Reducing Inefficiency in Carbon Auctions with Imperfect Competition

Kira Goldner 💿

Department of Computer Science, Columbia University, New York, NY, USA http://www.kiragoldner.com/kgoldner@cs.columbia.edu

Nicole Immorlica

Microsoft Research, New York, NY, USA http://www.immorlica.com/nicimm@microsoft.com

Brendan Lucier

Microsoft Research, Cambridge, MA, USA https://www.microsoft.com/en-us/research/people/brlucier/brlucier@microsoft.com

— Abstract -

We study auctions for carbon licenses, a policy tool used to control the social cost of pollution. Each identical license grants the right to produce a unit of pollution. Each buyer (i.e., firm that pollutes during the manufacturing process) enjoys a decreasing marginal value for licenses, but society suffers an increasing marginal cost for each license distributed. The seller (i.e., the government) can choose a number of licenses to put up for auction, and wishes to maximize the societal welfare: the total economic value of the buyers minus the social cost. Motivated by emission license markets deployed in practice, we focus on uniform price auctions with a price floor and/or price ceiling. The seller has distributional information about the market, and their goal is to tune the auction parameters to maximize expected welfare. The target benchmark is the maximum expected welfare achievable by any such auction under truth-telling behavior. Unfortunately, the uniform price auction is not truthful, and strategic behavior can significantly reduce (even below zero) the welfare of a given auction configuration.

We describe a subclass of "safe-price" auctions for which the welfare at any Bayes-Nash equilibrium will approximate the welfare under truth-telling behavior. We then show that the better of a safe-price auction, or a truthful auction that allocates licenses to only a single buyer, will approximate the target benchmark. In particular, we show how to choose a number of licenses and a price floor so that the worst-case welfare, at any equilibrium, is a constant approximation to the best achievable welfare under truth-telling after excluding the welfare contribution of a single buyer.

2012 ACM Subject Classification Theory of computation \rightarrow Algorithmic game theory and mechanism design; Theory of computation \rightarrow Computational pricing and auctions

Keywords and phrases welfare, price of anarchy, mechanism design, equilibrium, costs

Digital Object Identifier 10.4230/LIPIcs.ITCS.2020.15

Related Version Full version: http://www.kiragoldner.com/papers/carbonauctions.pdf

Funding *Kira Goldner*: Supported in part by NSF CCF-1420381 and by a Microsoft Research PhD Fellowship. Supported in part by NSF award DMS-1903037 and a Columbia Data Science Institute postdoctoral fellowship.

Nicole Immorlica: Supported in part by NSF Award 1841550, "EAGER: Algorithmic Approaches for Developing Markets."

© Mira Goldner, Nicole Immorlica, and Brendan Lucier; licensed under Creative Commons License CC - BY

11th Innovations in Theoretical Computer Science Conference (ITCS 2020). Editor: Thomas Vidick; Article No. 15; pp. 15:1–15:21

Leibniz International Proceedings in Informatics

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

1 Introduction

Licenses for carbon and other emissions are a market-based policy tool for reducing pollution and mitigating the effects of climate change. Roughly speaking, a government agency distributes pollution licenses to firms according to some mechanism. At the end of a period of time (e.g., a year), firms must submit licenses to cover their pollution or else face severe penalties. Different ways of distributing licenses are possible. For example, if licenses are simply sold at a fixed price to anyone who wishes to pay, then this is equivalent to a *carbon tax* where polluters must pay a linear fee to offset their emissions. Alternatively, in a *cap-and-trade* mechanism, the agency releases a fixed number C of pollution licenses, either via a predetermined allocation at no price (e.g., based on prior pollution or industry averages) or via auction, and then polluters can trade these licenses on the open market. Many emission trading systems take features of both schemes, combining a cap-and-trade market with a price floor p, where each license must be sold above the reserve price of p, and/or a price ceiling p, where an extra license beyond the cap can always be purchased at a price—of p. In general, such a pollution license market is referred to as an Emission Trading System (ETS).

There are many important ETS in effect today. The EU ETS has thousands of participating firms and has raised billions of dollars in auction revenue over the past 10 years [2]. The Western Climate Initiative runs a license auction that is linked between California and Quebec [4, 1], and the Regional Greenhouse Gas Initiative (RGGI, pronounced "Reggie") serves New England and the New York region [3]. These markets differ in the details of their implementations, but in each case, licenses are distributed on a regular schedule, with a significant quantity of those licenses sold at auction.

One can model such a license auction as a multi-unit auction; that is, an auction for multiple identical goods. There is no bound on the supply of goods, but there is a cost of production corresponding to the social cost of more pollution. The social cost is typically assumed to be increasing and convex, and the value for licenses for each buyer is commonly assumed to be increasing and concave [27]. Each of the ETS described above uses a uniform-price auction rule to resolve this multi-unit auction. Such auctions proceed roughly as follows. Participating firms declare bids, which take the form of a non-decreasing concave function that describes their willingness-to-pay for varying quantities of licenses. This can alternatively be viewed as a list of non-increasing marginal bids for each successive license. Any marginal bids below the reserve price p are removed, and licenses are then distributed to the highest remaining marginal bids while supplies last. Each firm then pays a fixed price p per license, where p is set to some value between the lowest marginal winning bid and the highest marginal losing bid. Such auctions are not truthful, but are common in practice due to their many advantageous properties; see, e.g., Chapter 7 of [22].

Implementing a uniform-price auction for carbon licenses presents an optimization challenge: good outcomes require that the system designer correctly sets the quantity of licenses to distribute and/or the price (either direct price or auction reserve) at which they will be sold. The goal is to maximize *social welfare*: the aggregate value that firms receive for their licenses (i.e., by producing goods) minus the externality on society caused by polluting the corresponding amount during production. This optimization problem is complicated by uncertainty. Even if the social cost of pollution is fully known, the designer may not know what the demand for licenses will be, which makes it hard to predict the optimal level of pollution to allow. The goal of the designer, then, is to set the parameters of the uniform price auction to maximize the expected outcome over uncertainty in the market. Notably, the presence of social costs significantly increases the complexity of multi-unit auctions,

since efficiency becomes a mixed-sign objective. Even a slight misallocation of the licenses, resulting in a small reduction of received value, might have a disproportionately large effect on net welfare. In practice, inefficient allocations may be partially resolved in the trade phase of cap-and-trade systems. However, there are significant trading frictions, so it is imperative to choose an initial allocation, via the auction, that is as efficient as possible.

Since uniform price auctions are not truthful, one must account for strategic behavior of participating firms. A notable feature of ETS markets in practice is that, despite their size, they commonly have a small number of participants with significant market power.¹ The presence of dominant players suggests imperfect competition and raises the issue of strategic manipulation, in which large individual firms try to influence the market in their favor. One particular concern is demand reduction, where firms reduce their bids to suppress the price determined by the auction. We ask: how should one set auction parameters to (approximately) maximize welfare, in the face of strategic bidding?

1.1 Our Results

In this paper we apply ideas from theoretical computer science to approach the design of carbon license auctions with strategic firms. In this new domain, we use facts regarding uniform price auctions to understand what is happening in these markets and probabilistic analysis to handle the uncertainty of valuation realizations. We then apply a Price of Anarchy analysis to give strategic guarantees. As a result, we provide a concrete recommendation for how to set the parameters of the mechanism used in practice – a recommendation with provable approximation guarantees.

A market instance is described by a convex social cost curve and, for each participating firm, an independent distribution over concave valuations. To capture the relevant space of auctions, we formalize a class of allocation rules that we call *cap-and-price auctions*.² These auctions are parameterized by a maximum number of licenses to allocate, a price floor, and a price ceiling. They proceed by running a uniform-price auction for the given number of licenses, subject to the price floor and ceiling. Given a market instance, there is some choice of cap-and-price auction that maximizes the expected welfare generated under non-strategic (truth-telling) behavior. This optimal non-strategic solution will serve as a benchmark, which we denote opt. Our main result is an auction that approximates opt at any equilibrium of strategic behavior.

We show that, in general, the welfare obtained at an equilibrium of a cap-and-price auction can be negative, even for the auction that optimizes welfare under non-strategic bidding. This motivates our main question: given a cap-and-price mechanism that achieves a certain expected welfare W under truth-telling behavior, can one modify the auction parameters so that the expected welfare is an O(1) approximation to W at any Bayes-Nash equilibrium? To answer this question, we describe a subclass of cap-and-price auctions that use $safe\ prices$. A $safe\ price\ auction$ is a cap-and-price auction in which the price floor is at least the average social cost, per license, of selling all of that auction's licenses. We show that for any such auction the worst-case welfare at any Bayes-Nash equilibrium is within a constant factor of the welfare under truth-telling. This result proceeds by transforming the market instance into a corresponding instance without social costs, and makes use of bounds on the Bayes-Nash price of anarchy for uniform-price auctions [20, 13].

