

Efficiency Guarantees in Large Markets

Michal Feldman* Nicole Immorlica† Brendan Lucier‡
Tim Roughgarden§ Vasilis Syrgkanis¶

Abstract

We present an analysis framework for bounding inefficiency in markets with many agents. We use this framework to demonstrate that, in many markets, the efficiency in the large is much smaller than the worst-case bound. Our framework also differentiates between markets with similar worst-case performance, such as simultaneous uniform-price auctions and greedy combinatorial auctions, thereby providing new insights about which markets are likely to perform well in realistic settings.

1 Introduction

Markets are the basic mechanism which allocate scarce resources to agents. Carefully designed markets have seen a wide range of applications, including classic settings like the Dutch flower auctions and used car auctions, as well as modern settings like advertising auctions and FCC spectrum auctions. The prevalence of designed markets raises a pertinent question: which designs guarantee efficient outcomes? Standard results demonstrate that in the worst-case, many designs are highly inefficient. These examples typically include a small number of agents while, in reality, most applications have a large number of agents.

Are markets with many agents more or less efficient than those with few? One answer is that inefficiencies persist in the large: small markets can often be embedded in large markets through appropriate replication techniques. But for many natural market settings, we might hope that the answer is “more efficient:” perhaps the influence of each agent on the outcome is small. For example, in many market settings with a large number of small agents, the action of a single agent has only negligible effect on prices. This suggests that agents can be accurately modeled as “price-takers.” While it turns out to be difficult to characterize exact equilibrium behavior, we can use this intuition to bound the efficiency of equilibria by arguing about the payoffs of feasible deviation strategies.

The goal of this paper is to develop an analytical framework for bounding the inefficiency of equilibria — the *price of anarchy (POA)* — in fundamental classes of markets when there

*Tel-Aviv University, mfeldman@post.tau.ac.il

†Microsoft Research, nicimm@microsoft.com

‡Microsoft Research, brlucier@microsoft.com

§Stanford University, tim@cs.stanford.edu

¶Microsoft Research, vasy@microsoft.com

is a large number of agents. The worst-case POA is well understood in the models we study. These bounds, like any worst-case bounds, tend to be overly pessimistic and are determined by pathological examples. Our goal is to prove qualitatively better POA bounds, assuming only that the markets are “large.” Importantly, our framework differentiates between markets with similar worst-case performance, thereby providing new insights about which markets are likely to perform well in realistic settings.

1.1 Summary of Results

We present a general analysis framework and several instantiations for well-studied applications.

A General Analysis Framework. We define (λ, μ) -smooth in the large market sequences, and show that the POA of markets in such a sequence approaches $\lambda/(1-\mu)$. This notion inherits the generality and robustness of previous smoothness definitions [18, 19, 24, 25]: many different applications are amenable to a smoothness-type analysis, and the resulting POA bounds apply to a wide range of equilibria, such as correlated and Bayes-Nash equilibria.

We also define an intuitive and easy-to-apply sufficient condition for a market sequence to be smooth in the large. The condition formalizes the idea that equilibria in large markets should be more efficient because no individual can significantly affect the market’s outcome. Precisely, one defines for each agent an *approximate utility*, which represents the utility the agent would receive were he “infinitely small.”¹ The sufficient condition requires that the approximate utility is (λ, μ) -smooth with respect to the actual market. We prove that if this condition is satisfied, then the market sequence is (λ, μ) -smooth in the large.

We apply our framework to obtain POA bounds for “large versions” of a number of well-studied models. Our flagship application is a comparison of two well-studied combinatorial auction mechanisms: simultaneous single-item auctions and greedy combinatorial auctions. We first study simultaneous single-item auctions, where we obtain full efficiency in the large even with general combinatorial valuations. In contrast, we prove that greedy combinatorial auctions are nearly as inefficient in the large as in worst-case instances. This distinction is especially notable because these two auction formats have comparable worst-case performance when valuations exhibit complements.

Simultaneous Uniform Price Auctions. We prove that in a “large” combinatorial auction setting, running a separate simultaneous uniform price auction for each type of good leads to fully efficient equilibria in the limit. Specifically, we consider a setting where a fixed set of m different goods, each of some supply, are auctioned off to a set of n bidders. Bidders have combinatorial valuations over sets of allocated units of different items and we assume that they only want at most some fixed number r of units of each type of good. We grow the market by letting the number of bidders increase, the number of units of supply of each good increase with the number of bidders, and letting each bidder fail to arrive in the market with some probability $\delta > 0$. Under these conditions we show that the worst-case expected welfare of Bayes-Nash equilibria of these markets converges at a rate of $1 - O(\frac{1}{\sqrt{n}})$ to the

¹This is purely for the sake of analysis; the resulting POA bound applies to the equilibria of the actual market.

expected optimal welfare. We make no assumptions about the distributions of preferences, other than that they are independent. For example, we do not require that the market grows via the addition of new bidders with valuations distributed identically to the original bidders.

In the worst case, without a largeness assumption, the POA of simultaneous second-price auctions (a special case of simultaneous uniform-price auctions) with general combinatorial valuations is $\Omega(\sqrt{m})$. A striking feature of our result is that full efficiency is obtained at equilibrium even though bidders are forced to report far fewer parameters (linear in m) than are present in their (combinatorial) valuations. These near-optimal equilibria are not “truthful outcomes” in any sense. This highlights the power of adopting a smoothness-based framework — we can prove convergence to full efficiency without characterizing these near-optimal equilibria.

Greedy Combinatorial Auctions. We next use our framework to analyze greedy combinatorial auctions. Similar to the previous setting, there is a fixed set of m different goods and each good is in some supply. Each bidder wants a specific bundle of at most d items. A bidder might want up to r copies of his bundle, for some fixed constant r . In the greedy auction, each bidder submits a set of r marginal bids for his interest set. Bids are allocated in decreasing order as long as they are satisfiable, i.e., no item in the desired set has run out of supply. The worst-case POA of greedy combinatorial auctions is at most $d + 1$ [15].

We prove that if the number of bidders grows and the supply of each good grows in expectation, but is sufficiently uncertain, then the POA of Bayes-Nash equilibria of greedy combinatorial auctions converges to exactly d . This bound matches the algorithmic approximation of the greedy algorithm when there are no incentive constraints.

Discussion of Results. Our results demonstrate that in many of the game-theoretic models in which the POA has been studied, the POA in large games is much smaller than the worst-case bound. In some (but not all) cases, the POA approaches 1 as the game grows large. We suspect that our better POA bounds are more relevant for many computer science settings, which often feature a large number of “small” players in uncertain environments.

Our results have interesting implications for mechanism design. Theoretically optimal mechanisms can be relatively complex — for example, when bidders demand multiple units of a good, the welfare-maximizing (i.e., VCG) mechanism charges different prices for different units of the good. Some of our results give a new sense in which “simple” mechanisms can be near-optimal. For example, we prove that the (theoretically suboptimal) uniform-price mechanism has welfare approaching that of the VCG mechanism as the market grows large.

Finally, our approach provides a novel way to differentiate between competing mechanisms, by comparing equilibrium performance as the market grows large. For example, in the model we study for greedy combinatorial auctions, both greedy and simultaneous item auctions have worst-case POA linear in the size d of the desired bundles [15, 9]. Our large-game analysis suggests that single-item auctions (which are fully efficient in the limit) may be preferable to greedy auctions (which are not) in large, realistic instances.

1.2 Our Techniques

There are many ways to “let a market grow large,” and it is clear that some of these will not result in POA bounds better than the worst-case bounds. For example, in auctions, it is not enough to merely assume that the number of bidders is sufficiently large — we do not assume that bidders’ valuations are IID, so one can always add “dummy bidders” to a worst-case instance without affecting the equilibria (or the POA). A similar comment applies even in an IID Bayesian setting if we assume only that the number of bidders and available items tend to infinity. This is illustrated in the following example, taken from [23].

Example. (Inefficiency without randomness) Consider a setting where k units of a good are auctioned off to $n = k$ bidders. Each bidder wants at most two units. The value of each bidder for the first unit is the maximum of two random samples drawn uniformly from the interval $[2, 3]$ and independently for each bidder. The marginal value for the second unit is the minimum of the two samples. The units are sold via a uniform price auction: each bidder submits two marginal bids, the k highest marginal bids win and each bidder pays the highest losing marginal bid for each unit he won. We next show that the following is an equilibrium of the market, for any k : each bidder submits his higher marginal bid truthfully and 0 as his second marginal bid. The uniform price is 0 and each bidder derives utility v_i^1 . For any bidder to get a second unit he needs to bid at least 2, which would increase the uniform price to at least 2 on both units. The increase in payment is 4, while the increase in value is at most 3, and hence is not profitable. The expected welfare of this equilibrium is $\approx 2.67 \cdot k$. The expected optimal welfare is the expected sum of the highest k of $2k$ samples from $U[2, 3]$, which is approximately $2.75 \cdot k$ as k grows large. ■

The bad equilibrium in the above example relies on an exact match between the number of bidders and of units. An important ingredient of our large market model is *probabilistic demand*, meaning that a small random sample of the bidders fail to show up [23]. In addition to circumventing the “knife edge” example above, this assumption arguably increases the verisimilitude of the model. In many real-world auction settings — especially in advertising auctions, where the participants might well be chosen by a search engine’s heuristic matching algorithm — bidders cannot be sure who else will choose (or be chosen) to participate in an auction.²

At a technical level, the primary challenge in applying our framework is to define the approximate utility of each bidder so that it both approximates the actual utility and is smooth with respect to the actual market. Designing these approximate utilities requires understanding the approximations bidders may reasonably sustain in large markets (again, this is for the analysis only, not a behavioral assumption).³ For example, in our analysis, we think of bidders as ignoring the impact of their bids on the prices.

Finally, proving that the approximate utilities are smooth with respect to the actual market requires some technical finesse: the smoothness arguments work by defining deviations for bidders that are functions of the types. Often, the deviations in auctions ask bidders to

²We also consider probabilistic supply in Section 5.

³Though as a side benefit, our results also imply that playing an equilibrium with respect to the approximate utilities yields an approximate equilibrium of the large market that retains all of the efficiency guarantees that we prove. Conceivably, this could be an accurate model of how bidders behave in large and complex markets.

bid high values on their optimal allocation. However, in the presence of noise, the optimal allocation is a random variable. One technical contribution is to circumvent this difficulty by re-interpreting the noise as type uncertainty in a Bayesian game.

