[u_i(\mathcal{M}(a); \theta_i)\right] - c_i(\psi_i, \theta_i)\right]\right] \\
&= \mathrm{Wel}(\mathcal{M}, (\bar{\sigma}, \hat{\sigma}), F) - \mathbf{E}_{a\sim G}[\mathrm{Rev}(a; \mathcal{M})]\,. \qquad (2)
\end{aligned}$$

Multiplying both sides of the equality (2) with factor μ and combining it with the inequality (1) above, and recalling that $\mu \geq 1$ and the equilibrium utility is non-negative, we have

$$\mu \cdot \mathrm{Wel}(\mathcal{M}, (\bar{\sigma}, \hat{\sigma}), F) \geq \lambda \cdot \mathrm{Wel}(F),$$

i.e., $\mathrm{PoA}(\mathcal{M}, \mathcal{F}^{\Pi}) \leq \frac{\mu}{\lambda}$. □

A.3 Other Smooth Auctions

In Table 1 we list some additional examples of smooth auctions for allocation problems that fall within our framework. For each auction format we note the smoothness bound as well as the implied welfare bound when the auction is used in a combined market.

B LOW WELFARE SEQUENTIAL EQUILIBRIUM

B.1 Sequential Equilibrium in Combined Markets

We now formally define perfect Bayesian equilibrium (PBE) and sequential equilibrium (SE). A fully general definition is beyond the scope of this paper, so we will provide a definition tailored to our setting of combined mechanisms.

Recall our description of the two-stage combined market game $\mathcal{G}(\mathcal{M}^1, \Gamma, \mathcal{M}^2)$. Recall also that a BNE consists of a profile of agent strategies σ^1 for the auction $\mathcal{M}^1$, where σ_i^1 maps agent i's type to an action, and a profile of strategies σ^2 for the secondary auction $\mathcal{M}^2$, where σ_i^2 maps agent i's type and the realization of observations $(x^{\mathcal{M}^1}(a^1), p_i^{\mathcal{M}^1}(a^1), s_i)$ to an action.

A perfect Bayesian equilibrium (PBE) additionally includes a belief function β_i for each agent i, which maps the realization of observations $(x^{\mathcal{M}^1}(a^1), p_i^{\mathcal{M}^1}(a^1), s_i)$ for agent i into a distribution over the types of the other agents. We think of β_i as agent i's posterior belief about the other agents'

Auction Environment	**Mechanism $\mathcal{M}^1$**	**Smoothness**	**PoA in Combined Market $\mathcal{M}^C$**
single-item	first-price auction	$(1 - 1/e, 1)^*$	$e/(e-1)$
	all-pay auction	$(1/2, 1)^*$	2
combinatorial, submodular	simultaneous first price	$(1 - 1/e, 1)^\dagger$	$e/(e-1)$
	simultaneous all pay	$(1/2, 1)^\dagger$	2

* [Roughgarden, 2012] † [Syrgkanis and Tardos, 2013]

Table 1. The first column lists the auction environment and the second the auction mechanism. The third column lists the (λ, μ)-smoothness results from the literature, and by Theorem 1, this implies the price of anarchy upper bound for the combined market, when every valuations are independently distributed, as listed in the last column.

types given the observed outcome of the auction game. We write $\beta = (\beta_1, \dots, \beta_n)$ for the tuple of belief functions.

The collection $(\sigma^1, \sigma^2, \beta)$ forms a perfect Bayesian equilibrium (PBE) if

- β_i is the correct posterior distribution over agent types, given the observations $(x^{\mathcal{M}^1}(a^1), p_i^{\mathcal{M}^1}(a^1), s_i)$, assuming agents behave according to σ^1 in the auction $\mathcal{M}^1$, given the prior type distribution.
- For any realization of auction outcomes and observations, strategy σ_i^2 maximizes the expected utility of agent i in the secondary market $\mathcal{M}^2$ under beliefs β_i.
- Given that agents apply strategies σ^2 in the secondary market, following strategy σ_i^1 in the auction $\mathcal{M}^1$ maximizes the expected utility of agent i.

In words, a PBE satisfies subgame perfection (as agents must best-respond in the secondary market for any realization of the initial auction), and moreover agents update their beliefs rationally and consistently given the strategies of others.