¹ For example, in the EU ETS, the top 10 firms together control approximately 30 percent of all licenses allocated and traded; see [9], page 127.

Note, we are focused on the auction phase in this paper and do not model the trade phase of these systems.

To construct an auction that approximates opt at equilibrium, it therefore suffices to find a safe-price auction whose non-strategic welfare approximates that of the best unconstrained cap-and-price auction. However, as it turns out, there are market instances for which no constant approximation by a safe-price auction is possible. Such a situation can occur if there is a single firm that drives most of the demand for licenses, and this demand has very high variance. We show that this is essentially the only barrier to our desired price of anarchy result: one can always construct an auction that achieves a constant approximation to opt minus the welfare obtainable by allocating licenses to any one single auction participant. In other words, as long as no single firm accounts for most of the expected welfare of the carbon license market, one can find a cap and price floor so that the welfare at any equilibrium is within a constant of the optimal outcome achievable without strategic behavior. We argue that this assumption is reasonable, since presumably the motivation for running a carbon license market in the first place is that the demand for emission licenses is distributed across multiple firms. An alternative way to view this result is that one can always achieve a constant approximation to our benchmark by either running a safe-price auction, or by selling to only a single firm.

Along the way, we also show that for any cap-and-price auction that sells all of its licenses with probability q > 0, there is a safe-price auction whose welfare (under truth-telling) is within a factor O(1/q) of the original auction's welfare. In particular, if our benchmark is implemented by a cap-and-price auction that distributes all of its licenses with constant probability, then we can construct a safe-price auction whose welfare at any Bayes-Nash equilibrium is a constant approximation to the benchmark.

Note that our results are applicable for any setting with imperfect competition where a uniform price auction is used to allocate goods with a production or social cost. For instance, they may apply to allocating queries to a database, where the social cost is privacy. Another example might be allocating medallions to taxi and rideshare drivers in New York City, where the social cost is congestion.

1.2 Related Work

Our work is related to rich line of literature in economics comparing emission licence auction formats and flat "carbon tax" pricing methods to control the externalities of pollution. Weitzman [27] proposed a model of demand uncertainty and initiated a study comparing price-based vs quantity-based screening in the context of pollution externalities. Kerr and Cramton [10] noted that auctions tend to generate more efficient outcomes than grandfathered contracts (i.e., pre-determined allocations based on prior usage), which distort incentives to reduce pollution and efficiently redistribute licenses. Cramton, McKay, Ockenfels and Stoft [11] subsequently lay out arguments in favor of tax-based approaches. Murray, Newell, and Pizer [23] analyze the use of price ceilings in emission license auctions, and argue that they provide benefits of both auction-based and tax-based systems, improving efficiency in dynamic markets with intertemporal arbitrage. For a recent overview of auction-based systems used in practice, from both the economic and legal perspectives, we recommend [9].

A similar line of literature studied alternative approaches to the related electricity markets, where individual providers sell electricity into a central grid. Whereas the main issue in emission license sales is the social externality of production, the main focus in electricity markets was incentivizing participation of small firms. The primary discussion focused on using uniform pricing versus discriminatory pricing in the resulting procurement auction [12, 24].

Within the theoretical computer science community, there have been numerous studies of the efficiency of auction formats at equilibrium. For multi-unit auctions without production costs (or, equivalently, no social cost of pollution), the price of anarchy of the uniform-price auction was shown to be constant for full-information settings, and this was subsequently extended to Bayes-Nash equilibria [20, 13]. Our work can be seen as an extension of that work to the setting with a convex cost of allocation. Auctions with production costs have also been studied [6, 15], but primarily from the perspective of mechanism design, where the goal is to develop allocation and payment rules to achieve efficient outcomes. In the present work we do not take a mechanism design approach; we instead restrict our attention to (non-truthful) uniform price auctions, as these are used in practice, and study bounds on the worst-case welfare at equilibrium under different choices of the auction parameters.

Uniform-price auctions with costs can be viewed as games for which the designer has a mixed-sign objective (i.e., total value generated minus the social cost of production). Prior work on the price of anarchy under mixed-sign objectives has focused primarily on routing games [25, 7, 8]. Results in this space tend to be negative, motivating alternative measures of performance (such as minimizing a transformed measure of total cost) that avoid the pitfalls of mixed-sign optimization. Our work shows that in an auction setting, use of an appropriately-chosen thresholding rule (in the form of a reserve price) can enable a constant approximation to our mixed-sign objective of total value minus a convex social cost.

There is a vast amount of work on auctions with externalities, such as the seminal work of [16], and other work including externalities in advertising auctions [14], characterizations of equilibria in auctions with externalities [19], and more. However, these externalities are private and held by the buyers, as opposed to public and suffered by the seller as in this paper. In this case, the externality functions more like a production cost, as described above.

Kesselheim, Kleinberg and Tardos [17] study the price of anarchy of an energy market auction, which is a similar application to the carbon license auction that we study. Their focus is on the uncertainty of the supply and temporal nature of the auction, and not on externalities of production, and hence their technical model is quite different.

The theory of bidding in uniform-price auctions with imperfect competition is well-developed in the economic theory literature. Much of this work focuses specifically on the efficiency and revenue impact of demand reduction, and how modifications to the auction format or context might impact it. Demand reduction occurs at equilibrium even in very simple settings of full information, can dramatically reduce welfare, and is a concern in practice [28, 26, 5]. One way to reduce the inefficiency of demand reduction is to perturb the supply, either by allowing the seller to adjust the supply after bids are received [21], or by randomizing the total quantity of goods for sale (or otherwise smoothing out the allocation function) [18]. Such results are typically restricted to full-information settings. Moreover, these approaches are not necessarily appropriate in the sale of government-issued licenses, where one typically expects commitment and certainty about the quantity being sold. In contrast, we forego a precise equilibrium analysis and instead argue that setting a sufficiently high price floor can likewise mitigate the impact of demand reduction.

1.3 Roadmap

Our main result is a cap-and-price auction whose welfare in equilibrium approximates the maximum welfare of a cap-and-price auction with non-strategic reporting in markets without a single dominant bidder. In Section 2, we introduce our model and formally define cap-and-price auctions. In Section 3, we first show that we can restrict attention to cap-and-price auctions with an infinite price ceiling (i.e., ones who never sell more licenses than the cap).

We then derive a class of such cap-and-price auctions, which we call safe-price auctions, whose performance in equilibrium approximates the performance in a non-strategic setting. Thus it suffices to show that the best safe-price auction approximates the best cap-and-price auction in a non-strategic setting. However, this is not true: we give an example where it can be unboundedly worse. In Section 4, we demonstrate that the only barrier to this is the existence of a dominant bidder, yielding our main result.

2 Preliminaries

There are n firms seeking to purchase carbon licenses. Licenses are identical, and each permits one pollution unit. Each firm i has a monotone non-decreasing concave valuation curve $V_i()$ that maps a number of pollution units $x \in Z_{\geq 0}$ to a value $R_{\geq 0}$; this is the value they enjoy for polluting this amount. The profile of valuation curves $\mathbf{V} = (V_1, \ldots, V_n)$ is drawn from a publicly-known distribution F over profiles, where the draw of $V_i()$ is the private information of firm i. Valuations are drawn independently across firms, so $F = F_1 \times \ldots \times F_n$ is a product distribution and $V_i \square F_i$.

There is a publicly known monotone non-decreasing convex *cost function Q*() that maps a number of pollution units $x \in Z_{\geq 0}$ to the value of externality that society faces for having this much pollution. Given a valuation profile (V_1, \ldots, V_n) , the *welfare* of a given allocation rule $\mathbf{x} = (x_1, \ldots, x_n)$ is $V_i(x_i) \square Q(V_{i=1}^n x_i)$. Our objective is to maximize expected welfare.

Given some integer $x \ge 0$, we write V(x) for the maximum aggregate value that could be obtained by optimally dividing x licenses among the firms. That is,

$$V(x) = \max_{y \in \mathbb{Z}_+^n: ||y||_1 = x} X_i(y_i).$$

We refer to V as the *combined valuation curve*. As each $V_i()$ is weakly concave, so is V(). We will sometimes abuse notation and write $V \square F$ to mean that V is distributed as the aggregate value when $(V_1, \ldots, V_n) \square F$. We'll write $W(V, x) = V(x) \square Q(x)$ for the welfare generated by allocating x licenses optimally among the firms.

We'll write $v_i(j) = V_i(j) \square V_i(j \square 1)$ for firm i's marginal value for aquiring license j, for each $j \ge 1$. By the monotonicity and concavity of $V_i()$, $v_i()$ is non-increasing. We'll also write $V_i(j|k) = V_i(j+k) \square V_i(k)$ for the marginal value of j additional items given k already allocated. We'll write $d_i(p) = \max\{j : v_i(j) \ge p\}$ for the number of units demanded by bidder i at price p.