1.3 Related Work

The worst-case POA (without the largeness assumption) is well understood in all of the models that we study. For the POA of uniform-price auctions and simultaneous item auctions, see [5, 3, 11, 10, 16, 7]. For the POA of greedy combinatorial auctions, see [15].

Most previous work in computer science that concentrates specifically on games with many players is motivated by complexity concerns. For example, the literature on “compact representations” of games proposes succinct descriptions, with size polynomial in the number of players, that are well structured but still rich enough to capture many interesting applications. See [14, 4, 21] and the references therein for many examples. These references also discuss the well-studied problem of computing equilibria efficiently in compactly represented multi-player games. These results suggest that computing an exact equilibrium in the markets we study is computationally difficult, suggesting that a succinct characterization of these equilibria is unattainable. This motivates our approach of bounding the efficiency of markets without explicitly characterizing the equilibria.

Large markets have also been considered in the economics literature; see [17, 22] for some early examples and Kalai [13] for work on the robustness of equilibria in large games. The closest work to ours in this literature is Swinkels [23], who studies a single-good uniform price auction with decreasing marginal valuations and with demand and supply uncertainties under the same large market assumptions that we make. Our simultaneous uniform price auction result generalizes the supply uncertainty result of [23] to allow for heterogeneous goods and our greedy auction result generalizes the demand uncertainty result of [23] by allowing bids on bundles of items, rather than single items (the uniform price auction is a special case of the greedy auction for $d = 1$). Our framework also allows us to relax some technical assumptions made in [23] regarding the valuation distributions. In another closely related work, Jackson and Manelli [12] study conditions under which the outcome of a market at equilibrium approximates the fully efficient outcome and identify a crucial property to be insensitivity of prices with respect to individual reports. They show that such insensitivity limits the gains that any participant can obtain via strategic manipulation, and argue that natural markets achieve this property as they approach a large-game limit.

The perspective and goals of all of these works in the economics literature differ from ours in several predictable ways. Their emphasis has been on understanding what equilibria “look like”, and ideally solving for them explicitly (if not in large finite games, then at least “in the limit”); a technically difficult subproblem that often arises in this approach is to prove that the equilibria of large finite games approach the equilibria of a “limit game”. Because we care about equilibria only through their objective function values, we can bypass the problem of characterizing equilibria and their limits, and instead argue directly about the approximation guarantees obtained. In addition, all previous work in economics on efficiency in large games considered only the special case of full efficiency in the limit, as in our result on uniform-price auctions. No previous work considered models where inefficiency persists in the limit, as with our other result on greedy combinatorial auctions. Our smoothness-based

framework is general enough to cover both situations with a common analysis.

In a different direction, Alvarez and Budish [1] defined “strategyproof in the large” mechanisms, where truth-telling constitutes an approximate equilibrium as many players arrive in the market and submit bids drawn from the same distribution. Our work differs, in that we don’t need to make such symmetric strategy assumptions, and we can accommodate mechanisms where efficiency is achieved in the limit even though participants are not truthful in the limit.

Finally, independently of but subsequent to our work, Cole and Tao [6] analyzed the POA of the Walrasian mechanism under large market conditions. They apply smoothness bounds to show that the Walrasian mechanism approaches full efficiency in the limit as the market grows large. Whereas [6] provides a thorough analysis of a complex and general market mechanism, our primary focus is to develop a general smoothness framework and apply it to analyze and compare many different types of large games.

2 Preliminaries: Mechanisms, Equilibria and Price of Anarchy

In this work, we present a framework to study the efficiency of markets in the large. Consider a market with n bidders and m items. Each bidder $i \in [n]$ has a valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$, that assigns a value for each possible allocation of items. We will denote the set of possible valuations for bidder i with \mathcal{V}_i and the set of valuation profiles with $\mathcal{V} = \mathcal{V}_1 \times \dots \times \mathcal{V}_n$.

A mechanism \mathcal{M} consists of a triple $(\{S_i\}_{i=1}^n, \{x_i\}_{i=1}^n, \{P_i\}_{i=1}^n)$. S_i is a strategy space for each bidder (and $S = S_1 \times \dots \times S_n$). $x_i : S \rightarrow 2^{[m]}$ is an allocation function that maps a strategy profile to an allocation of items to bidder i , such that $x(s) = (x_1(s), \dots, x_n(s))$ is feasible (no two items are allocated to different bidders). $P_i : S \rightarrow \mathbb{R}_+$ is a payment function. A bidder’s utility for an allocation is his value minus his payment, i.e., $u_i(s; v_i) = v_i(x_i(s)) - P_i(s)$. We will be interested in analyzing the social welfare of an equilibrium strategy profile $s \in S$, which is the total value of the resulting allocation:

$$SW(s; v) = \sum_{i=1}^n v_i(x_i(s)) \tag{1}$$

The optimal feasible allocation for valuation profile v will be denoted by $\text{OPT}(v)$, i.e. $\text{OPT}(v) = \max_{x \text{ is feasible}} \sum_{i=1}^n v_i(x_i)$. The revenue of a mechanism is the sum of the payments, i.e., $\mathcal{R}(s) = \sum_{i=1}^n P_i(s)$.

We will consider a Bayesian setting in which each bidder’s valuation v_i is drawn independently from some distribution \mathcal{F}_i . A strategy function for bidder i is a (possibly randomized) mapping μ_i from V_i to S_i , which we think of as a specification of the strategy to use given a valuation. A Bayesian Nash equilibrium (BNE) is a profile of strategy functions such that no single bidder can increase his expected utility (over randomization in types and strategies) by unilaterally modifying his strategy. Formally, the profile of strategy functions $\mu = (\mu_1, \dots, \mu_n)$ is a BNE if for all i , all valuations $v_i \in \mathcal{V}_i$, and all alternative strategies

$s'_i \in S_i$, we have

$$\begin{aligned} & \mathbb{E}_{v_{-i} \sim \mathcal{F}_{-i}} [u_i(\mu_i(v_i), \mu_{-i}(v_{-i}); v_i)] \\ & \geq \mathbb{E}_{v_{-i} \sim \mathcal{F}_{-i}} [u_i(s'_i, \mu_{-i}(v_{-i}); v_i)]. \end{aligned}$$

Note that the non-Bayesian notion of Nash Equilibrium is a special case of the above, in which every distribution \mathcal{F}_i is a point mass.

The Bayes-Nash Price of Anarchy (BNE-POA) of a mechanism \mathcal{M} is the worst-case ratio between the expected optimal welfare and the expected welfare at equilibrium, over all type distributions and all BNE. That is,

$$\text{BNE-POA} = \max_{\mathcal{F}} \max_{\mu} \frac{\mathbb{E}_{v \sim \mathcal{F}}[\text{OPT}(v)]}{\mathbb{E}_{v \sim \mathcal{F}}[\text{SW}(\mu(v); v)]}, \quad (2)$$

where the maximum over strategy functions μ is taken over all BNE for distribution profile \mathcal{F} .

3 Smoothness in the Large

Sequences of mechanisms. We will typically work with a sequence of mechanisms $\{\mathcal{M}^n\}_{n=1}^{\infty}$, indexed by the number of participating players n . For shorter notation we will write $\{x^n\}$ to denote the sequence $\{x^n\}_{n=1}^{\infty}$.

In a sequence of mechanisms $\{\mathcal{M}^n\}$, everything will be changing parametrically with the number of players, such as the set of items m^n , the strategy spaces $S^n = (S_1^n, \dots, S_n^n)$, the allocation functions $x^n = (x_1^n, \dots, x_n^n)$, the payment functions $P^n = (P_1^n, \dots, P_n^n)$ and the valuation profile space $\mathcal{V}^n = (\mathcal{V}_1^n, \dots, \mathcal{V}_n^n)$. We will also denote with $\text{OPT}^n(\cdot)$ the optimal welfare, with $\mathcal{R}^n(\cdot)$ the revenue and with u_i^n the utility of player i in mechanism \mathcal{M}^n . For the moment one can imagine arbitrary ways for the mechanism to grow; in subsequent sections we give specific conditions for how the market should grow for our framework to be applicable.

Smoothness in the large. For finite games, Roughgarden [18] introduced the notion of *smoothness* as a method for bounding inefficiency of equilibria. The smoothness approach proceeds by exploring specific deviations, instead of characterizing the (potentially complex) structure of equilibria. This approach was specialized to the mechanism design setting via the notion of smooth mechanisms by Syrgkanis and Tardos [25]. We extend the notion of smoothness to large games. In what follows we present the specific extension in the context of mechanism design (i.e., large mechanisms), but the framework is widely applicable and the reader is directed to the full version of the paper for the formulation of the framework in general games. Intuitively, a sequence of mechanisms is said to be (λ, μ) -smooth in the large if for any ϵ , and a sufficiently large number of players, each player i has a special strategy that allows him to acquire a $\lambda \cdot (1 - \epsilon)$ fraction of his valuation for his optimal set of items, by paying no more than μ times the current price paid for these items.

Definition 1 (Smooth in the large). *A sequence of mechanisms $\{\mathcal{M}^n\}$ is (λ, μ) -smooth in the large if for any $\epsilon > 0$, there exists $n(\epsilon) < \infty$, such that for any $n > n(\epsilon)$, for any $v^n \in \mathcal{V}^n$, for each $i \in [n]$, there exists a strategy $s_i^{*,n} \in S_i^n$, such that for any $s^n \in S^n$:*

$$\sum_{i=1}^n u_i^n(s_i^{*,n}, s_{-i}^n; v_i^n) \geq \lambda(1 - \epsilon)\text{OPT}^n(v^n) - \mu \cdot \mathcal{R}^n(s^n) \quad (3)$$

The following theorem shows that if a sequence of mechanisms is (λ, μ) -smooth in the large, for some $\lambda, \mu \geq 0$, then its price of anarchy as $n \rightarrow \infty$ is at most $\frac{\max\{1, \mu\}}{\lambda}$. Moreover, it implies that for any sufficiently large but finite market of size n the price of anarchy of all Bayes-Nash equilibria is at most a $1 + \epsilon(n)$ multiplicative factor away from the limit price of anarchy, where the rate of convergence of $\epsilon(n)$ to 0 will depend on the application and can be derived from the proof of smoothness in the large.