An important subtlety is how Bayes' rule should be applied to events with 0 probability, such as under deviations from the equilibrium. The definition of PBE does not specify what beliefs are to be held by an agent if they observe an auction outcome that has probability 0 given the type distribution and strategy profile σ^1. In sequential equilibrium, these beliefs are constrained by thinking of each strategy as a limit of 'trembling' strategies in which all possible actions have a positive chance of being observed. More formally, a sequential equilibrium (SE) is a PBE such that there exists a sequence of totally mixing strategy profiles $\{\sigma^{1,k}\}$ for the auction mechanism converging to σ^1 and a sequence of beliefs $\{\beta^k\}_{k\geq 1}$ converging to β such that β^k is consistent with each agent applying Bayes' rule to all observations, under the assumption that agents are applying strategy profile $\sigma^{1,k}$.

B.2 Proof of Proposition 1

Proof of Proposition 1. To show that this behavior forms a sequential equilibrium, we must first describe the beliefs of the agents about the valuation profile when entering the secondary market. Specifically, we claim that each agent will have a posterior distribution that is simply equal to the prior distribution over all other agents' valuations. That is, no agent receives an informative signal about the valuations of other agents. These beliefs are indeed consistent with the suggested strategies, since the proposed bidding behavior of each agent in the auction does not depend on her valuation. Furthermore, we assume that agents would rationalize any observed off-path behavior as independent of valuations, which is consistent with retaining the prior beliefs.

Now consider behavior in the secondary market given these beliefs. Whichever price speculator C chooses, it is a dominant strategy for agents A and B to purchase licenses as long as their marginal values exceed the listed price. Given this, it is dominant for agent C to post a price that maximizes revenue, in expectation over her beliefs about the agent values. One can verify that the suggested price of 1 will always maximize revenue for agent C, given the prior valuation distribution, no matter what allocation is generated by the auction. We can therefore assume that, no matter what outcome is generated by the auction, agent C will list all acquired licenses at a price of 1.

Given this behavior in the secondary market, we verify that the proposed bidding behavior in the auction is indeed an equilibrium. Indeed, any unilateral change in an agent's bid will either not change the outcome of the auction (when the acquired quantity is the same or lower), or turn the clearing price to 1 or higher (if the agent obtains more licenses). This results in lower utility to every agent in the combined market, since licenses sell at a price of 1 in the secondary market. □

We note that this sequential equilibrium has an additional desirable property. To describe it, we first note that since the agents have uniquely optimal behavior in the secondary market, it is straightforward define payoffs in the initial uniform-price auction with respect to the utility that will be obtained assuming this secondary market behavior. Thinking of payoffs in this way, our property is that the agents are not employing *weakly dominated strategies* in the primary auction.

Definition 6. *A strategy σ_i^1 is weakly dominated by strategy $\overline{\sigma}_i^1$ if, for every action profile a_{-i}^1 of the other agents, the total expected utility of agent i is no greater when using σ_i^1 than when using $\overline{\sigma}_i^1$, and there exists some action profile for which $\overline{\sigma}_i^1$ is strictly better. Strategy σ_i^1 is* not weakly dominated *if there is no other strategy that weakly dominates it.*

For agents A and B, it is easy to see that the chosen bid profiles are not weakly dominated. Indeed, any bid less than 2 on the first item would be worse against a competing bid strictly between the placed bid and 2 on each item, and any non-zero bid on a second item would be strictly worse against a competitor that bids 2 on $m - 1$ items, as this would only increase the price paid.

To see why the bidding strategy of C is not weakly dominated, notice that if agent A wins two units in the primary market, then the speculator cannot obtain utility by selling to agent A in the secondary market. So any bid on fewer than $m - 2$ units would be strictly worse than bidding 1 on $m-2$ units if agent A bids ϵ on two units and agent B bids 1 on a single unit. Likewise, bidding $b < 1$ on any of the $m - 2$ units would be strictly worse than bidding 1 on $m - 2$ units if agent A bids $b + \epsilon$ on two units and agent B bids 1 on a single item, since the per-unit price is effectively unchanged but the speculator wins one fewer unit and agent A wins one more, causing the speculator to lose the opportunity to sell a unit to agent A. Finally, increasing the bids above 0 on any of the remaining two units can only increase the price paid by agent C, which would strictly reduce her utility.