We will study equilibria and outcomes of license allocation auctions. An auction takes as input a reported valuation V_i from each firm i. The auction then determines an allocation $\mathbf{x} = (x_1, \ldots, x_n)$ and a price $p_i \geq 0$ that each firm must pay. The auctions we consider will be uniform price auctions, where the auction determines a per-license price p and each firm i pays $p_i = p$ x_i . Given an implicit uniform-price auction, we will tend to write $x_i(\mathbf{V})$ (resp., $p(\mathbf{V})$) for the allocation to agent i (resp., per-unit price) when agents report according to \mathbf{V} . For a given valuation profile \mathbf{V} , we'll also write $U_i(\mathbf{V})$ for the utility enjoyed by firm i when agents bid according to \mathbf{V} :

$$U_i(\tilde{\mathbf{V}}) = V_i(x_i(\tilde{\mathbf{V}})) \, \Box p(\tilde{\mathbf{V}}) \, x_i(\tilde{\mathbf{V}}).$$

Finally, given an auction M and a distribution F over input profiles, we will write W(M, F) for the expected welfare of M on input distribution F. We will sometimes drop the dependence of F and simply write W(M) when F is clear from context. We emphasize that W(M) is a

non-strategic notion of welfare; it is the expected welfare of M with respect to inputs drawn from F, or equivalently the welfare of M under truthful reporting when agent valuations are distributed according to F. Let $\mathbf{x}^{M}(V) = (x_{1}^{M}(V), \dots, x_{n}^{M}(V))$ be the allocation rule of mechanism M for realization V. Then

$$W(M) = E_{V \square F} \qquad V_i \quad x_i^M(V) \quad \square Q \qquad x_i^M(V) \quad .$$

A Bayes-Nash equilibrium of a uniform-price auction is a choice of bidding strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ for each agent, mapping each realized valuation curve V_i to a (possibly randomized) reported valuation $\sigma_i(V_i)$, so that each agent maximizes their expected utility by bidding according to σ_i given that other agents are bidding according to σ_{\square} . That is, for all V_i and all \tilde{V}_i , we have

$$\mathsf{E}_{V_{\square i}\square F_{\square i}}[U_i(\sigma_i(V_i),\,\sigma_{\square i}(V_{\square i}))] \geq \; \mathsf{E}_{V_{\square i}\square F_{\square i}}[U_i(V_i,\,\sigma_{\square i}(V_{\square i}))].$$

We make a standard assumption of no overbidding on behalf of the firms uniform-price auction, which is motivated by the fact that overbidding is a weakly-dominated strategy.

Motivated by license auctions used in practice, we study a particular form of uniform price auction that we call a cap-and-price auction which forces prices to be within some fixed interval.3

- **Definition 1.** A cap-and-price auction $M(C, p, \overline{p})$ is parameterized by a quantity $C \ge 1$, a price floor p, and a price ceiling $\bar{p} > p$. The allocation and price are determined as follows:
- **1.** If $p \mid d_i(p) \ge C$, then $x_i = d_i(p)$ for all i and p = p. **2.** If $p \mid d_i(p) < C$, then $x_i = d_i(p)$ for all i and p = p.
- $_{i}d_{i}(\underline{p}) < C$, then $x_{i} = d_{i}(\underline{p})$ for all i and $p = \underline{p}$.
- **3.** Otherwise, choose $p = V(C) \square V(C \square 1)$, the C^{th} highest bid, and choose χ to be any optimal allocation of C licenses among the firms: $x \in \arg\max_{y \in \mathbb{Z}_i^n : ||y||_1 = C} \int_i^n V_i(y_i)$.

We can think of this as a uniform price auction of up to C licenses (case 3 above), where the price is set to the lowest winning marginal bid, with two modifications. First, there is a reserve price p, so that possibly fewer than C licenses are sold if there is not enough demand at price p; this is case 2. Second, if the lowest winning marginal bid would be larger than \bar{p} , then the price is lowered to \bar{p} and firms can purchase as many licenses as they like at this price: this is case 1.

Given social cost function Q and valuation distribution F, we will write opt for the optimal expected welfare obtained by any cap-and-price auction under truthful reporting. That is, opt = $\max_{C,p,\overline{p}} \{W(M(C,p,\overline{p}))\}$. We will also tend to write C^{\square} , p^{\square} , and \overline{p}^{\square} for the parameters that achieve this maximum. We emphasize that this is a non-strategic notion: opt is the maximum expected welfare attainable when bidders report truthfully. We view opt as a benchmark against which we will compare performance at Bayes-Nash equilibrium. Note also that, in general, opt may be strictly less than the expected welfare of the unconstrained welfare-optimal allocation; an example is provided in Appendix С.

Safe-Price Auctions 3

We will derive a class of cap-and-price auctions, which we will call safe-price auctions, whose performance in equilibrium approximates their non-strategic welfare. The hope is that the non-strategic welfare of this class then approximates the non-strategic welfare of the larger class of cap-and-price auctions.

³ Note the cap is not a hard constraint, but rather governs which prices bind.

Definition 2. A safe-price auction is a cap-and-price auction M(C, p, p) parameterized by a license quantity C, a price floor p = Q(C)/C, and price ceiling $p = \infty$. We will write M(C) for the safe-price auction with license quantity C, and P(C) = Q(C)/C for the associated safe price.

This definition restricts cap-and-price auctions in two ways. The first is that the price ceiling is now infinite. The second is that we have imposed a lower bound on the price floor.⁴ The first restriction is for convenience; the following lemma shows that it does not cause much loss in welfare. Motivated by this lemma, we will restrict our attention from now on to auctions with no price ceiling, that is, $\bar{p} = \infty$. We notate such mechanisms M = (C, p). Note that, in the language of definition 1, case (1) can no longer occur. That is, when is there is no price ceiling, only exactly C or less than C licenses will be allocated.

Lemma 3. For any distribution F and any cap-and-price auction $M(C, \underline{p}, \overline{p})$, there exist a cap C^0 and price floor p^0 such that $W(M(C^0, p^0, \infty)) \ge \frac{1}{2}W(M(C, p, \overline{p}))$.

The proof of Lemma 3 appears in Appendix A.1. The main idea is to decompose the expected welfare of a cap-and-price auction $M(C, \underline{p}, \overline{p})$ into the welfare attained from the first (at most) C licenses sold plus the incremental welfare attained from any licenses sold after the first C. The first can be approximated by $M(C, \underline{p}, \infty)$; the latter, if non-negative, can be approximated by $M(\infty, \overline{p}, \infty)$.

The second way safe-price auctions restrict price-and-cap auctions is with a lower bound on the price floor. As the following example shows, the equilibrium welfare of cap-and-price auctions, even those with an infinite price ceiling, suffer from a problem known as demand reduction in which a bidder improves her price, and hence utility, by asking for fewer units. This can have drastic consequences on welfare, causing it to become negative, due to the existence of the social cost. In particular, this means the approximation of such auctions, relative to the best welfare attainable in non-strategic settings, is unbounded.

Example 4. Consider two agents, 1 and 2. Their distributions over valuations will be point-masses, so that the valuations are actually deterministic. These valuation functions are given by the following marginals: $v_1(1) = 10$, $v_1(2) = 10$, $v_2(1) = 6$, $v_2(2) = 1$, and all other marginals are 0. We will have C = 2, p = 0, $p = \infty$, and the social cost function is given by Q(x) = 9x. Under truthful reporting, the auction M(C, p) allocates 2 licenses to agent 1 at a price of 6 each, resulting in a welfare of $V_1(2) \square Q(2) = 20 \square 18 = 2$. However, we note that firm 1 could improve their utility from 8 to 9 by instead reporting a modified valuation \tilde{V} given by $\tilde{V}_1(1) = 10$ and $\tilde{V}_1(2) = 1$. If agent 2 continues to report truthfully, M(C, p) will allocate 1 license to each agent at a price of 1 each, resulting in a welfare of $V_1(1) + V_2(1) \square Q(2) = 16 \square 18 = \square 2$. One can verify that this is indeed a pure Nash equilibrium of the auction, and hence a Bayes-Nash equilibrium as well.

The issue illustrated in Example 4 is that strategic behavior can cause licenses to be allocated to agents whose marginal values for these licenses are below Q(C)/C, causing the aggregate value derived from allocating C licenses to fall below Q(C). Safe-price auctions seek to prevent this by imposing a sufficiently high price floor.

Indeed, the following lemma shows that price floors do indeed circumvent the issue in the above example. Namely, we show safe-price auctions have good equilibria, compared to their own non-strategic welfare. This is the main result of this section, and motivates us to focus on analyzing the non-strategic welfare of safe-price auctions.

⁴ In fact we set p to be equal to P(C), but our results still hold if this is relaxed to $p \ge P(C)$.