Theorem 2. *If a sequence of games is (λ, μ) -smooth in the large then*

$$\limsup_{n \rightarrow \infty} \text{BNE-POA}^n \leq \frac{\max\{1, \mu\}}{\lambda}.$$

I.e., for any ϵ there exists a market size $n(\epsilon)$ such that for any $n \geq n(\epsilon)$, every Bayes-Nash equilibrium of the mechanism \mathcal{M}^n with value distributions $\mathcal{F}_1 \times \dots \times \mathcal{F}_n$ has expected social welfare at least $(1 - \epsilon) \frac{\lambda}{\max\{1, \mu\}}$ of the expected optimal welfare.

Proof. By (λ, μ) -smoothness in the large, for any ϵ there exists a market size $n(\epsilon)$ such that for any $n \geq n(\epsilon)$ the mechanism \mathcal{M}^n is a $(\lambda(1 - \epsilon), \mu)$ -smooth mechanism, as defined in [25]. Therefore, by the results in [25], the BNE-POAⁿ is at most $\frac{\max\{1, \mu\}}{\lambda(1 - \epsilon)}$. The theorem then follows. ■

3.1 Main Technique: Smooth Approximate Utility Functions

We present the notion of a (λ, μ) -smooth approximate utility function sequence with respect to a sequence of mechanisms $\{\mathcal{M}^n\}$.

Definition 3 (Smooth approximate utility). *Let $U_i^n : S^n \times V_i \rightarrow \mathbb{R}_+$ be a utility function for player $i \in [n]$, and let $U^n = (U_1^n, \dots, U_n^n)$ be a vector of utility functions. A sequence $\{U^n\}$ is a sequence of (λ, μ) -smooth approximate utility functions for the sequence of mechanisms $\{\mathcal{M}^n\}$ if the following two properties are satisfied:*

1. **(Approximation)** *The approximate utility U_i^n converges to the true utility u_i^n uniformly over $s^n \in S^n$ and $v_i \in \mathcal{V}_i$. I.e., for any ϵ , there exists $n(\epsilon) < \infty$, such that for any $n > n(\epsilon)$, for any $i \in [n]$ and $v_i \in \mathcal{V}_i$, and for any $s^n \in S^n$:*

$$\|u_i^n(s^n; v_i) - U_i^n(s^n; v_i)\| < \epsilon. \quad (4)$$

2. **(Smoothness)** *For each mechanism \mathcal{M}^n in the sequence, the approximate utility satisfies the following (λ, μ) -smoothness property with respect to \mathcal{M}^n : For any n , for any $v \in \mathcal{V}^n$, for any $i \in [n]$, there exists a strategy $s_i^{*,n} \in S_i^n$, such that for any strategy profile $s^n \in S^n$:*

$$\sum_{i=1}^n U_i^n(s_i^{*,n}, s_{-i}^n; v_i) \geq \lambda \text{OPT}^n(v) - \mu \cdot \mathcal{R}^n(s^n) \quad (5)$$

We show that if a sequence of mechanisms admits a (λ, μ) -smooth approximate utility sequence, and if its optimal social welfare increases at least at the same asymptotic rate as the number of players, then this sequence of mechanisms is (λ, μ) -smooth in the large.

Theorem 4. *If a sequence of mechanisms $\{\mathcal{M}^n\}$ admits (λ, μ) -smooth approximate utility functions, and $\text{OPT}^n(t) = \Omega(n)$, then $\{\mathcal{M}^n\}$ is (λ, μ) -smooth in the large.*

Proof. Since $\{\mathcal{M}^n\}$ admits (λ, μ) -smooth approximate utility functions $\{U^n\}$, we have that for any n and $v \in \mathcal{V}^n$ there exists strategies $s_i^{*,n}$ for each $i \in [n]$ such that, for any $s \in S^n$,

$$\sum_{i=1}^n U_i^n(s_i^{*,n}, s_{-i}; v_i) \geq \lambda \text{OPT}^n(v) - \mu \cdot \mathcal{R}^n(s).$$

By the approximation property of U_i^n we have that for any ϵ , there exists $n(\epsilon) < \infty$ such that for any $n > n(\epsilon)$: $u_i^n(s; v_i) \geq U_i^n(s; v_i) - \epsilon$ for any $v_i \in \mathcal{V}_i^n$ and $s^n \in S^n$. Thus:

$$\sum_{i=1}^n u_i^n(s_i^{*,n}, s_{-i}; v_i) \geq \lambda \text{OPT}^n(v) - \mu \cdot \mathcal{R}^n(s) - n \cdot \epsilon.$$

Since $\text{OPT}^n(v) = \Omega(n)$, for any $v \in \mathcal{V}^n$, we can write $\text{OPT}^n(v) \geq \rho \cdot n$ for some $\rho > 0$ and for sufficiently large n . Thus we get:

$$\sum_{i=1}^n u_i^n(s_i^{*,n}, s_{-i}; v_i) \geq \left(\lambda - \frac{\epsilon}{\rho} \right) \text{OPT}^n(t^n) - \mu \cdot \mathcal{R}^n(s).$$

Therefore, for any $\delta > 0$, we can pick ϵ appropriately small, such that $\lambda - \frac{\epsilon}{\rho} \geq \lambda(1 - \delta)$, which would then yield the theorem. ■

4 Simultaneous Uniform Price Auctions

We consider a setting with a growing number of n bidders and a fixed number of m different (types of) goods. There are k_j^n units of each good $j \in [m]$ which grows as $\Omega(n)$ with the number of bidders. Each bidder $i \in [n]$ has a valuation function $v_i : \mathbb{N}^m \rightarrow [0, H]$, that assigns a value for each possible bundle, depending on the number of units of each good. These functions are bounded in their demand for the number of units of each good. Specifically, let x_i^j denote the number of units of good j allocated to bidder i , and let $x_i = (x_i^1, \dots, x_i^m)$ be an allocation vector for bidder i . Then there is a publically known constant r such that: $v_i(x_i) = v_i(\min\{x_i^1, r\}, \dots, \min\{x_i^m, r\})$. We will also assume that these valuations are bounded away from zero for any non-empty allocation, i.e. $v_i(x_i) \geq \rho > 0$ for every non-zero x_i .

The units of each good $j \in [m]$ are simultaneously and independently sold via the means of a uniform price auction. The auctioneer solicits r bids $b_i^{j,1} \geq \dots \geq b_i^{j,r}$ from each bidder i for each good j , referred as marginal bids. All bids of good j (from all bidders) are ordered in a decreasing order, and each of the first k_j^n bids wins a unit. In the case of ties, bidders are processed in a random order, and all tying bids of a bidder are allocated sequentially in order until the supply of the good runs out. Every bidder is charged the highest losing marginal bid for good j for every unit of good j allocated to him. We will assume that no bid exceeds some fixed number B ; i.e., $b_i^{j,x} < B$ for every i, j, x . Since we assumed that $v_i(x_i) \leq H$, it is a weakly dominated strategy for a bidder to bid more than H on an

individual marginal bid, though our formulation allows even for $B > H$, as long as B doesn't grow with the market. We will denote by \mathcal{M}^n an instance of the simultaneous uniform price auction among n bidders.

Notably, the above auction is not truthful for many reasons. First, the auction format is not even rich enough to allow bidders to express their true valuations, as they are forced to place additively separable bids on the different goods. Second, even for a single type of good, a uniform-price auction is not truthful for bidders with multi-unit demands. Nonetheless, we will show that in large markets, under a particular type of demand uncertainty — where each bidder “fails to arrive” with constant probability — all equilibria achieve full efficiency.

Theorem 5 (Full Efficiency in the Limit). *In the setting described above, if each bidder fails to arrive in the market with probability δ , then the implied sequence of mechanisms is $(1, 1)$ -smooth in the large; hence full efficiency is achieved in the limit. Moreover, the fraction of the optimal welfare achieved at equilibrium converges to 1 at a rate of $1 - O\left(\frac{1}{\sqrt{n}}\right)$.*

Crucially, the fact that we recover full efficiency in the large is not trivial in our setting. For instance, if one removes the noisy arrival assumption, then existing examples in [23] show that even when there is only one type of good, inefficiency can persist in the limit.

Sketch of proof of Theorem 5. At a high level, Theorem 5 is established by showing that the simultaneous uniform price auction where each bidder fails to arrive with probability δ admits $(1, 1)$ -smooth approximate utility functions. The full proof is deferred to Appendix A.

The approximate utility functions we define will have the following intuitive interpretation: each bidder i looks at the $(k + 1)$ -th highest bid at each auction excluding his own bids. Denote this with P_j^{-i} . This is the price that the other bidders would have paid for each unit of good j had bidder i not been in the market. In bidder i 's approximate utility, he has the delusion that this is also the price he faces; i.e., any marginal bid that he submits that surpasses the price P_j^{-i} will win a unit at price P_j^{-i} .⁴ In the actual market this is obviously not true: to win $x \in \{1, \dots, r\}$ units, bidder i actually needs to exceed the x -th lowest winning bid in his absence, and his price will be equal to this bid which may be greater than his imagined price of P_j^{-i} . However, as we shall soon show, with the proposed noise in the system, the price P_j^{-i} is “sufficiently random” that it is distributed almost identically to this x -th lowest winning bid for any constant x .

In what follows we present some of the technical challenges and techniques in our proof. Following the framework of smooth approximate utilities, we first sketch the proof of the approximation and then the smoothness of the approximate utility functions described above.

Approximation. We first show (in Lemma 11) that the bidder's utility from any bid vector b converges to his approximate utility, as the market grows large. Technically, the two utilities differ either when the allocation is different or when the price paid is different. The allocation differs only when some of the bidder's marginal bids are among the $k + 1 - r$ and $k + 1$ highest bids,⁵ since this is the only case where the bidder may believe that his marginal bid is a winning bid (under his delusional utility) while it is actually a losing bid (under his true utility). However, due to the random arrival, for any bid b , the probability

⁴A bid that is equal to P_j^{-i} will pass through the tie-breaking rule.