C MISSING PROOFS FOR BALANCED PRICES

C.1 Proof of Theorem 3

Proof of Theorem 3. Our argument is an adaptation of a proof method due to Dutting et al. [2020]. Our approach is to use the definition of balancedness to derive lower bounds on the utility obtained by the buyers and on the revenue collected by the primary auction mechanism. Since agents achieve non-negative utility at equilibrium, and since the secondary market is budget-balanced, we can add these together to obtain a lower bound on the welfare generated in the combined market.

Since $\mathcal{M}^1$ is a posted-price mechanism that uses (α, β)-balanced prices, we know that it uses a static pricing function p defined by $p(x) = \frac{\alpha}{1+\alpha\beta}\mathbf{E}_{\mathbf{v}}[p^{\mathbf{v}}(x)]$ where for each possible type profile $\mathbf{v}$ the price $p^{\mathbf{v}}$ is (α, β)-balanced for $\mathbf{v}$. Fix some Nash equilibrium of the combined mechanism $\mathcal{M}^C$, and write $\mathbf{x}(\mathbf{v})$ for the allocation obtained by the posted-price mechanism $\mathcal{M}^1$, at this equilibrium, when the valuation profile is $\mathbf{v}$. We emphasize that $\mathbf{x}(\mathbf{v})$ is the allocation from the posted-price mechanism, and does not account for any transfer of items that might occur in the secondary market. We will also write $u_i(\mathbf{v})$ for the utility obtained by agent i in the combined market (including any transfers that occur in the secondary market) when valuations are $\mathbf{v}$.

We first bound the consumer surplus. We will have each player consider a strategy that purchases a certain collection of items in the primary market and avoids participating in the secondary market. This strategy will always be feasible, since $\mathcal{M}^2$ is a trade mechanism.

To describe these strategies, we first sample "phantom" valuations $\mathbf{v}' \sim F$. Buyer i will consider buying the set of items $\mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}), M\backslash\mathbf{x}(v'_i, \mathbf{v}_{-i}))$ at the posted prices. This is buyer i's part of the optimal allocation, under valuations $(v_i, \mathbf{v}'_{-i})$, of all items that are not allocated (at equilibrium) by the posted price mechanism under valuations $(v'_i, \mathbf{v}_{-i})$. This strange choice of valuation profiles is carefully chosen to make the valuation of the allocation independent of the items available, while still coupling with the outcomes under valuation profile $\mathbf{v}$. Since $\mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}), M\backslash\mathbf{x}(v'_i, \mathbf{v}_{-i}))$ is a subset of the items actually available to agent i when agents have valuations $\mathbf{v}$ (as the set of items available to agent i only depends on the valuation profile $\mathbf{v}_{-i}$ in a posted pricing mechanism), her expected utility $\mathbf{E}_{\mathbf{v}}[u_i(\mathbf{v})]$ is at least the expected utility of following this strategy (and not participating in the secondary market). That is,

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}[u_i(\mathbf{v})] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}), M\backslash\mathbf{x}(v'_i, \mathbf{v}_{-i}))) - p(\mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}), M\backslash\mathbf{x}(v'_i, \mathbf{v}_{-i})))] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}[v'_i(\mathrm{OPT}_i(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v}))) - p(\mathrm{OPT}_i(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))].\end{aligned}$$

Here the inequality holds because agent i is choosing her utility-optimal allocation given the prices, and the equality holds by a change of variables (swapping the role of v'_i and v_i).

Summing the previous inequality over all buyers gives

$$\mathbf{E}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'(\mathrm{OPT}(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))\right] - \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_i p(\mathrm{OPT}(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))\right]. \tag{3}$$

Recalling that we set p_i to be $\frac{\alpha}{1+\alpha\beta}$ times $\mathbf{E}_v[p_i^v]$, we can use linearity of expectation plus the definition of balanced prices (applied pointwise to each realization of $\mathbf{v}'$ and $\mathbf{x}(\mathbf{v})$) to conclude that

$$\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_i p(\mathrm{OPT}(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))\right] \leq \beta \cdot \frac{\alpha}{\alpha\beta + 1} \cdot \mathbf{E}_{\mathbf{v},\mathbf{v}'}[\mathbf{v}'(\mathrm{OPT}(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))].$$

Substituting into (3) yields

$$\mathbf{E}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] \geq \left(1 - \frac{\alpha\beta}{1+\alpha\beta}\right)\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'(\mathrm{OPT}(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))\right] = \frac{1}{1+\alpha\beta}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'(\mathrm{OPT}(\mathbf{v}', M\backslash\mathbf{x}(\mathbf{v})))\right]. \tag{4}$$

We next provide a bound on the revenue generated in the posted-price mechanism at equilibrium. Recall that $\mathbf{x}(\mathbf{v})$ is the equilibrium allocation in the posted-price mechanism under valuations $\mathbf{v}$.