Lemma 5. For any safe-price auction M(C), the expected welfare at any Bayes-Nash equilibrium is at least $(\frac{1}{3 \cdot 15})$ W(M(C)).

Proof. Fix a license cap C and write $\underline{p} = Q(C)/C$ for the corresponding safe price, and recall that $M(C) = M(C, \underline{p})$. We claim that this auction is strategically equivalent to a modified auction, as follows. First, since each agent pays at least \underline{p} per license, we can define $\hat{v}_i(j) = v_i(j) \Box p$ for the residual marginal utility of agent i for their j^{th} license, after taking into account that they must pay at least \underline{p} . Then from each agent's perspective, playing in M(C, p) with marginal values $v_i(j)$ gives the same utility as playing in M(C, p) with marginal values $\hat{v}_i(j)$. The welfare generated by the original auction is $V(x) \Box Q(x)$, which is equal to $V(\hat{x}) \Box (Q(x) \Box px)$. So the welfare generated by M is equivalent to the welfare generated with modified valuations \hat{V} , where the marginal social cost of each unit of pollution is reduced by p. But now note that from the definition of p, $Q(x) \Box px \leq 0$ for all $x \leq C$. So in an auction with a cap of C, the welfare is at least the welfare obtained with a social cost of 0; it is always non-negative.

We conclude that the welfare generated by $M(C, \underline{p})$, at any equilibrium, is equal to the welfare generated at an equilibrium of the standard uniform price auction with modified valuations \hat{V} . Theorem 3 of [13] (which bounds the welfare in Bayes-Nash equilibria of uniform auctions without costs) therefore applies⁵, and the welfare at any equilibrium is at least $\frac{1}{3.1}$ times the optimum achievable under the modified valuations, which is at most W(M(C)).

4 Welfare Approximation

Motivated by Lemma 5, we would like to show that the non-strategic welfare of safe-price auctions approximates that of cap-and-price auctions. Unfortunately, there are cases where this does not hold, as demonstrated by Example 6. The example consists of just a single firm that dominates the market and whose demand for licenses has high variance. The firm's value is always large enough that even at a low price and high license cap, the resulting allocation has high net welfare. However, safe-price auctions will generate much lower welfare. Roughly speaking, for any number of licenses C that the auction designer selects as the cap for a safe-price auction, significant welfare will be lost from realizations where the firm's demand is much higher than C, and the corresponding safe price will exclude welfare gains from realizations where the firm's demand is much lower than C. If the variance is high enough, this will cause overall welfare to be low regardless of the choice of C.

Example 6. We will present an example for which the expected welfare of any safe-price auction M(C) is at most an O(1/n)-approximation to opt. Let $Q(x) = x^2$. There is a single firm participating in the auction. That firm's valuation curve is drawn according to a distribution F over n different valuation curves, which we'll denote $V^{(1)}, \ldots, V^{(n)}$. For all i, valuation $V^{(i)}$ is defined by $V^{(i)}(x) = \max\{2^{i+1} \ x, 2^{2^{i+1}}\}$. The probability that $V^{(i)}()$ is drawn from F is proportional to $\frac{1}{2^{2r}}$. That is, the firm has valuation $V^{(i)}$ with probability $\frac{1}{62^{2i}}$, where $\theta = \begin{bmatrix} n & 1 \\ i=1 \end{bmatrix} 2^{-2i} = \frac{1}{3}(1 \square 2^{-2n})$ is the normalization constant. Note that a cap-and-price auction with cap $C = \infty$ and price floor p = 1 achieves welfare $\begin{bmatrix} n & 1 \\ i=1 \end{bmatrix} 2^{2i} = \frac{1}{6}n$, so opt is at least this large.

The proof of the 3.15 bound in [13] is actually for a uniform-second-price auction that charges the highest losing bid, or $V(C+1) \square V(C)$; however, the proof is also correct for charging a price of the lowest winning bid $V(C) \square V(C \square 1)$.

⁶ In fact, this is the "first-best" welfare obtainable at every possible realization, so this is actually the best possible cap-and-price auction and hence the exact value of opt.

Consider a safe-price auction with cap C^0 . Suppose that $C^0 \in (2^k, 2^{k+1}]$ for some $k \ge 1$. Then the safe price is $P(C^0) = Q(C^0)/C^0 = (C^0)^2/C^0 = C^0 > 2^k$. If the firm has valuation $V^{(i)}$ with $i \le k \square 1$, then all marginal values are strictly below $P(C^0)$, then no licenses are allocated and the welfare generated is 0. On the other hand, if the firm has valuation $V^{(i)}$ for any $i \ge k$, then since at most $C^0 \le 2^{k+1}$ licenses can be purchased, the auction generates welfare at most $2^{i+1}(2^{k+1}) \square (2^{k+1})^2 = 2^{2k+2}(2^{i\square k} \square 1)$.

The total expected welfare of $M(C^0)$ is therefore at most

$$\sum_{i=k}^{X^{n}} 2^{2k+2} (2^{i \square k} \square 1) \qquad \frac{1}{62^{2i}} = \frac{4}{6} \sum_{i=k}^{X^{n}} 2^{\square 2(i \square k)} (2^{i \square k} \square 1) \leq \frac{4}{6} \bigcirc O(1).$$

Since opt $\geq \frac{1}{\theta}$ n but $W(M(C^0)) = O(1)$ for all C^0 , we conclude that no auction with a safe price achieves an o(n)-approximation to opt. This concludes the example.

Example 6 is driven by a single firm that dominates the market and has high variance in their demand. We next show that this is the *only* barrier to a good approximation: either a safe-price auction is constant-approximate, or else one can approximate the optimal welfare by selling to just a single firm. To this end, we will be interested in the expected maximum welfare attainable by allocating licenses to just one firm, which we will denote $W^{(1)}$. That is,

$$W^{(1)} = \mathsf{E}_{V \, \square F} \quad \max_{i, x \geq 0} \{ V_i(x) \, \square \, Q(x) \} \quad .$$

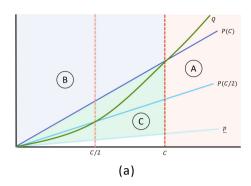
The following theorem is our main result.

Theorem 7. There exists a constant c such that, for any cap C and price floor \underline{p} , there exists C^0 such that $c W(M(C^0)) + W^{(1)} \ge W(M(C, p))$.

Note that a corollary of Theorem 7 (combined with Lemmas 3 and 5) is that for any cap-and-price auction $M(C, \underline{p}, \overline{p})$, there exists some C^0 such that the expected welfare of any Bayes-Nash equilibrium of $M(C^0)$ is at least a constant factor of $W(M(C, \underline{p}, \overline{p})) \square W^{(1)}$. That is, the worst-case equilibrium welfare is a constant-factor approximation to the best welfare achievable by any cap-and-price auction under truth-telling, excluding the welfare contribution of a single firm.

Also, $W^{(1)}$ is the expected welfare of a truthful mechanism that we will call $M^{(1)}$. In $M^{(1)}$, the firms first participate in a second-price auction for the right to buy licenses. The firm with the highest bid wins, and they pay the second-highest bid. The winning firm can then purchase any number of licenses $x \geq 0$, for which they pay Q(x) (in addition to their payment from the initial second-price auction). Since the utility obtained by firm i if they win the initial auction is precisely $\max_x \{V_i(x) \mid Q(x)\}$, and since a second-price auction is truthful and maximizes welfare, we can conclude that $W(M^{(1)}) = W^{(1)}$. Then another corollary of Theorem 7 is that one can approximate the non-strategic welfare of any cap-and-price auction using either a safe-price auction (in which case the approximation holds at any Bayes-Nash equliibrium), or using the (truthful) mechanism $M^{(1)}$ that allocates licenses to at most one firm.

We are now ready to prove Theorem 7. Fix a cost function Q and distribution F, and choose the optimal cap C and price floor p such that W(M(C, p)) is maximized. Let d() and x() be the demand and allocation functions for M(C, p). For any given realization of preferences V, there is a total demand $d(V) = \bigcap_{i=1}^{N} d_i(p)$, where again, $d_i(p) = \arg\max_x V_i(x) \Box px$. Recall also that $W(x, V) = V(x) \Box Q(x)$ is the welfare generated by allocation X given aggregate valuation V.



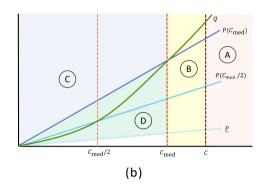


Figure 1 A visualization of the proofs of Theorem 8 (a) and Theorem 7 (b). The x-axis represents quantity of licenses; the y-axis represents value. The expected welfare of benchmark $W(M(C, \underline{p}))$ is divided into separate contributions, based on the location of allocation outcome (x, V(x)); these are depicted as shaded regions. Each of these contributions is bounded by either the welfare of a safe-price auction or the welfare obtained from a single buyer. Figure (a) is the (simpler) partition used in Theorem 8, and (b) is the partition used in Theorem 7.