⁵In the actual proof we also take care of tie-breaking.

that the number of bids above b is equal to some number x goes to 0 as $x \rightarrow \infty$. Thus, the probability of any of these events goes to 0. Now, since there is only a constant number of these events (by the assumption that r is constant), the probability of any bad event occurring goes to 0 (by the union bound). Finally, we show that the difference in price paid also goes to zero. Technically, the distributions of the $(k + 1)$ -th and the $(k + 1 + x)$ -th highest bids are identical for any $x \in [-r, r]$. Thus, their expectation converges to zero as well (as they are bounded random variables).

Smoothness. The other part of the proof (Lemmas 9, 10) shows that these delusional utilities satisfy the $(1, 1)$ -smoothness property. Observe that under this delusion a bidder believes that he can always grab his optimal set of items at the current price in which they are sold. This is essentially the $(1, 1)$ -smoothness property. However, there are two crucial subtleties that need to be handled carefully. First, the prices of the goods are random, thus unknown to the bidder. Second, the optimal set of items for a bidder is also random, as it depends on who arrives in the market (which is not observed by the bidder when he decides his bid vector). The first problem is bypassed by observing that since these are threshold price mechanisms, the bidder can simply bid sufficiently high (even overbid). Specifically, if a bidder's optimal allocation is $x_i = (x_i^1, \dots, x_i^m)$, where x_i^j denotes the number of units of good j , then by bidding sufficiently high on the x_i^j highest marginal bids on each good j , he will almost surely win the items, or otherwise some price must be so high that we can charge the welfare loss to some other allocated bidder. We will show that bidding $v_i(x_i)$ as the first x_i^j marginal bids on each good j is sufficiently high to establish our $(1, 1)$ -smoothness argument. To bypass the second problem, we observe that the utility of any bidder under this game is lower bounded by the utility if he can bid even when he doesn't arrive, but has value of 0. The latter game is a simultaneous uniform price auction with no noisy demand, but with a Bayesian uncertainty on the values. Thus we will use a technique similar to the one used to show that smoothness for complete information games implies smoothness for games with Bayesian uncertainty in the values [19, 24, 25]. In particular, the smoothness deviation samples an arrival vector from the distribution and uses this random sample as a proxy for the true arrival vector, targeting the optimal bundle under this random sample.

4.1 Constant Inefficiency in the Limit under Supply Uncertainty

The noisy arrival of bidders can be seen as a type of demand uncertainty. In prior work of Swinkels [23], and in Section 5 which generalizes the prior work of Swinkels to combinatorial auctions with fixed demand sets, the uncertainty instead regards the supply: namely the probability that the number of units of the good equals any fixed number goes to 0 as the market grows large. For these settings, supply uncertainty is sufficient to recover full efficiency in the limit. In contrast, for simultaneous uniform price auctions, supply uncertainty can lead to a constant factor inefficiency even in the limit. In particular, it sustains a “search friction” in the limit: bidders do not know which items will have higher supply and thereby cannot decide which items to target. At equilibrium, their supply prediction ends up leading to constant factor inefficiencies that do not vanish. The following is a concrete counter example showing that supply uncertainty may lead to constant inefficiency in simultaneous uniform price auctions.

Consider a simultaneous uniform price auction game with two types of goods A and B .

For each market size n , there are $t = n/3$ unit-demand bidders that have only a value of $1/2$ for each unit of good A and no value for a unit of good B , i.e. $v_i(x_i) = \frac{1}{2}\mathbf{1}\{x_i^1 \geq 1\}$. We refer to these bidders as type a bidders. There are also t unit-demand bidders that have value $1/2$ only for units of good B and not of good A and we refer to them as type b bidders. Finally, there are t unit-demand bidders each having a value of 1 for each unit of each good and desiring only one unit from some of the two goods, i.e. $v_i(x_i) = \mathbf{1}\{x_i^1 + x_i^2 \geq 1\}$. We refer to them as type c bidders. The supply of each good is distributed uniformly in $[0, t]$. Obviously, as $t \rightarrow \infty$ this supply distribution satisfies the property that the supply being equal to any fixed number goes to 0.

Equilibrium. We argue that the following is an equilibrium: Each type c bidder picks uniformly at random one good A or B and submits a bid of 1 at the uniform price auction for that good. Each type a bidder submits a bid of $1/2$ at the uniform price auction for good A and each type b bidder submits a bid of $1/2$ for good B .

Equilibrium verification. This is obviously an equilibrium for the type a and b bidders, since they are essentially unit-demand bidders in a single uniform price auction, hence the mechanism is dominant strategy truthful from their perspective. Thus it remains to argue that type c bidders don't want to bid on both items. At each item, for them to win a unit they have to bid at least $1/2$, or otherwise they will lose to the type a or type b bidders. Moreover, the uniform price that they will have to pay on each good is always at least $1/2$, since there are always t bidders bidding $1/2$.

Moreover, observe that in the limit of many bidders we can essentially assume that exactly half of the type c bidders go for item A and exactly half go for item B (in fact we could have also analyzed the less natural equilibrium where this happens deterministically for every n , assuming $n/3$ is even). A type c bidder's utility at the current equilibrium strategy is $1/2$ when the supply of the good that he chose is less than or equal to $t/2$, since he pays $1/2$, and it is 0 when the supply is more than $t/2$, since he pays 1. Thus his expected utility is $1/4$.

When he bids on both items, then observe that whenever he wins a unit at both auctions he has to pay a price of $2 \cdot 1/2$, thus getting 0 utility. Thus he gets any utility only when he wins a unit at exactly one of the two auctions and only when he wins it at a price of $1/2$. There are only two possible reasonable bids for the bidder 1 or $1/2$. When he bids $1/2$ he is tying with the type a or b bidders and ties are broken at random. Let $p(b)$ be the probability that a bidder wins at auction A or B with a bid of b . Observe that $p(1/2) < p(1) = 1/2$. Thus if $b_A, b_B \in \{1/2, 1\}$ are the bids of the bidder in each auction, then his utility is: $p(b_A)(1 - p(b_B))\frac{1}{2} + p(b_B)(1 - p(b_A))\frac{1}{2} = \frac{1}{2}(p(b_A) + p(b_B) - p(b_A)p(b_B))$. For $0 \leq p(b) \leq 1/2$, the latter is maximized at $p(b_A) = p(b_B) = 1/2$, leading to a utility of $1/4$. In essence, the only reasonable bid of a type c bidder is to submit 1 on one of the two goods.

Sub-optimality. Now we argue about the suboptimality of this equilibrium as $n \rightarrow \infty$. The optimal allocation is to give as many units of either good A or B as possible to type c bidders and all remaining units to type a or b bidders. There are always enough remaining

type a or b bidders for all units to be allocated. Thus the expected optimal welfare is:

$$\begin{aligned}
\mathbb{E}[\text{OPT}(k_A, k_B)] &= \mathbb{E} \left[\min\{k_A + k_B, t\} + \frac{1}{2} (k_A + k_B - t)^+ \right] \\
&= \mathbb{E} \left[k_A + k_B - \frac{1}{2} (k_A + k_B - t)^+ \right] \\
&= t - \frac{1}{2} \mathbb{E} [(k_A + k_B - t)^+] \\
&= t \left(1 - \frac{1}{2} \mathbb{E} \left[\left(\frac{k_A}{t} + \frac{k_B}{t} - 1 \right)^+ \right] \right)
\end{aligned}$$

As $t \rightarrow \infty$, then $\frac{k_A}{t}$ and $\frac{k_B}{t}$ are distributed uniformly in $[0, 1]$. Thus by simple integrations for two $U[0, 1]$ random variables x, y : $\mathbb{E}[(x + y - 1)^+] = 1/6$. Hence, $\mathbb{E}[\text{OPT}(k_A, k_B)] \approx \frac{11 \cdot t}{12}$.

On the other hand the expected welfare at equilibrium is simply:

$$\begin{aligned}
\mathbb{E}[SW(b)] &= 2 \cdot \mathbb{E} \left[\min\{k_A, t/2\} + \frac{1}{2} (k_A - t/2)^+ \right] \\
&= 2 \cdot \mathbb{E} \left[k_A - \frac{1}{2} (k_A - t/2)^+ \right] = t - \mathbb{E} [(k_A - t/2)^+]
\end{aligned}$$

For large enough t , $\mathbb{E} [(k_A - t/2)^+] \approx \frac{t}{8}$. Therefore $\mathbb{E}[SW(b)] \approx \frac{7 \cdot t}{8}$. Therefore, the ratio of the expected optimal welfare over the expected equilibrium welfare converge to $\frac{22}{21} > 1$. Hence the limit price of anarchy is strictly greater than 1.

It is worth noting that this example can be taken to the extreme, when there are m goods, t type c bidders are interested in all of the goods and each good has a set of t price setters interested only in that good, with value $1/2$. The supply of each good is distributed uniformly in $[0, \frac{2t}{m}]$. One equilibrium is for the type c bidders to pick one item uniformly at random and bid 1, while the price setting people bid truthfully on their good. As m grows large, then the total supply $\sum_j k_j$ is with high probability concentrated around its expected value, which is t . Thus the expected optimal welfare converges to t . On the other hand at equilibrium each good has approximately t/m type c bidders and the supply of that good is distributed $U[0, \frac{2t}{m}]$. Thereby the expected welfare from each good from calculation similar to the two good case, is $\frac{7t}{8m}$. Hence, the price of anarchy converges to $8/7$. The essence is that bidders cannot take advantage of the concentration of total supply, which the optimal welfare can.

5 Greedy Auctions

As in Section 4, we consider a setting with n bidders and a fixed number of m different (types of) goods. For this section, we focus on a restricted class of multi-unit single-minded valuation functions, which take the following form: each agent i has a desired set of items $S_i \subseteq [m]$ and a non-convex function $v_i : \mathbb{N} \rightarrow [0, H]$, where $v_i(\ell)$ denotes agent i 's value for receiving ℓ copies of set S_i , up to a maximum of r . Write d for the maximum size of any set S_i .

The goods will be sold via a greedy auction. Bidders submit bids, in the form of a desired set T_i and a list of marginal values $b_i^1 \geq \dots \geq b_i^r$. These marginal bids are then considered in decreasing order by value.⁶ When a bid b_i^ℓ is considered, one unit of each item in T_i will be allocated to bidder i if there are remaining units of *all* items in T_i , otherwise the bid is rejected.