From the definition of balanced prices we have

$$\begin{aligned}\sum_i p(x_i(\mathbf{v})) &= \frac{\alpha}{1+\alpha\beta} \sum_i \mathbf{E}_{\mathbf{v}'} \left[p^{\mathbf{v}'}(x_i(\mathbf{v})) \right] \\ &\geq \frac{1}{1+\alpha\beta} \mathbf{E}_{\mathbf{v}'} [\mathbf{v}'(\mathrm{OPT}(\mathbf{v}')) - \mathbf{v}'(\mathrm{OPT}(\mathbf{v}', M \backslash \mathbf{x}(\mathbf{v})))].\end{aligned}$$

Taking expectations over $\mathbf{v} \sim F$ yields

$$\mathbf{E}_{\mathbf{v}} \left[\sum_i p(x_i(\mathbf{v})) \right] \geq \frac{1}{1+\alpha\beta} \mathbf{E}_{\mathbf{v}} [\mathbf{v}(\mathrm{OPT}(\mathbf{v}))] - \frac{1}{1+\alpha\beta} \mathbf{E}_{\mathbf{v}',\mathbf{v}} [\mathbf{v}'(\mathrm{OPT}(\mathbf{v}', M \backslash \mathbf{x}(\mathbf{v})))]. \tag{5}$$

We are finally ready to bound the expected welfare obtained at equilibrium. We first argue that the expected welfare obtained at equilibrium is at least the sum of the expected utility bound (4) and expected revenue bound (5) calculated above. Recall that the utility bound is a lower bound on the expected utility obtained from the buyers in the combined mechanism. Moreover, since the secondary market $\mathcal{M}^2$ is a trade mechanism, it is budget-balanced and hence no additional revenue is extracted or lost. Thus our bound on revenue collected in the auction also bounds the total revenue raised for the combined mechanism. The total welfare obtained by the combined mechanism $\mathcal{M}^C$ is therefore the sum of the consumer surplus and the revenue generated by the posted price mechanism.

Summing up the utility and revenue bounds (4) and (5), we conclude that the total expected welfare of $\mathcal{M}^C$ is at least

$$\mathbf{E}_{\mathbf{v}} \left[\sum_i u_i(\mathbf{v}) \right] + \mathbf{E}_{\mathbf{v}} \left[\sum_i p(x_i(\mathbf{v})) \right] \geq \frac{1}{1+\alpha\beta} \mathbf{E}_{\mathbf{v}} [\mathbf{v}(\mathrm{OPT}(\mathbf{v}))]$$

as claimed. ◻

C.2 Proof of Theorem 4

Proof of Theorem 4. Our analysis is similar to the proof of Theorem 3. Fix some Bayes-Nash equilibrium of the combined mechanism. Let $\mathbf{b}(\mathbf{v})$ be the bids made in the auction at this equilibrium and $\mathrm{OPT}(\mathbf{v}; k)$ be the welfare-optimal allocation of k units when the valuation profile is $\mathbf{v}$. We denote the allocation of the auction under bids $\mathbf{b}$ as $\mathbf{x}(\mathbf{b})$, and the number of items left unallocated by the auction when agents bid according to $\mathbf{b}$ as $Z(\mathbf{b})$. That is, $Z(\mathbf{b}) = m - \sum_i x_i(\mathbf{b})$.