We will begin by proving a simpler version of Theorem 7, which is parameterized by $q := \Pr[d(V) \ge C]$, the probability that the auction sells C licenses. When q is bounded away from 0, we show that there is a safe-price auction that achieves an O(1/q)-approximation to W(M(C, p)). We note that this result does not require firm valuations to be independent, and holds even if the valuations are drawn from an arbitrarily correlated distribution F. We will explain in Section 4.1 how to extend the argument to the general case of Theorem 7, with details deferred to the appendix.

Theorem 8. For the welfare-optimal cap C and price floor p, there exists a cap C^0 such that (1 + 2/q) $W(M(C^0)) \ge W(M(C, p))$.

Proof. The welfare generated by M(C, p) can be broken down as follows:

$$W(M(C, p)) = \Pr_{V \square F} [d(V) \ge C] E[W(x(V), V) \mid d(V) \ge C] + \Pr_{V \square F} [d(V) < C] E[W(x(V), V) \mid d(V) < C].$$

The first term is the contribution to the welfare from events where C licenses are sold (or, equivalently, at least C licenses are demanded). The second term is the contribution from events where the total demand for licenses is less than C.

where the final inequality used the fact that $Q(x) \ge Q(x^{<}) + Q(x^{>})$, due to the convexity of Q. We conclude that

We claim that the welfare obtained from the first two terms, (1A) and (1B), are covered by the safe-price auction with cap C. The intuition is that a price floor of P(C) does not interfere with the welfare contribution due to licenses with marginal values greater than P(C). One subtlety is that some of the licenses in the summation (1A) might have marginal values less than P(C), but it turns out that it can only improve welfare to exclude

so that $W(M(C, p)) \le (1A) + (1B) + (1C)$. See Figure 1(a) for an illustration.

B Claim 9. $W(M(C)) \ge (1A) + (1B)$.

such licenses from the allocation.

Proof. Write d^0 and x^0 for the demand and allocation under auction M(C), respectively. We have

$$W(M(C)) = \Pr_{V \cap F} [d(V) \ge C] \ E[W(x^{0}(V), V) \mid d(V) \ge C] +$$

$$\Pr_{V \cap F} [d(V) < C] \ E[W(x^{0}(V), V) \mid d(V) < C].$$

Note that we intentionally use d(V), the demand for M(C, p), rather than $d^0(V)$ in the expressions above. The second term is precisely (1B), since $x^0 = x^2$ from the definition of $x^{>}$. We claim that the first term is at least (1A). To see why, fix any V with $d(V) \geq C$ (and corresponding allocation x = x(V) where necessarily |x| = C, and note that $x^0 \le x_i$ for all i. Furthermore, $V_i(x_i^0) \ge V_i(x_i) \square (x_i \square x_i) \mathcal{P}(C)$, since x_i 4s simply x_i after possibly excluding some items with marginal value less than P(C). Thus, since |x| = C,

$$V(x^0) \Box x^0 | P(C) \ge V(x) \Box x | P(C) = V(x) \Box Q(|x|).$$

Also, by convexity, we have $Q(y) \le y P(C)$ for all $y \in [0, |x|]$. In particular, we have $V(x^0) \square Q(|x^0|) \ge V(x^0) \square x^0 | P(C) \ge V(x) \square Q(|x|)$, and hence

$$\Pr_{V \sqcup F}[d(V) \ge C] \ E[W(x^{0}(V), V) \mid d(V) \ge C]$$

 $\ge \Pr_{V \sqcup F}[d(V) \ge C] \ E[W(x(V), V) \mid d(V) \ge C]$

as claimed. С The more difficult part of the proof of Theorem 8 is to account for term (1C). This represents the contribution of licenses whose marginal values lie below P(C), and are therefore excluded when the price floor is set to $\underline{p} = P(C)$. What we show is that the welfare contribution from all such licenses is approximated by W(M(C/2)), the welfare obtained by the safe auction with cap C/2. We do this in three steps. First, we show that any optimal cap-and-price auction $M(C,\underline{p})$ must generate non-negative expected welfare conditional on selling C licenses. The intuition is that if the expected welfare from selling C licenses is negative, then it must be strictly preferable to reduce the quantity of licenses.

Lemma 10. If cap C and price floor p are chosen to maximize W(M(C, p)), then

$$E_{V \square F}[W(M(C, p)) \mid d(V) \geq C] \geq 0$$

where d(V) is the total quantity of licenses demanded at price p.

The proof of Lemma 10 appears in Appendix B.1. Next, for any cap-and-price auction M, we provide an upper bound on the total welfare contribution that comes from events where x < C licenses are sold, and the average marginal value of the sold licenses, V(x)/x, is less than P(C). In fact, it will be useful to state the lemma more generally: we provide a more general bound that applies to any $C \le C$ and any valuation that is at most P(C) X (not necessarily Y(X)). The lemma bounds the total welfare contribution, P(C) $X \square Q(X)$, by the difference in price between (a) C licenses sold at the safe price for C, and (b) the same C licenses sold at the safe price for C/2.

Lemma 11. Choose some quantity \tilde{C} and a number of licenses $x \leq \tilde{C}$. Then $P(\tilde{C}) \times Q(x) \leq \tilde{C}$ ($P(\tilde{C}) \square P(\tilde{C}/2)$).

The proof of Lemma 11 appears in Appendix B.2. We are now ready to bound the contribution of (1C). The intuition is as follows. By Lemma 11, the contribution from (1C) is at most (a constant times) the gap between the line y = P(C)x and the curve Q(x) at point x = C/2. But recall that $\Pr[d(V) \ge C] = q$, and Lemma 10 implies that, on average, the expected marginal value of licenses allocated subject to this event is at least P(C). Therefore, if we set a license cap of C/2, then with probability at least q all C/2 licenses will be sold at an average expected marginal value of at least P(C). The welfare generated in this case covers the "gap" at C/2, and hence covers the loss due to excluding licenses with marginal values at most P(C).

B Claim 12. $W(M(C/2)) \ge \frac{1}{2}q$ (1C).

Proof. Write d^0 and x^0 for the demand and allocation under mechanism M(C/2), respectively. Choose any V such that $d(V) \ge C$. For any such V, $d^0(V) \le d(V)$, and is formed by removing items with marginal value at most P(C/2). In particular, since $x^0(V) \le d^0(V)$, we have $V(x^0(V)) \ge V(C) \square (C \square x^0(V)) P(C/2)$. Noting that $Q(|x^0(V)|) \le |x^0(V)| P(C)$ since $x^0(V) \le x(V)$, we have

$$V(x^{0}(V)) \square Q(|x^{0}(V)|) \ge V(x^{0}(V)) \square x^{0}(V)| P(C)$$

$$\ge V(C) \square (C \square x^{0}(V)|) P(C/2) \square x^{0}(V)| P(C).$$

For convenience we will assume C is even for the remainder of this section. When C is odd, the result holds for at least one of the floor or the ceiling of C/2. Details appear in the full version.

Taking an expectation over all V such that $d(V) \ge C$, we note that Lemma 10 implies $E[V(C)] \ge C$ P(C). We therefore have

```
E[V(x^{0}(V)) \square Q(|x^{0}(V)|)] \ge C P(C) \square (C \square E[|x^{0}(V)|]) P(C/2) \square E[|x^{0}(V)|] P(C) = (C \square E[|x^{0}(V)|])(P(C) \square P(C/2))
\ge (C/2) (P(C) \square P(C/2))
```

where in the last inequality we used that $|x^0(V)| \le C/2$ by definition. Since $\Pr[d(V) \ge C] = q$, we therefore have that

```
W(M(C/2)) \ge \Pr[d(V) \ge C] \ \mathbb{E}[V(x^0(V)) \square Q(|x^0(V)|) \mid d(V)
 \ge C] \ge q(C/2) \ (P(C) \square P(C/2).
```

We now claim that (C/2) $(P(C) \square P(C/2) \ge (1C)/2$, completing the proof. To see why, note that for any $x^\square < C$ and any V, $V^<(x^\square) \square Q(|x^\square|) \le |x^\square| P(C) \square Q(|x^\square|)$. Thus $(1C) \le \max_{x^\square} \{|x^\square| P(C) \square Q(|x^\square|)\}$. But note that for any $x^\square > C/2$, we have $Q(x^\square) > |x^\square| P(C/2)$, and hence $|x^\square| P(C) \square Q(|x^\square|) \le |x^\square| (P(C) \square P(C/2)) \le (C) (P(C) \square P(C/2))$. Also, for any $x^\square < C/2$, $Q(|x^\square|)$ lies above the line joining (C/2, Q(C/2)) and (C, Q(C)), and hence $|x^\square| P(C) \square Q(|x^\square|)$ is at most the distance between $|x^\square| P(C)$ and this line, which is at most (C) $(P(C) \square P(C/2))$. So in either case we have $\max_{x^\square} \{|x^\square| P(C) \square Q(|x^\square|)\} \le C (P(C) \square P(C/2))$ as required.