For payments, we will charge each bidder i an amount, per unit of set T_i received, equal to the largest bid that a shadow bidder could have placed on set T_i and been rejected. We formalize this as follows: choose an item j , fix the quantity k_{-j} of all other items, and imagine that there are infinitely many copies of item j . Denote with $\theta_j^t(b)$ the t -th highest bid for a set containing item j that would be allocated on input b . Write $\theta^k(T_i, b) = \max_{j \in T_i} \{\theta_j^{k_j+1}(b)\}$. Then bidder i 's payment will be $x_i^k(b) \cdot \theta^k(T_i, b)$. The utility of a bidder in the greedy auction, given supply profile k , is then $u_i^{n,k}(b; v_i) = v_i(x_i^k(b)) - x_i^k(b) \cdot \theta^k(T_i, b)$.

We show that even though the price of anarchy guarantee of the auction can improve from $d+1$ to d , as the market grows large, the inefficiency does not drop below d . This is in contrast to the simultaneous uniform price auction where full efficiency is recovered in the limit.⁷ This negative result holds even if we assume additional noise in the form of supply uncertainty: namely, for any fixed number, the probability for getting this number of items goes to zero as the market grows large.

Definition 6. *We say that the sequence of markets satisfies supply uncertainty if the quantity k_j^n of item j is a random variable, and moreover for any $\epsilon > 0$ there exists some $n(\epsilon)$ such that, for all $n > n(\epsilon)$, $\Pr[k_j(n) = t] < \epsilon$ for all j and all values t .*

Proposition 7. *For any $\epsilon > 0$, there exists a greedy combinatorial auction with demand uncertainty, supply uncertainty, and $k_j(n) = \Omega(n)$ for each item j , for which there exist equilibria achieving no more than a $\frac{1}{d} + \epsilon$ fraction of the optimal welfare.*

Proof. Choose R sufficiently large and $\lambda > 0$ sufficiently small, as functions of ϵ . There are d^2 items, labeled $\{a_{ij} : i, j \in [d]\}$. There are $n = R(d+1)$ buyers. For each $j \in [d]$, there are R buyers who each want the set $\{a_{ij} : i \in [d]\}$ for value 1; we will refer to these as type- j bidders. There are also R buyers who each want the set $\{a_{1j} : j \in [d]\}$ for value $1 + \lambda$; we will refer to these as type-0 bidders.

For each item a_{ij} , the number of copies is uniformly distributed in the range $\frac{r}{2} \pm \sqrt{r}$. Also, each buyer fails to arrive with probability δ , where $\delta < 0.5$ is a fixed constant. Chernoff and union bounds imply that, with probability at least $(1 - (d+1)e^{-cR})$ for some constant $c = c(\delta)$, at least $\frac{R}{2}$ buyers of each type arrive. Under this event, the optimal welfare is at least $d \cdot (\frac{R}{2} - \sqrt{R})$, for any realization of supply (by not allocating to any type-0 bidders). On the other hand, the greedy auction will allocate to type-0 bidders first, after which at most $2\sqrt{R}$ bidders of any other type can be allocated (again, under any realization of the supply). So the welfare obtained by the greedy algorithm is at most $(R/2 + \sqrt{R})(1 + \lambda) + 2d\sqrt{R}$.

⁶In the same way as in Simultaneous Uniform Price Auctions, we can handle ties in such a way that it is without loss to assume all bids are distinct.

⁷Roughly speaking, this difference is due to the greedy algorithm being inherently d -approximate even in a fractional market. Whether or not there are small-market frictions, a single bid of value $1 + \epsilon$ can “block” up to d bids of value 1. The construction in Proposition 7 makes this intuition explicit.

Accounting for the exponentially unlikely event that there are fewer than $R/2$ bidders of any type, the optimal welfare is at least $\frac{dR}{2}(1 - o(1))$ and the greedy welfare is at most $(\frac{R(1+\lambda)}{2})(1 + o(1))$, where the asymptotic notation is with respect to R growing large. Taking $\lambda < \epsilon$ and R sufficiently large completes the proof. ■

As a positive result we show that the proposed framework can be used to show that supply uncertainty alone is a sufficient condition for the greedy combinatorial auction to be d -approximately efficient in the limit. Supply uncertainty was also assumed in prior work of Swinkels [23]. We remark that for simultaneous uniform price auctions (studied in Section 4), supply uncertainty alone does not lead to full efficiency; an example appears in the full version of the paper.

Theorem 8 (Approximate efficiency in the Limit). *The greedy combinatorial auction under supply uncertainty admits a $(1, d)$ -smooth approximation in the large. In particular, if $k_j(n) = \Omega(n)$ for each item j , then the implied sequence of mechanisms is $(1, d)$ -smooth in the large, and hence achieves a $1/d$ fraction of the optimal welfare.*

One might hope to prove an analogous result to Theorem 8 under demand uncertainty as well as under supply uncertainty. However, it turns out that under demand uncertainty, a bidder's proposed approximate utility and actual utility may fail to converge; an example appears in the full version of the paper. Thus, to apply our framework to prove smoothness in the large under demand uncertainty, one would need to find an alternate approximate utility sequence.

Sketch of proof of Theorem 8. We first define a notion of approximate utility, then establish that this approximation satisfies the properties of a $(1, d)$ -smooth approximation in the large. The full proof of Theorem 8 appears in Appendix B.

To define the approximate utility, consider $\theta^k(T_i, b_{-i})$, which is the critical value for set T_i if agent i were not present. Write $X_i^k(b) = \max\{\ell : b_{i,\ell} > \theta^k(T_i, b_{-i})\}$. That is, $X_i^k(b)$ is the number of bids made by agent i that are strictly greater than $\theta^k(T_i, b_{-i})$. Then the approximate utility is:

$$U_i^{n,k}(b; v_i) = v_i(X_i^k(b)) - X_i^k(b) \cdot \theta^k(T_i, b_{-i}). \quad (6)$$

This is the utility of the original game, not taking into account the effect of bidder i 's bid upon the critical value of T_i . We denote by u_i^n and U_i^n the expected utility and approximate utility, respectively, in expectation over the distribution of k .

We must show that U_i^n satisfies the conditions of being a $(1, d)$ -smooth approximation to the critical greedy auction, in the large. The fact that U_i^n approximates u_i^n follows from the supply uncertainty: the variation in critical price calculation is smoothed over by uncertainty in the number of units of each item. The smoothness condition follows in a manner similar to the Simultaneous Uniform Price auctions: under U_i^n , each agent effectively views herself as a price-taker; the factor of d is effectively due to the approximation factor of the greedy allocation algorithm.

References

- [1] E. M. Azevedo and E. Budish. Strategy-proofness in the large. *Working paper*, 2011.

- [2] A. C. Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):pp. 122–136, 1941.
- [3] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*, 2011.
- [4] K. L. Brown. Resource allocation in competitive multiagent systems. Ph.d. thesis, Stanford University, 2003.
- [5] G. Christodoulou, A. Kovacs, and M. Schapira. Bayesian Combinatorial Auctions. In *ICALP '08 Proceedings of the 35th international colloquium on Automata, Languages and Programming*, 2008.
- [6] R. Cole and Y. Tao. The price of anarchy of large walrasian auctions. *CoRR*, abs/1508.07370, 2015.
- [7] B. de Keijzer, E. Markakis, G. Schäfer, and O. Telelis. Inefficiency of standard multi-unit auctions. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, pages 385–396, 2013.
- [8] C.-G. Esseen. On the liapunoff limit of error in the theory of probability. *Arkiv fr Matematik, Astronomi och Fysik.*, A28:1–19, 1942.
- [9] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *In Proceedings of 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [10] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *STOC*, 2013.
- [11] A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *EC'11*, 2011.
- [12] M. O. Jackson and A. M. Manelli. Approximately Competitive Equilibria in Large Finite Economies. *Journal of Economic Theory*, 77(2):354–376, December 1997.
- [13] E. Kalai. Large robust games. *Econometrica*, 72(6):1631–1665, 2004.
- [14] M. Kearns. Graphical games. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [15] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *SODA*, 2010.
- [16] E. Markakis and O. Telelis. Uniform price auctions: Equilibria and efficiency. In *SAGT*, 2012.
- [17] D. J. Roberts and A. Postlewaite. The incentives for price-taking behavior in large exchange economies. *Econometrica*, 44(1), 1976.

- [18] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, 2009.
- [19] T. Roughgarden. The price of anarchy in games of incomplete information. In *EC*, 2012.
- [20] I. Shevtsova. On the absolute constants in the Berry-Esseen type inequalities for identically distributed summands. *ArXiv e-prints*, Nov. 2011.
- [21] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations*. Cambridge University Press, 2008.
- [22] R. M. Starr. Quasi-Equilibria in Markets with Non-Convex Preferences. *Econometrica*, 37(1):25–38, January 1969.
- [23] J. M. Swinkels. Efficiency of large private value auctions. *Econometrica*, 69(1):pp. 37–68, 2001.
- [24] V. Syrgkanis. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012.
- [25] V. Syrgkanis and E. Tardos. Composable and efficient mechanisms. In *STOC*, 2013.

A Simultaneous Uniform Price Auctions

This section is dedicated to the proof of Theorem 5. We will view the simultaneous uniform price auction with random arrivals as an ex-ante mechanism $\mathcal{M}^{n,\delta}$, where the noise is endogenized in the rules of the mechanism and then we will show that mechanism $\mathcal{M}^{n,\delta}$ is $(1,1)$ -smooth in the limit. We will refer to this mechanism as *simultaneous uniform price auction with endogenous δ -noisy demand*.

Basic Notation. We first introduce some useful notation. We will denote with $u_i^n(b; v_i)$ the expected utility from a simultaneous uniform price auction where $b = (b_1, \dots, b_n)$ and b_i is a vector of marginal bids $b_i^{j,x}$, with $j \in [m]$ and $x \in [r]$, satisfying the decreasing marginal bid property, i.e. $b_i^{j,x}$ is decreasing in x . We will denote with $x_i(b)$ the allocation of player i under bid profile b in the simultaneous uniform price auction, which is a random variable (due to tie-breaking). For any vector x , we will denote with $\theta_t(x)$, the t -th highest element in x . Thus $\theta_t(b^j)$ is the t -th highest marginal bid at the uniform price auction for good j . Thus we can write:

$$u_i^n(b; v_i) = \mathbb{E} \left[v_i(x_i(b)) - \sum_{j \in [m]} x_i^j(b) \cdot \theta_{k_j^n + 1}(b^j) \right] \quad (7)$$

where expectation is taken over $x_i(b)$.