We first bound the buyer surplus at equilibrium. To this end, fix a valuation profile $\mathbf{v}$ and consider a possible deviation by buyer i. Sample phantom valuations $\mathbf{v}' \sim F$. Let $z(v_i)$ be the largest index j such that $v_{ij} \geq p$. Note that due to the reserve price, agent i would obtain negative marginal utility for any items won in excess of $z(v_i)$. Let $y_i = \min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); m)\}$ be agent i's optimal allocation when others' valuations are $\mathbf{v}'_{-i}$, excluding any items for which her marginal value is less than p. Our proposed deviation for agent i is to place an auction bid of b'_i where $b'_{ij} = p$ for $j \leq y_i$, and $b_{ij} = 0$ for $j > y_i$, then not participate in the secondary market. Importantly, this deviation depends on v_i but not v_{-i}. Note also that the utility obtained under this deviation can only be non-negative since $y_i \leq z(v_i)$ and the price paid per item obtained is exactly p.

Under this deviation, either agent i wins y_i items or all items are sold, and in the latter case agent i receives all items not allocated to the other agents under bid profile $b'_i(v'_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})$. Note that this quantity should be at least the number of items such that the bids of agents from $-i$ are strictly below p, which again is at least $Z(b_i(v'_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))$, the number of items unallocated if we also include agent i bidding at equilibrium as if her valuation is v'_i. Thus we conclude that

$$x_i(b'_i, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \geq \min\{y_i, Z(\mathbf{b}(v'_i, \mathbf{v}_{-i}))\}.$$

The right hand side term of the above inequality is at least $\min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))\}$, agent i's part of the optimal allocation of items left unallocated, up to a maximum of $z(v_i)$. We conclude that

$$\begin{aligned}
\mathbb{E}_{\mathbf{v}}[u_i(\mathbf{v})] &\geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))\}) \\
&\qquad - p \times \min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))\}] \\
&\geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))) - p \times \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))] \\
&= \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}')))) - p \times \mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}')))]
\end{aligned}$$

where the second inequality follows because dropping the minimum with $z(v_i)$ can only introduce items with negative marginal utility at price p, and the final equality is a change of variables.

Summing over all agents, we have

$$\begin{aligned}
\sum_i \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{v})] &\geq \sum_i \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}'))))) - p \times \sum_i \mathbb{E}_{\mathbf{v},\mathbf{v}'}[\mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}')))] \\
&\geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\frac{Z(\mathbf{b}(\mathbf{v}'))}{m} \times \mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] - p \times \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \\
&= \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \times \mathbb{E}_{\mathbf{v}}\left[\frac{1}{m}\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] - p \times \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \\
&= p \times \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]
\end{aligned}$$

On the other hand, the expected revenue raised is at least p times the expected number of items sold in the auction, due to the reserve price. That is, the revenue is at least

$$p \times (m - \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]).$$

As in the proof of Theorem 3, the total welfare of the combined mechanism is at least the sum of buyer surplus and the revenue generated by the auction. Thus, summing up expected buyer surplus and expected revenue, we have that the total welfare generated at equilibrium is at least

$$p \cdot \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))] + p(m - \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]) = p \cdot m = \frac{1}{2} \cdot \mathbb{E}_{\mathbf{v}}[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))]$$

as required. □

D UNIFORM-PRICE AUCTION WITH IMPRECISE RESERVES

We now state and prove a generalized version of Theorem 4 in which the reserve price used in the uniform-price auction may differ from the suggested value, proving that the equilibrium welfare guarantee degrades gracefully with the error in price. In the special case that the price is exactly the suggested value, we get Theorem 4 as a special case.

Theorem 5. *For a multi-unit allocation problem with non-increasing marginal values, consider a uniform-price auction subject to a per-unit reserve price of p' such that $|p' - p| < \epsilon$ where $p = \frac{1}{2m}\mathbf{v}(\mathrm{OPT}(\mathbf{v}))$. Then for any trade mechanism, and at any equilibrium of the combined market, the expected welfare is at least $\frac{1}{2}E[v(\mathrm{OPT}(v))] - m\epsilon$.*

Proof. Fix some Bayes-Nash equilibrium of the combined mechanism. Let $\mathbf{b}(\mathbf{v})$ be the bids made in the auction at this equilibrium and $\mathrm{OPT}(\mathbf{v}; k)$ be the welfare-optimal allocation of k units when the valuation profile is $\mathbf{v}$. We denote the allocation of the auction under bids $\mathbf{b}$ as $\mathbf{x}(\mathbf{b})$, and the number of items left unallocated by the auction when agents bid according to $\mathbf{b}$ as $Z(\mathbf{b})$. That is, $Z(\mathbf{b}) = m - \sum_i x_i(\mathbf{b})$.