Combining Claim 9 and Claim 12 we have that $W(M(C)) + (2/q)W(M(C/2)) \ge W(M(C,p))$, which completes the proof of Theorem 8.

As a corollary, Theorem 8 combined with Lemma 5 and Lemma 3 implies that for any cap and price auction $M(C, \underline{p}, \overline{p})$, there exists a safe-price auction $M(C^0)$ such that, at any Bayes-Nash equilibrium of $M(C^0)$, the expected welfare generated is at leas $\frac{1}{1.15}$ $\frac{1}{2}$ $\frac{1}{1+2/q}$ $\frac{1}{2}$ $\frac{1}{2}$ $\frac{1}{1+2/q}$ $\frac{1}{2}$ $\frac{$

4.1 The General Case

We complete the proof of Theorem 7 in Appendix B.3. Here we describe at a high level what steps are needed to complete the argument. We will focus on the case where $\Pr[d(V) \ge C] < 1 \square 1/e$, since if $\Pr[d(V) \ge C] \ge 1 \square 1/e$ then Theorem 7 follows from Theorem 8.

Recall that in Claim 12, we used the assumption that $\Pr[d(V) \geq C] = q$ to argue that the welfare gained in M(C/2) in the event that $d(V) \geq C$ covers the welfare lost from marginals lying below P(C), up to a constant factor. If the probability that $d(V) \geq C$ is small, then this may no longer be true. To handle this, we consider a reduced cap $C_{\text{med}} < C$ set to be the largest integer such that $\Pr[d(V) \geq C_{\text{med}}] \geq 1 \Box 1/e$. Our hope is to reproduce the argument from Claim 12, but substituting C_{med} for C. To this end, we divide the welfare of M(C, p) into four parts: all welfare under the event that d(V) > C; all welfare from individual agents whose demand \overline{l} is at least C_{med} ; the contribution of marginal values greater than $P(C_{\text{med}})$ (but with individual firms demanding at most C_{med}) when d(V) < C, and the contribution of marginal values less than $P(C_{\text{med}})$ when d(V) < C. See Figure 1(b) for an illustration.

As in the proof of Theorem 8, the contribution due to events where d(V) > C can be covered by the welfare of M(C, P(C)).

The contribution from agents who individually demand at least $C_{\rm med}$ licenses is a new case that we didn't have to handle in Theorem 8. It is here that we use $M^{(1)}$, allocating to any single agent. Because the *total* demand is at most than $C_{\rm med}$ with probability at least

1/e, independence implies that the expected number of agents who demand more than C_{med} licenses is at most 1, given that the probability that none have demand more C_{med} must be at least 1/e. So the total contribution to welfare of all such events is at most $W(M^{(1)})$.

If $d(V) \leq C_{\text{med}}$, then the contribution from marginal values that are at least $P(C_{\text{med}})$ can be covered by the welfare of $M(C_{\text{med}}, P(C_{\text{med}}))$, precisely as in Claim 9. We must also handle the case that $d(V) \in [C_{\text{med}}, C]$, and consider the welfare contribution of agents that do not (individually) demand more than C_{med} licenses. Here we use independence: the total quantity demanded by such "small" agents is likely to concentrate, so it is unlikely that the total demand will be larger than $2C_{\text{med}}$. Thus, by imposing a cap of C_{med} , we lose at most a constant factor of the welfare from marginal values greater than $P(C_{\text{med}})$.

The final step is to show that $W(M(C_{\rm med}/2))$ obtains at least a constant fraction of the welfare generated by marginal values less than $P(C_{\rm med})$, similarly to Claim 12. When proving Claim 12, we used Lemma 10 to argue about the welfare generated by events where d(V) > C. Unfortunately, Lemma 10 does not extend to $C_{\rm med}$: it could be that the expected welfare generated by $M(C_{\rm med}, P(C_{\rm med}))$, conditional on selling $C_{\rm med}$ licenses, is negative. However, we can prove an upper bound on how negative this expected welfare can be. After all, if the expected welfare is sufficiently negative sufficiently often, it would be welfare-improving to increase the price floor of M(C, p) from p to $P(C_{\rm med})$, contradicting the supposed optimality of M(C, p). This turns out to be enough to prove a bound similar to Claim 12.

Combining these bounds, we can conclude that each of the four parts of the welfare of $M(C, \underline{p})$ can be covered by either a safe-price auction or by $M^{(1)}$, which completes the proof of the theorem.

- References -

- 1 California Air Resources Board | Quarterly Auction Information. https://ww3.arb.ca.gov/cc/capandtrade/auction/auction.htm. Accessed: 2019-09-06.
- 2 EU Emissions Trading System (EU ETS). https://ec.europa.eu/clima/policies/ets_en. Accessed: 2019-09-06.
- 3 RGGI, Inc. https://www.rggi.org/. Accessed: 2019-09-06.
- 4 Western Climate Initiative, Inc. http://www.wci-inc.org/. Accessed: 2019-09-06.
- 5 Lawrence M. Ausubel, Peter Cramton, Marek Pycia, Marzena Rostek, and Marek Weretka. Demand Reduction and Inefficiency in Multi-Unit Auctions. *Review of Economic Studies*, 81(4):1366–1400, 2014. URL: https://EconPapers.repec.org/RePEc:oup:restud:v:81:y: 2014:i:4:p:1366-1400.
- 6 A. Blum, A. Gupta, Y. Mansour, and A. Sharma. Welfare and Profit Maximization with Production Costs. In 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, pages 77–86, October 2011. doi:10.1109/FOCS.2011.68.
- 7 Chi Kin Chau and Kwang Mong Sim. The price of anarchy for non-atomic congestion games with symmetric cost maps and elastic demands. *Operations Research Letters*, 31(5):327–334, 2003. doi:10.1016/S0167-6377(03)00030-0.
- 8 Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Bottleneck Links, Variable Demand, and the Tragedy of the Commons. *Netw.*, 60(3):194–203, October 2012. doi:10.1002/net.21458.
- 9 Deborah Cotton. 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, 2015. doi:10.1111/1467-8489.12094.
- Peter Cramton and Suzi Kerr. Tradeable carbon permit auctions: How and why to auction not grandfather. *Energy Policy*, 30(4):333–345, 2002. URL: https://EconPapers.repec.org/RePEc:eee:enepol:v:30:y:2002:i:4:p:333-345.

- Peter Cramton, David JC MacKay, Axel Ockenfels, and Steven Stoft. *Global Carbon Pricing:*The Path to Climate Cooperation. The MIT Press, June 2017. doi:10.7551/mitpress/10914.
 001.0001.
- Peter Cramton and Steven Stoft. Why We Need to Stick with Uniform-Price Auctions in Electricity Markets. *The Electricity Journal*, 20(1):26–37, 2007. URL: https://ideas.repec.org/a/eee/jelect/v20y2007i1p26-37.html.
- Bart de Keijzer, Evangelos Markakis, Guido Schäfer, and Orestis Telelis. Inefficiency of Standard Multi-unit Auctions. In Hans L. Bodlaender and Giuseppe F. Italiano, editors, *Algorithms ESA 2013*, pages 385–396, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- Arpita Ghosh and Mohammad Mahdian. Externalities in Online Advertising. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, pages 161–168, New York, NY, USA, 2008. ACM. doi:10.1145/1367497.1367520.
- 2 Zhiyi Huang and Anthony Kim. Welfare Maximization with Production Costs: A Primal Dual Approach. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 59–72, Philadelphia, PA, USA, 2015. Society for Industrial and Applied Mathematics. URL: http://dl.acm.org/citation.cfm?id=2722129.2722135.
- Philippe Jehiel, Benny Moldovanu, and Ennio Stacchetti. How (not) to sell nuclear weapons. The American Economic Review, pages 814–829, 1996.
- 17 Thomas Kesselheim, Robert Kleinberg, and Eva Tardos. Smooth Online Mechanisms: A Game-Theoretic Problem in Renewable Energy Markets. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 203–220, New York, NY, USA, 2015. ACM. doi:10.1145/2764468.2764487.
- 18 Ilan Kremer and Kjell Nyborg. Underpricing and Market Power in Uniform Price Auctions. Review of Financial Studies, 17(3):849–877, 2004. URL: https://EconPapers.repec.org/RePEc:oup:rfinst:v:17:y:2004:i:3:p:849-877.
- 19 Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.
- Evangelos Markakis and Orestis Telelis. Uniform Price Auctions: Equilibria and Efficiency. *Theory of Computing Systems*, 57(3):549–575, October 2015. doi:10.1007/s00224-014-9537-9.
- David McAdams. Adjustable supply in uniform price auctions: Non-commitment as a strategic tool. *Economics Letters*, 95(1):48–53, April 2007. URL: https://ideas.repec.org/a/eee/ecolet/v95y2007i1p48-53.html.
- Paul Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004. doi:10.1017/CBO9780511813825.
- Brian Murray, Richard Newell, and William Pizer. Balancing Cost and Emissions Certainty: An Allowance Reserve for Cap-and-Trade. *Review of Environmental Economics and Policy*, 3(1):84–103, 2009. URL: https://EconPapers.repec.org/RePEc:oup:renvpo:v:3:y:2009: i:1:p:84-103.
- Stephen J Rassenti, Vernon L Smith, and Bart J Wilson. Discriminatory Price Auctions in Electricity Markets: Low Volatility at the Expense of High Price Levels. *Journal of Regulatory Economics*, 23(2):109–123, March 2003. URL: https://ideas.repec.org/a/kap/regeco/v23y2003i2p109-23.html.
- Adrian Vetta. Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. In 43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings, page 416, 2002. doi:10.1109/SFCS.2002.1181966.
- Robert J. Weber. Making More from Less: Strategic Demand Reduction in the FCC Spectrum Auctions. *Journal of Economics & Management Strategy*, 6(3):529–548, 1997. doi:10.1111/j. 1430-9134.1997.00529.x.
- 27 Martin L Weitzman. Prices vs. quantities. *The review of economic studies*, 41(4):477–491, 1974.
- 28 Robert Wilson. Auctions of Shares. *The Quarterly Journal of Economics*, 93(4):675–689, 1979. URL: http://www.jstor.org/stable/1884475.