We will denote with $u_i^{n,\delta}(b; v_i)$ the expected utility of player i in the simultaneous uniform price auction with noisy arrivals. Concretely, let z_i be a $\{0, 1\}$ random variable that equals 1 with probability $1 - \delta$, indicating whether player i arrived in the market and let $z = (z_1, \dots, z_n)$. Then

$$u_i^{n,\delta}(b; v_i) = \mathbb{E}[z_i \cdot u_i^n(b \cdot z; v_i)]. \quad (8)$$

Approximate Utility. We denote with $U_i^n(b; v_i)$ an approximate utility associated with the non-noisy sequence of mechanisms, defined by the following allocation and payment rules (due to heavy notation we avoid giving an algebraic description of U_i^n and only describe it in words). We remind the reader the approximate utility is not the utility associated with any mechanism, and in fact would not be feasible for all bidders simultaneously. It is simply a construct for the proof of smoothness, and can be interpreted as an intuition for what's guiding bidder behavior. To construct the approximate utility, for each uniform price auction $j \in [m]$, let $\theta_{k_j^n+1}(b_{-i}^j)$ be the $k_j^n + 1$ -highest marginal bid excluding player i 's bids. Every marginal bid $b_i^{j,x} > \theta_{k_j^n+1}(b_{-i}^j)$ wins a unit at auction j and bids with $b_i^{j,x} = \theta_{k_j^n+1}(b_{-i}^j)$ win with some probability that follows from the random tie-breaking rule described in the beginning of the section. We will denote with $X_i(b)$ the allocation function that is implied by the above description, which is also a random variable due to the tie-breaking rule. For every unit that a player i wins at auction j , she pays $\theta_{k_j^n+1}(b_{-i}^j)$. Thus we can write the approximate utility as:

$$U_i^n(b; v_i) = \mathbb{E} \left[v_i(X_i(b)) - \sum_{j \in [m]} X_i^j(b) \cdot \theta_{k_j^n+1}(b_{-i}^j) \right] \quad (9)$$

Then denote with $U_i^{n,\delta}(b; v_i)$ an approximate utility for the noisy arrival mechanism, which is simply defined as:

$$U_i^{n,\delta}(b; v_i) = \mathbb{E} [z_i \cdot U_i^n(b \cdot z; v_i)] \quad (10)$$

(1, 1)-Smoothness of Approximate Utility. We will first show that the approximate utility $U_i^{n,\delta}$ satisfies the (1, 1)-smoothness property with respect to the sequence of mechanisms $\mathcal{M}^{n,\delta}$. To achieve this we will break it into two parts. First we will show that the approximate utility U_i^n , satisfies the (1, 1)-smoothness property with respect to the non-noisy sequence of mechanism \mathcal{M}^n . Then we show generically, that if a sequence of utility functions $U_i^n(s; v_i)$, satisfy the (λ, μ) -smoothness property with respect to a sequence of mechanisms \mathcal{M}^n , then the sequence of utility functions $U_i^{n,\delta}(s; v_i) = \mathbb{E} [z_i \cdot U_i^n(s \cdot z; v_i)]$ satisfies the (λ, μ) -smoothness property with respect to the sequence of mechanisms $\mathcal{M}^{n,\delta}$, which is the version of \mathcal{M}^n where each player arrives with probability δ . This completes the first part of the proof.

Lemma 9. U_i^n satisfies the (1, 1)-smoothness property with respect to the sequence of simultaneous uniform price auctions \mathcal{M}^n .

Proof. Consider a mechanism \mathcal{M}^n in the sequence, valuation profile $v \in \mathcal{V}^n$ and let $\text{OPT}^n(v)$ be the optimal allocation. For each player i let x_i^* denote his allocation in the welfare maximizing allocation. Consider the following deviation b_i^* for each player i : at each auction $j \in [m]$, bid $v_i(x_i^*)$ as the first $x_i^{j,*}$ marginal bids and 0 on the remaining marginal bids.

Consider any bid profile b . There are two cases: either player i wins at least his optimal allocation in his delusion in which case he gets approximate utility:

$$U_i^n(b_i^*, b_{-i}) \geq v_i(x_i^*) - \sum_{j \in [m]} x_i^{j,*} \cdot \theta_{k_j^n+1}(b_{-i}^j)$$

or otherwise, there is at least one $q \in [m]$ with $x_i^{q,*} > 0$, for which $\theta_{k_q^n+1}(b_{-i}^q) \geq v_i(x_i^*)$ and at which player i wins strictly less than $x_i^{q,*}$ units. In that case, player i 's approximate utility is at least:

$$\begin{aligned} U_i^n(b_i^*, b_{-i}; v_i) &\geq - \sum_{j \in [m]} x_i^{j,*} \cdot \theta_{k_j^n+1}(b_{-i}^j) + \theta_{k_q^n+1}(b_{-i}^q) \\ &\geq v_i(x_i^*) - \sum_{j \in [m]} x_i^{j,*} \cdot \theta_{k_j^n+1}(b_{-i}^j) \end{aligned}$$

Hence, the latter inequality holds always. Summing up the inequality for each player and observing that $\theta_{k_j^n+1}(b_{-i}^j) \leq \theta_{k_j^n+1}(b^j)$ and $\sum_{i=1}^n x_i^{j,*} \leq k_j^n$, we get:

$$\sum_{i=1}^n U_i^n(b_i^*, b_{-i}; v_i) \geq \text{OPT}^n(v) - \sum_{j \in [m]} k_j^n \cdot \theta_{k_j^n+1}(b^j)$$

Now it is easy to see that $\mathcal{R}^n(b) = \sum_{j \in [m]} k_j^n \cdot \theta_{k_j^n+1}(b^j)$, since at each uniform price auction, either k_j^n units were sold at a price of $\theta_{k_j^n+1}(b^j)$ or $\theta_{k_j^n+1}(b^j) = 0$. This completes the proof. \blacksquare

Lemma 10. *If U_i^n satisfies the (λ, μ) -smoothness property with respect to a sequence of mechanisms \mathcal{M}^n , then $U_i^{n,\delta} = \mathbb{E}_z [z_i \cdot U_i^n(s \cdot z; v_i)]$ satisfies the (λ, μ) -smoothness property with respect to the sequence of mechanisms $\mathcal{M}^{n,\delta}$.*

Proof. By smoothness of U_i^n with respect to \mathcal{M}^n we know that for any valuation vector v , there exists for each player i a deviation $s_i^*(v)$ such that for any strategy profile s :

$$\sum_{i=1}^n U_i^n(s_i^*(v), s_{-i}; v_i) \geq \lambda \text{OPT}^n(v) - \mu \mathcal{R}^n(s)$$

Observe that for any strategy profile s :

$$U_i^n(s_i^*(v \cdot z), s_{-i}; v_i) \cdot z_i \geq U_i^n(s_i^*(v \cdot z), s_{-i}; v_i \cdot z_i)$$

Observe that:

$$\begin{aligned} &\mathbb{E}_{z, \tilde{z}} [U_i^n(s_i^*(v \cdot (z_i, \tilde{z}_{-i})), s_{-i} \cdot z_{-i}; v_i) \cdot z_i] \\ &= \mathbb{E}_{z, \tilde{z}} [U_i^n(s_i^*(v \cdot \tilde{z}), s_{-i} \cdot z_{-i}; v_i) \cdot \tilde{z}_i] \\ &\geq \mathbb{E}_{z, \tilde{z}} [U_i^n(s_i^*(v \cdot \tilde{z}), s_{-i} \cdot z_{-i}; v_i \cdot \tilde{z}_i)] \end{aligned}$$

Thus, for any strategy profile s , valuation profile v and arrival vector z :

$$\begin{aligned} &\sum_{i \in [n]} \mathbb{E}_{z, \tilde{z}} [U_i^n(s_i^*(v \cdot (z_i, \tilde{z}_{-i})), s_{-i} \cdot z_{-i}; v_i) \cdot z_i] \\ &\geq \mathbb{E}_{z, \tilde{z}} \left[\sum_i U_i^n(s_i^*(v \cdot \tilde{z}), s_{-i} \cdot z_{-i}; v_i \cdot \tilde{z}_i) \right] \\ &\geq \mathbb{E}_{z, \tilde{z}} [\lambda \text{OPT}^n(v \cdot \tilde{z}) - \mu \mathcal{R}^n(s \cdot z)] \end{aligned}$$

Observe that: $\mathbb{E}_{\tilde{z}} [U_i^n(s_i^*(v \cdot (z_i, \tilde{z}_{-i})), s_{-i} \cdot z_{-i}; v_i) \cdot z_i]$ corresponds to the utility of a player under the following deviation: random sample an arrival vector \tilde{z}_{-i} , deviate assuming the arrival vector (z_i, \tilde{z}_{-i}) . This is a valid deviation for the noisy arrival mechanism $\mathcal{M}^{n,\delta}$ and hence the above inequality shows that $U_i^{n,\delta}$ satisfies the (λ, μ) -smoothness property with respect to $\mathcal{M}^{n,\delta}$. ■

Approximation. Now we move on to showing that $U_i^{n,\delta}$ approximates $u_i^{n,\delta}$ as $n \rightarrow \infty$.