We first bound the buyer surplus at equilibrium. To this end, fix a valuation profile $\mathbf{v}$ and consider a possible deviation by buyer i. Sample phantom valuations $\mathbf{v}' \sim F$. Let $z(v_i)$ be the largest index j

such that $v_{ij} \geq p$. Note that due to the reserve price, agent i would obtain negative marginal utility for any items won in excess of $z(v_i)$. Let $y_i = \min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); m)\}$ be agent i's optimal allocation when others' valuations are $\mathbf{v}'_{-i}$, excluding any items for which her marginal value is less than p'. Our proposed deviation for agent i is to place an auction bid of b'_i where $b'_{ij} = p'$ for $j \leq y_i$, and $b_{ij} = 0$ for $j > y_i$, then not participate in the secondary market. Importantly, this deviation depends on v_i but not v_{-i}. Note also that the utility obtained under this deviation can only be non-negative since $y_i \leq z(v_i)$ and the price paid per item obtained is exactly p'.

Under this deviation, either agent i wins y_i items or all items are sold, and in the latter case agent i receives all items not allocated to the other agents under bid profile $b'_i(v'_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})$. Note that this quantity should be at least the number of items such that the bids of agents from $-i$ are strictly below p', which again is at least $Z(b_i(v'_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))$, the number of items unallocated if we also include agent i bidding at equilibrium as if her valuation is v'_i. Thus we conclude that

$$x_i(b'_i, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \geq \min\{y_i, Z(\mathbf{b}(v'_i, \mathbf{v}_{-i}))\}.$$

The right hand side term of the above inequality is at least $\min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))\}$, agent i's part of the optimal allocation of items left unallocated, up to a maximum of $z(v_i)$. We conclude that

$$\begin{aligned}
\mathbb{E}_{\mathbf{v}}[u_i(\mathbf{v})] &\geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))\}) \\
&\qquad\qquad - p' \times \min\{z(v_i), \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))\}] \\
&\geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))) - p' \times \mathrm{OPT}_i((v_i, \mathbf{v}'_{-i}); Z(\mathbf{b}(v'_i, \mathbf{v}_{-i})))] \\
&= \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}')))) - p' \times \mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}')))]
\end{aligned}$$

where the second inequality follows because dropping the minimum with $z(v_i)$ can only introduce items with negative marginal utility at price p, and the final equality is a change of variables.

Summing over all agents, we have

$$\begin{aligned}
\sum_i \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{v})] &\geq \sum_i \mathbb{E}_{\mathbf{v},\mathbf{v}'}[v_i(\mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}'))))) - p' \times \sum_i \mathbb{E}_{\mathbf{v},\mathbf{v}'}[\mathrm{OPT}_i(\mathbf{v}; Z(\mathbf{b}(\mathbf{v}')))] \\
&\geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\frac{Z(\mathbf{b}(\mathbf{v}'))}{m} \times \mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] - p' \times \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \\
&= \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \times \mathbb{E}_{\mathbf{v}}\left[\frac{1}{m}\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] - p' \times \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \\
&\geq \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \times \mathbb{E}_{\mathbf{v}}\left[\frac{1}{m}\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] - (p + \epsilon) \times \mathbb{E}_{\mathbf{v}'}[Z(\mathbf{b}(\mathbf{v}'))] \\
&= (p - \epsilon) \times \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]
\end{aligned}$$

where the last inequality uses the assumption that $|p' - p| < \epsilon$, and the final equality uses the definition of p.

On the other hand, the expected revenue raised is at least p' times the expected number of items sold in the auction, due to the reserve price. That is, the revenue is at least

$$p' \times (m - \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]) \geq (p - \epsilon) \times (m - \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]).$$

As in the proof of Theorem 3, the total welfare of the combined mechanism is at least the sum of buyer surplus and the revenue generated by the auction. Thus, summing up expected buyer surplus and expected revenue, we have that the total welfare generated at equilibrium is at least

$$(p - \epsilon) \cdot \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))] + (p - \epsilon)(m - \mathbb{E}_{\mathbf{v}}[Z(\mathbf{b}(\mathbf{v}))]) = (p - \epsilon) \cdot m = \frac{1}{2} \cdot \mathbb{E}_{\mathbf{v}}[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))] - m\epsilon$$

as required. □