A Omitted Proofs from Section 3

A.1 Proof of Lemma 3

Recall the statement of Lemma 3: for any distribution F and any cap-and-price auction $M(C, \underline{p}, \overline{p})$, there exist a cap C^0 and price floor \underline{p}^0 such that $W(M(C^0, \underline{p}^0, \infty)) \geq \frac{1}{2}W(M(C, p, \overline{p}))$.

Proof. We show that given any mechanism $M = (C_1, p, p)$, we can construct a mechanism M^0 with price ceiling ∞ such that $W(M^0) \ge \frac{1}{2}W(M)$.

We can decompose the expected welfare of M into (a) the welfare attained from the first (at most) C_1 licenses sold, and (b) the incremental welfare attained from any licenses sold after the first C_1 .

Note that the auction $M^1 = (C_1, \underline{p}, \infty)$, which is M but with price ceiling set to ∞ , achieves welfare precisely equal to the former of these two parts. This is because M^1 always allocates at most C_1 licenses, and will allocate them efficiently subject to all marginal values being at least p.

Next consider the second of these two parts of the welfare of M. If the expected welfare in the second part is negative, then we are already done, so suppose not. Whenever more than C_1 licenses are sold, all licenses are sold at price p, and therefore have marginal value at least p. Auction $M^2 = (\infty, p, \infty)$, with no license cap and with a price floor of p, will also sell all such licenses that have marginal value at least p. We note, however, that M^2 additionally also includes the marginal contribution of the first C_1 licenses. But we claim that this contribution is non-negative: when the event occurs that more than C_1 licenses are sold, the marginal contribution of the first C_1 licenses to the welfare can only be greater than that of those beyond the first C_1 . Thus, since the expected welfare in the second part is non-negative, the welfare is only higher if we also include the contribution of the first C_1 licenses whenever more than C_1 licenses are sold. This is precisely the welfare of auction $M^2 = (\infty, p, \infty)$, so the welfare of M^2 is therefore at least that of the second of the two parts of the welfare of M.

We conclude that $W(M^1) + W(M^2) \ge W(M)$, and hence either $W(M^1)$ or $W(M^2)$ is at least $\frac{1}{2}W(M)$.

B Omitted Proofs from Section 4

B.1 Proof of Lemma 10

First recall the statement of the lemma. If C and p are chosen to maximize W(M(C, p)), then $E[W(V, x(V)) \mid d(V) \geq C] \geq 0$.

Proof. Suppose not. Then it must be that $E[W(V,x(V)) \mid d(V) \geq C] < 0$. We will show this implies $W(M(C \square 1,p)) > W(M(C,p))$, contradicting the optimality of C. To see why, note that when d(p) < C, the welfare of the two auctions is identical. Write x and x^0 for the allocation functions from M(C,p) and $M(C \square 1,p)$, respectively. When $x \geq C$, we have $x^0 = C \square 1$, and $V(x^0) \geq \frac{C \square 1}{x} V(x)$ by concavity. Similarly, $Q(x^0) \leq \frac{C \square 1}{x} Q(x)$ by convexity. Thus, for any V such that $d(V) \geq C$ and hence x(V) = C, we have

$$V(x^{0}(V)) \square Q(x^{0}(V)) \geq \frac{C \square 1}{x(V)} (V(x(V)) \square Q(x(V))) = \frac{C \square 1}{C} (V(x(V)) \square Q(x(V))).$$

Taking expectations, we therefore have

$$E[V(x^{0}(V)) \square Q(x^{0}(V)) \mid d(V) \geq C] \geq \frac{C \square 1}{C} (E[V(x(V)) \square Q(x(V)) \mid d(V) \geq C])$$

$$> E[V(x(V)) \square Q(x(V)) \mid d(V) \geq C]$$

where the last inequality follows because $E[V(x(V)) \square Q(x(V)) \mid d(V) \ge C] < 0$ by assumption. This implies $W(M(C \square 1, p)) > W(M(C, p))$, which is the desired contradiction.

B.2 Proof of Lemma 11

First recall the statement of the lemma. Choose some quantity \tilde{C} and a number of licenses $x \leq \tilde{C}$. Then $P(\tilde{C}) \times Q(x) \leq C (P(C) P(C/2))$.

Proof. Suppose $x \ge \tilde{C}/2$. Then $Q(x) \ge \frac{x}{\tilde{C}/2}$ $Q(\tilde{C}/2)$ by convexity. So $P(\tilde{C})$ $x \square Q(x) \le x$ $P(C) \square \tilde{Q}(C/2)\tilde{/}(C/2) \stackrel{\sim}{\le} C$ $(P(\tilde{C}) \square \tilde{P}(C/2))$ as claimed.

Next suppose $x < \tilde{C}/2$. Then $(x, P(\tilde{C}) \ x)$ and (x, Q(x)) both lie between the line through the origin with slope $P(\tilde{C})$, and the line between $(\tilde{C}, Q(\tilde{C}))$ and $(\tilde{C}/2, Q(\tilde{C}/2))$. Their difference is therefore at most twice the difference between those two lines at x-coordinate $\tilde{C}/2$, which is $(\tilde{C}/2)$ $(P(\tilde{C}) \square P(\tilde{C}/2))$.

B.3 Omitted Details from the Proof of Theorem 7

Recall the statement of Theorem 7: for any cap C and price floor \underline{p} , there exists C^0 and a constant c such that c $W(M(C^0)) + W(M^{(1)}) \ge W(M(C,p))$.

As in the proof of Theorem 8, the welfare generated by M(C, p) can be broken down as follows:

$$W(M(C, p)) = \Pr_{V \square F}[d(V) \ge C] E[W(x(V), V) \mid d(V) \ge C] +$$

 $\Pr_{V \square F}[d(V) < C] E[W(x(V), V) \mid d(V) < C]$

We will break down the second term into three sub-terms. Recall that we can assume $\Pr[d(V) \geq C] \leq 1 \Box 1/e$. Choose $C_{\text{med}} < C$ so that $\Pr[d(V) \geq C_{\text{med}}] = 1 \Box 1/e$. We will write $x_i = x^\Box + x_i^> + x_i^<$. If $x_i \geq C_{\text{med}}$ then we set $x^\Box = x_i$ and $x^> = x^< = 0$, otherwise we set $x^\Box = 0$. In this latter case, we set $x^>$ and $x^<$ similarly as in case 1, except that we consider marginal values above and below $P(C_{\text{med}})$ rather than P(C). That is, if ϑ_i is the largest $j \geq 1$ such that $v_i(j) \geq P(C_{\text{med}})$ (or $\vartheta_i = 0$ if $v_i(1) < P(C_{\text{med}})$), we have $x^> = \min\{x_i, \vartheta_i\}$ and $x^< = x_i \Box x^>$. And as before, we will define $V^<$ as $V^<(x) = V^<(x|\vartheta_i)$, the valuation of i counting only those marginal values less than $P(C_{\text{med}})$.