Lemma 11. *For any valuation v_i and for any bid profile sequence b^n :*

$$\lim_{n \rightarrow \infty} \left\| u_i^{n,\delta}(b^n; v_i) - U_i^{n,\delta}(b^n; v_i) \right\| = 0 \quad (11)$$

Proof. We need to show that for any ϵ , there exists $n(\epsilon) < \infty$ such that for any $n > n(\epsilon)$ and for any bid profile b and for any valuation v_i :

$$\Delta = \left\| u_i^{n,\delta}(b; v_i) - U_i^{n,\delta}(b; v_i) \right\| < \epsilon$$

By triangle inequality we can lower bound the left hand side by:

$$\begin{aligned} \Delta \leq & \left\| \mathbb{E} [v_i(x_i(b \cdot z)) - v_i(X_i(b \cdot z))] \right\| + \\ & \left\| \mathbb{E} \left[\sum_{j \in [m]} X_i^j(b^j \cdot z) \theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \right. \right. \\ & \left. \left. x_i^j(b^j \cdot z) \cdot \theta_{k_j^n+1}(b^j \cdot z) \right] \right\| \end{aligned}$$

The first part of the upper bound can be upper bounded by:

$$\begin{aligned} & \left\| \mathbb{E} [v_i(x_i(b \cdot z)) - v_i(X_i(b \cdot z))] \right\| \\ & \leq H \cdot \Pr [x_i(b \cdot z) \neq X_i(b \cdot z)] \\ & \leq H \cdot \sum_{j \in [m]} \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \end{aligned}$$

The second part can also be upper bounded by the summation of the following two quantities:

$$\begin{aligned} & \left\| \mathbb{E} \left[\sum_{j \in [m]} (X_i^j(b^j \cdot z) - x_i^j(b^j \cdot z)) \theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) \right] \right\| \\ & \left\| \mathbb{E} \left[\sum_{j \in [m]} x_i^j(b^j \cdot z) \cdot (\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z)) \right] \right\| \end{aligned}$$

The first quantity is upper bounded by:

$$B \cdot r \cdot \sum_{j \in [m]} \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)]$$

since by assumption all marginal bids fall in a range $[0, B]$ and the difference in the two allocations of a player is at most r , by the r -demand assumption.

The second quantity is upper bounded by:

$$r \cdot \sum_{j \in [m]} \left\| \mathbb{E} \left[\theta_{k_j^n + 1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n + 1}(b^j \cdot z) \right] \right\|$$

since a player is allocated at most r units of each good j .

Thus, by the above reasoning, it suffices to show that there exists a finite $n(\epsilon)$ such that for any $n > n(\epsilon)$ the following two properties hold for each uniform price auction $j \in [m]$ and for any bid profile b^j

$$\begin{aligned} \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] &\leq \frac{\epsilon}{2m(B \cdot r + H)} \\ \left\| \mathbb{E} \left[\theta_{k_j^n + 1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n + 1}(b^j \cdot z) \right] \right\| &\leq \frac{\epsilon}{2rm} \end{aligned}$$

Hence we break the proof in two lemmas:

Lemma 12. *For any uniform price auction $j \in [m]$ and for any $\epsilon > 0$, there exists $n(\epsilon) < \infty$ such that for any $n > n(\epsilon)$ and for any bid profile b^j :*

$$\Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] < \epsilon \quad (12)$$

Proof. We will show that the probability converges to 0 conditional on any draw of the random tie-breaking priority order. Moreover, conditional on the tie-breaking rule it is without loss of generality to assume that all marginal bids in b^j are distinct and that there are no ties. The reason is that conditional on the tie-breaking priority rule, we can add small quantities (much smaller than the smallest difference between any two different marginal bids), to the input bids of the players, so as to simulate the exact same allocation rule as would have been achieved by the original bid profile and with the priority rule drawn (e.g. if player i was ordered first by the tie-breaker then add to all his marginal bids $n \cdot \delta$, if he was ordered second then add $(n - 1) \cdot \delta$ etc., similarly if any of his own bids are identical then add even smaller δ 's to differentiate them).

So suffices to prove that the probability goes to zero assuming that there are no two identical bids in b^j . Observe that the two allocations are different only when any of the marginal bids of player i is among the $k + 1 - r$ and the $k + 1$ highest arriving marginal bids. If a marginal bid is not among the $k + 1 - r$ and the $k + 1$ highest arriving marginal bids, then it is either not allocated by both allocation rules because it is below the $k + 1$ highest arriving bid or is allocated by both rules because it is among the $k - r$ highest arriving bids and so adding player i 's bids will not push it out of the allocation.

Let $B(b^j \cdot z; x)$ denote the number of arriving marginal bids that are strictly above x . Thus we can upper bound the desired probability by the union bound as:

$$\begin{aligned} \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \\ \leq \sum_{t=1}^r \sum_{q=k+1-r}^{k+1} \Pr [B(b^j \cdot z; b_i^{j,t}) = q] \end{aligned}$$

Now we can use the re-interpretation of a Lemma by Swinkels re-written in our terminology:

Lemma 13 (Swinkels [23]). *For any $x \in [0, B]$ and for any ϵ , there exists a $q(\epsilon) \leq \infty$ such that for any $q > q(\epsilon)$ and for any b^j :*

$$\Pr[B(b^j \cdot z; x) = q] \leq \epsilon \quad (13)$$

Thus for sufficiently large n , k_j^n is sufficiently large that each probability in the double summation can be made smaller than any ϵ . Since the summation is over a constant number of quantities, the double summation can also be made smaller than any ϵ for sufficiently large n . This completes the proof of the Lemma. ■

Lemma 14. *For any uniform price auction $j \in [m]$ and for any $\epsilon > 0$, there exists $n(\epsilon) < \infty$ such that for any $n > n(\epsilon)$ and for any bid profile b^j :*

$$\left\| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right\| \leq \epsilon \quad (14)$$

Proof. Observe that $\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) \leq \theta_{k_j^n+1}(b^j \cdot z)$. Moreover, since player i submits at most r bids, the $k_j^n + 1$ highest bid among all bids except player i 's is at least the $k_j^n + 1 + r$ highest bid among all bids including player i 's. Thus:

$$\begin{aligned} & \left\| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right\| \\ & \leq \left\| \mathbb{E} \left[\theta_{k_j^n+1+r}(b^j \cdot z) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right\| \end{aligned}$$

We will now use the following reinterpretation of a Lemma of Swinkels, which we state in our terminology:

Lemma 15 (Swinkels [23]). *For any $x, x' \in [k_j^n, k_j^n + r + 1]$ the difference of the cumulative density functions of the x -th and the x' -th highest arriving bid in a single uniform price auction converges to 0 uniformly over x, x' and b^j , as $k_j^n \rightarrow \infty$.*

Since the CDFs of the random variables $\theta_{k_j^n+1+r}(b^j \cdot z_{-i})$ and $\theta_{k_j^n+1}(b^j \cdot z)$ converge and since the two quantities are bounded in $[0, B]$, their expectations also converge and therefore there exists $n(\epsilon)$ such that for any $n > n(\epsilon)$ and for any b^j :

$$\left\| \mathbb{E} \left[\theta_{k_j^n+1+r}(b^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right\| \leq \epsilon$$

which completes the proof of the lemma. ■

Combining Lemmas 12 and 14 establishes the assertion of Lemma 11. ■

A.1 Rates of Convergence for $r = 1$

We now establish the rate of convergence claimed in Theorem 5. We begin by studying the convergence rate in the special case $r = 1$, before moving on to the general case in Appendix A.2.

Theorem 16. *If $r = 1$ and $k_j^n \geq \frac{36 \cdot \rho^2 \cdot m^2 (B+H)^2}{\epsilon^2 \delta (1-\delta)}$, then:*

$$\sum_{i=1}^n u_i^{n,\delta}(s_i^{*,n}, s_{-i}^n; v_i^n) \geq (1 - \epsilon) \text{OPT}^n(v^n) - \mathcal{R}^n(s^n) \quad (15)$$

and therefore the robust price of anarchy is at most $\frac{1}{1-\epsilon}$.

Equivalently if $k_j^n = \Omega(n)$, and for constant ρ, m, B, H, δ , the welfare at every equilibrium (BNE, CCE etc) is at least $\left(1 - o\left(\frac{1}{\sqrt{n}}\right)\right)$ of the expected optimal welfare.

The proof follows the structure of the argument from Appendix A. We will modify some of the lemmas presented in Appendix A, to make more explicit the dependency between ϵ and the parameters of the model. The modified lemmas are presented below; Theorem 16 then follows directly.

Lemma 17. *If $r = 1$, then for any $x \in [0, B]$ and for any ϵ , if $q > \frac{4}{\epsilon^2 \cdot \delta \cdot (1-\delta)}$, then for any b^j :*

$$\Pr[B(b^j \cdot z; x) = q] \leq \epsilon \quad (16)$$

Proof. Consider any bid profile b^j , consisting of one bid per player. If less than q players are bidding above x in b^j , then $\Pr[B(b^j \cdot z; x) = q] = 0$ and the theorem follows. Thus in the bid profile that maximizes the probability that we want to upper bound, there are $t \geq q$ players bidding above x . Then the probability of the event of interest is equal to the probability that exactly q of these players remain after the random deletion. Observe that the number of players among these bidders that remain after the random deletion follows a Binomial distribution of t trials, each with success probability $(1 - \delta)$, denoted as $\mathcal{B}(t, 1 - \delta)$.

By the Berry-Esseen theorem [2, 8, 20] we know that the CDF of $\mathcal{B}(t, p)$ is approximated by the CDF of the normal distribution with mean $t \cdot p$ and variance $t \cdot p \cdot (1 - p)$, with an additive error that is upper bounded by $err \leq \frac{p^2 + (1-p)^2}{2\sqrt{np(1-p)}}$. Denote with $\Phi(\cdot)$ the CDF of the standard normal distribution. If X is a random variable distributed according to $\mathcal{B}(t, p)$, then

$$\begin{aligned} \Pr[X = k] &= \Pr[X \leq k] - \Pr[X \leq k - 1] \\ &\leq \Phi\left(\frac{k - t \cdot p}{\sqrt{t \cdot p \cdot (1 - p)}}\right) - \Phi\left(\frac{k - 1 - t \cdot p}{\sqrt{t \cdot p \cdot (1 - p)}}\right) + 2 \cdot err \\ &= \frac{1}{\sqrt{2\pi}} \int_{\frac{k-1-t \cdot p}{\sqrt{t \cdot p \cdot (1-p)}}}^{\frac{k-t \cdot p}{\sqrt{t \cdot p \cdot (1-p)}}} e^{-\frac{z^2}{2}} dz + 2 \cdot err \\ &\leq \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{t \cdot p \cdot (1 - p)}} + 2 \cdot err \\ &\leq \left(\frac{1}{\sqrt{2\pi}} + p^2 + (1 - p)^2\right) \frac{1}{\sqrt{t \cdot p \cdot (1 - p)}} \\ &\leq \frac{2}{\sqrt{t \cdot p \cdot (1 - p)}} \end{aligned}$$

By the above we get that:

$$\Pr[B(b^j \cdot z; x) = q] \leq \frac{2}{\sqrt{t \cdot \delta \cdot (1 - \delta)}} \leq \frac{2}{\sqrt{q \cdot \delta \cdot (1 - \delta)}}$$