Using convexity of Q as in case 1, we then have

$$W(M(C, p)) \leq \Pr_{V \subseteq F}[d(V) \geq C] \ E[W(x(V), V) \mid d(V) \geq C]$$

$$+ (V_{i}(x_{i}^{\square}(V)) \square Q(x^{\square}(V)))dFV$$

$$+ (V_{i}(x_{i}^{>}(V)) \square Q(x^{>}(V)))dFV$$

$$+ (V_{i}(x_{i}^{>}(V)) \square Q(x^{>}(V)))dFV$$

$$+ (V_{i}(x_{i}^{>}(V)) \square Q(x^{<}(V)))dFV$$

$$+ (V_{i}^{<}(x_{i}^{<}(V)) \square Q(x^{<}(V)))dFV$$

so that $W(M(C, p)) \le (2A) + (2B) + (2C) + (2D)$.

As in Theorem 8, we have that $W(M(C)) \ge (2A)$. For (2B), note that the probability that every bidder i has $x_i < C_{med}$ is at least 1/e, from the definition of C_{med} . If we write p_i for the probability that $x_i < C_{\text{med}}$, then we have $\bigcap_{i=1}^{Q} p_i \ge 1/e$. Subject to this condition, the value of $\bigcap_{i=1}^{Q} (1 \square p_i)$ is maximized by setting all p_i equal, in which case we have $(1 \square p_i) \le n(1 \square e^{\square 1/n}) \le 1$. But this means that the expected number of agents with $x_i \stackrel{'}{\geq} C_{\text{med}}$ is at most 1. Thus, (2B) is at most the value of allocating to just a single bidder with $x_i > C_{med}$, which is at most the maximum possible welfare attainable from allocating to any one bidder. We therefore have that $W(M^{(1)}) \ge (2B)$.

We next claim that $W(M(C_{\text{med}})) \ge (2C)/4$. To see why, note that

 $\leq 4W(M(C_{med}))$

where the last inequality follows because, when imposing a cap of C_{med} on an allocation of total size at most 2C_{med}, the value obtained is at least the top half of all marginal values, which is more than half of all marginal values.

Finally, we claim that $W(M(C_{\text{med}}/2)) \ge (2D)/21$. Define $\Psi := (C_{\text{med}}/2) (P(C_{\text{med}}) \square$ $P(C_{\text{med}}/2)$). Then Lemma 11, combined with the fact that the total demand is at most C_{med} with probability 1/e, implies that $(2D) \leq (2/e) \Psi$.

We claim that for $\theta = \frac{8}{3(e \square 1)} \approx 1.55$, we must have $\Pr[V(C_{\text{med}}) < Q(C_{\text{med}}) \square \theta \Psi \mid d(V) > 0]$ $C_{\text{med}} \leq 1/4$. That is, conditional on having total demand at least C_{med} , the probability that the welfare generated at C_{med} is more negative than $\Box 6 \, \Psi$ is at most 1/4. Suppose not: then consider the difference in welfare between M(C, p) and $M(C, P(C_{med}))$. The difference is that (2D) would be removed, as would any (negative) contribution with $V(C_{\text{med}}) < Q(C_{\text{med}})$ and $d(V) > C_{\text{med}}$. As we noted above, the total contribution of (2D) to the welfare is at most $2\Psi/e$. But the contribution due to the event $V(C_{\text{med}}) < Q(C_{\text{med}}) \square \Psi$ and $d(V) > C_{\text{med}}$ is at most $\Box(\theta\Psi)(1 \Box 1/e)(3/4)$. So as long as $\theta \geq (8/3(e \Box 1)) \approx 1.55$, we would have $W(M(C, P(C_{med}))) > W(M(C, p))$. This contradicts the optimality of M(C, p).

We can therefore assume $\Pr[V(C_{\text{med}}) \geq Q(C_{\text{med}}) | \theta \Psi | d(V) > C_{\text{med}}] \geq 1/4$. By concavity of V and convexity of Q, this means $\Pr[V(C_{\text{med}}/2) \geq Q(C_{\text{med}}/2) + (2 \square \theta)\Psi/2 \mid d(V) > 0]$ $C_{\text{med}}] \ge 1/4$. Since $d(V) \ge C_{\text{med}}$ with probability at least $1 \square 1/e$, the total welfare generated by $M(C_{\text{med}}/2)$ from the events $d(V) \geq C_{\text{med}}$ is at least $(2 \square \beta)(1 \square 1/e)\Psi/8$, and hence $W(M(C_{\text{med}}/2)) \ge (2 \square 6)(1 \square 1/e)\Psi/8$. Since $(2 \square 6)(1 \square 1/e)e/16 \ge 1/21$, the result follows. We conclude that

$$W(M(C)) + W(M^{(1)}) + 4W(M(C_{med})) + 21W(M(C_{med}/2)) \ge W(M(C, p)).$$

This implies Theorem 7, with c = 26.

Cap-and-Price Auction Outcomes are not Fully Efficient

We note that cap-and-price auctions cannot always implement the fully optimal allocation rule for every distribution F. For an allocation rule to be implementable by some M(C, p, p), it must be that on every input V, either the total allocation is C, or the total allocation is

at most C and each agent wins precisely their marginal bids above p, or the total allocation is at least C and each agent wins precisely their marginal bids above \bar{p} . When there is uncertainty about V, such a restricted allocation rule might return suboptimal allocations on some realizations.

In fact, we note that there are cases where, even under truthful reporting, no cap-and-price auction can achieve a non-vanishing approximation of the unrestricted welfare-optimal allocation.

Example 13. We present an example for which the expected welfare of any cap-and-price auction $M(C, \underline{p})$ is at most an O(1/n)-approximation to the unrestricted allocation that can optimize individually for each realization of the valuation curves, the welfare of which we call the "first-best welfare" in line with economics terminology.

Let $Q(x)=x^2$. There is a single firm participating in the auction. That firm's valuation curve is drawn according to a distribution F over n different valuation curves, which we'll denote $V^{(1)},\ldots,V^{(n)}$. For all i, valuation $V^{(i)}$ is defined by $V^{(i)}(x)=2^{i+1}$ x. These curves are depicted in Figure 2. The probability that $V^{(i)}()$ is drawn from F is proportional to $-\frac{1}{2}$. That is, the firm has valuation $V^{(i)}$ with probability $\frac{1}{62^{2i}}$, where $\theta=\frac{n}{i+1}2^{\square 2i}=\frac{1}{3}(1\square 2^{\square 2n})$ is the normalization constant.

Then in expectation over the realizations of V(), the first-best welfare is

$$\sum_{i=1}^{X^n} 2^{2i} \frac{1}{62^{2i}} = \frac{1}{6} n.$$

Consider a cap-and-price auction M(C, p). For any cap C, the optimal price floor is to set p=2C. This eliminates all possible negative welfare contributions at C. As a sanity check, any smaller price floor allows negative contributions, yet any larger price floor excludes positive contributions. For $C^0 \in (2^k, 2^{k+1}]$, , if the firm has valuation $V^{(i)}$ with $i \leq k$, then all marginal values are strictly below p, so no licenses are allocated and the welfare generate is 0. On the other hand, if the firm has valuation $V^{(i)}$ for any $i \geq k+1$, then since at most $C < 2^{k+1}$ licenses can be purchased, the auction generates welfare at most $2^{i+1}(2^{k+1}) \square (2^{k+1})^2 = 2^{2k+2}(2^{i\square k} \square 1)$. This yields expected welfare

$$\begin{array}{lll}
\chi^{n} \\
 & i=k+1
\end{array}$$

$$2^{2k+2}(2^{i\square k} \square 1) & \frac{1}{62^{2i}} = \frac{1}{6} \sum_{i=k+1}^{\chi^{n}} 2^{\square 2(i\square k\square 1)}(2^{i\square k} \square 1) \\
 & = \frac{1}{6} \sum_{j=1}^{\chi^{n} k} 2^{\square 2(j\square 1)}(2^{j} \square 1) \\
 & \leq \frac{1}{6} \sum_{j=1}^{\chi^{n} k} 2^{\square (j\square 1)} \\
 & = \frac{1}{6}(2 \square 2^{\square (n\square k\square 1)}) \\
 & \leq \frac{2}{6} = o(n).$$

In comparison to opt, any cap-and-price policy is off by an order of n, so no o(1/n)-approximation to the first-best welfare is possible. This concludes the example.

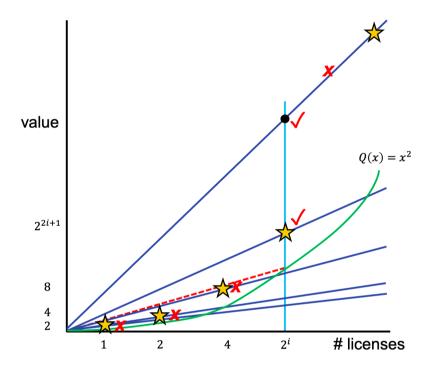


Figure 2 A depiction of Example 13, with the welfare achieve by the vertical cap and dashed price floor denoted by the checks and x's.