For $q \geq \frac{4}{\epsilon^2 \cdot \delta \cdot (1 - \delta)}$ the latter probability is at most ϵ as desired. \blacksquare

Lemma 18. For $r = 1$, for any uniform price auction $j \in [m]$ and for any $\epsilon > 0$, if $k_j^n \geq \frac{16B^2}{\epsilon^2 \delta (1 - \delta)}$, then for any bid profile b^j :

$$\left| \mathbb{E} \left[\theta_{k_j^n + 1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n + 1}(b^j \cdot z) \right] \right| \leq \epsilon \quad (17)$$

Proof. Observe that $\theta_{k_j^n + 1}(b_{-i}^j \cdot z_{-i}) \leq \theta_{k_j^n + 1}(b^j \cdot z)$. Moreover, since player i submits one bid, the $k_j^n + 1$ highest bid among all bids except player i 's is at least the $k_j^n + 2$ highest bid among all bids including player i 's. Thus:

$$\begin{aligned} & \left| \mathbb{E} \left[\theta_{k_j^n + 1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n + 1}(b \cdot z) \right] \right| \\ & \leq \left| \mathbb{E} \left[\theta_{k_j^n + 2}(b^j \cdot z) - \theta_{k_j^n + 1}(b^j \cdot z) \right] \right| \end{aligned}$$

Let $F_t(\cdot)$ denote the CDF of the t -th highest bid. Observe that if the number of arriving bids strictly above x are less than t , i.e., $B(b^j \cdot z; x) < t$, then both θ_t and θ_{t+1} are at most x and therefore the conditional CDFs of θ_t and θ_{t+1} evaluated at x are both 1. If $B(b^j \cdot z; x) > t + 1$ then both θ_t and θ_{t+1} are strictly above x and therefore the conditional CDFs evaluated at x are both 0. Thus the conditional CDFs differ only when $B(b^j \cdot z; x) \in [t, t + 1]$ and they differ by at most 1. Hence:

$$|F_t(x) - F_{t+1}(x)| \leq \Pr[B(b^j \cdot z; x) \in [t, t + 1]] \quad (18)$$

By Lemma 17, if $t \geq \frac{16B^2}{\epsilon^2 \cdot \delta \cdot (1 - \delta)}$, then $\Pr[B(b^j \cdot z; x) = t] \leq \frac{\epsilon}{2B}$ and $\Pr[B(b^j \cdot z; x) = t + 1] \leq \frac{\epsilon}{2B}$, so by the union bound $|F_t(x) - F_{t+1}(x)| \leq \frac{\epsilon}{B}$.

Last observe that:

$$\begin{aligned} & \mathbb{E} [\theta_t(b^j \cdot z) - \theta_{t+1}(b^j \cdot z)] \\ & = \int_0^B 1 - F_t(x) dx - \int_0^B 1 - F_{t+1}(x) dx \\ & = \int_0^B F_{t+1}(x) - F_t(x) dx \leq \epsilon \end{aligned}$$

\blacksquare

Lemma 19. If $k_j^n \geq \frac{36 \cdot m^2 (B+H)^2}{\epsilon^2 \delta (1 - \delta)}$ then for any valuation v_i and for any bid profile sequence b^n :

$$\left\| u_i^{n, \delta}(b^n; v_i) - U_i^{n, \delta}(b^n; v_i) \right\| \leq \epsilon \quad (19)$$

Proof. By same reasoning as in Lemma 11, the difference in utilities is upper bounded by the following quantity:

$$\sum_{j \in [m]} (B + H) \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \\ + \left| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right|$$

and:

$$\Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \\ \leq \Pr[B(b^j \cdot z; b_i^{j,t}) = k_j^n] + \Pr[B(b^j \cdot z; b_i^{j,t}) = k_j^n + 1].$$

By the previous lemmas, if $k_j^n \geq \frac{36 \cdot m^2 (B+H)^2}{\epsilon^2 \delta (1-\delta)}$ then:

$$\left| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right| \leq \frac{\epsilon}{3 \cdot m} \\ \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \leq \frac{2\epsilon}{3m(B+H)}$$

which subsequently gives that the utility difference is at most ϵ . ■

A.2 Convergence Rate for General r

We now present a version of Theorem 16 that generalizes to arbitrary r . The proof will follow a similar structure to that of Theorem 16.

Theorem 20. *If $k_j^n \geq \frac{16 \cdot m^2 (B+H)^2 r^8 \rho^2}{\epsilon^2 \delta (1-\delta)} + r$, then:*

$$\sum_{i=1}^n u_i^{n,\delta}(s_i^{*,n}, s_{-i}^n; v_i^n) \geq (1 - \epsilon) \text{OPT}^n(v^n) - \mathcal{R}^n(s^n) \quad (20)$$

and therefore the robust price of anarchy is at most $\frac{1}{1-\epsilon}$.

Equivalently if $k_j^n = \Omega(n)$, and for constant ρ, r, m, B, H, δ , the welfare at every equilibrium (BNE or CCE) is at least $\left(1 - o\left(\frac{1}{\sqrt{n}}\right)\right)$ of the expected optimal welfare.

Lemma 21. *For any $x \in [0, B]$ and for any ϵ , if $q > \frac{4r^2}{\epsilon^2 \cdot \delta \cdot (1-\delta)}$, then for any b^j :*

$$\Pr[B(b^j \cdot z; x) = q] \leq \epsilon \quad (21)$$

Proof. Consider any bid profile b^j , consisting of at most r bids per player. For $h \in [1, r]$ let N_h denote the subset of players that under b^j , they submit h bids above x and denote with $n_h = |N_h|$. Observe that there must exist at least one $h^* \in [1, r]$ such that $n_{h^*} \geq \frac{q}{r^2}$. Otherwise, we get: $\sum_{h=1}^r n_h \cdot r < q$ and therefore, there are in total less than q bids above x . Hence, the probability we want to upper bound is 0.

Let $Z_{-N_{h^*}}$ denote the number of arriving bids from players outside of N_h and $Z_{N_{h^*}}$ the number of arriving bids from players in N_h . By the independent arrival assumption, conditional on the bid profile b^j , these two random variables are independent. Thus:

$$\begin{aligned} \Pr[B(b^j \cdot z; x) = q] &= \sum_{q'=1}^q \Pr[Z_{-N_{h^*}} = q'] \cdot \Pr[Z_{N_{h^*}} = q - q'] \\ &\leq \max_{z \in [1, q]} \Pr[Z_{N_{h^*}} = z] \end{aligned}$$

Let X denote the number of players from N_{h^*} that end up arriving. Observe that: $\Pr[Z_{N_{h^*}} = z] = \Pr[X = \frac{z}{h^*}]$, if z is a multiple of h^* and 0 otherwise. Thus:

$$\Pr[B(b^j \cdot z; x) = q] \leq \max_{x \in [1, \lceil q/h^* \rceil]} \Pr[X = x]$$

If we denote with $\mathcal{B}(t, p)$ the binomial distribution of n trials each with success probability p , then observe that $X \sim \mathcal{B}(n_{h^*}, 1 - \delta)$.

By the Berry-Esseen theorem [2, 8, 20] we know that the CDF of $\mathcal{B}(t, p)$ is approximated by the CDF of the normal distribution with mean $t \cdot p$ and variance $t \cdot p \cdot (1 - p)$, with an additive error that is upper bounded by $err \leq \frac{p^2 + (1-p)^2}{2\sqrt{np(1-p)}}$. Denote with $\Phi(\cdot)$ the CDF of the standard normal distribution. If X is a random variable distributed according to $\mathcal{B}(t, p)$, then

$$\begin{aligned} \Pr[X = k] &= \Pr[X \leq k] - \Pr[X \leq k - 1] \\ &\leq \Phi\left(\frac{k - t \cdot p}{\sqrt{t \cdot p \cdot (1 - p)}}\right) - \Phi\left(\frac{k - 1 - t \cdot p}{\sqrt{t \cdot p \cdot (1 - p)}}\right) + 2 \cdot err \\ &= \frac{1}{\sqrt{2\pi}} \int_{\frac{k-1-t \cdot p}{\sqrt{t \cdot p \cdot (1-p)}}}^{\frac{k-t \cdot p}{\sqrt{t \cdot p \cdot (1-p)}}} e^{-\frac{z^2}{2}} dz + 2 \cdot err \\ &\leq \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{t \cdot p \cdot (1 - p)}} + 2 \cdot err \\ &\leq \left(\frac{1}{\sqrt{2\pi}} + p^2 + (1 - p)^2\right) \frac{1}{\sqrt{t \cdot p \cdot (1 - p)}} \\ &\leq \frac{2}{\sqrt{t \cdot p \cdot (1 - p)}} \end{aligned}$$

By the above we get that:

$$\Pr[B(b^j \cdot z; x) = q] \leq \frac{2}{\sqrt{n_{h^*} \cdot \delta \cdot (1 - \delta)}} \leq \frac{2r}{\sqrt{q \cdot \delta \cdot (1 - \delta)}}$$

For $q \geq \frac{4r^2}{\epsilon^2 \cdot \delta \cdot (1 - \delta)}$ the latter probability is at most ϵ as desired. \blacksquare

Lemma 22. *For any uniform price auction $j \in [m]$ and for any $\epsilon > 0$, if $k_j^n \geq \frac{4B^2 r^4}{\epsilon^2 \delta (1 - \delta)}$, then for any bid profile b^j :*

$$\left| \mathbb{E} \left[\theta_{k_j^n + 1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n + 1}(b^j \cdot z) \right] \right| \leq \epsilon \quad (22)$$

Proof. Moreover, since player i submits at most r bids, the $k_j^n + 1$ highest bid among all bids except player i 's is at least the $k_j^n + 1 + r$ highest bid among all bids including player i 's. Thus:

$$\begin{aligned} & \left| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b \cdot z) \right] \right| \\ & \leq \left| \mathbb{E} \left[\theta_{k_j^n+1+r}(b^j \cdot z) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right| \end{aligned}$$

Let $F_t(\cdot)$ denote the CDF of the t -th highest bid. Observe that if the number of arriving bids strictly above x are less than t , i.e., $B(b^j \cdot z; x) < t$, then the conditional CDFs of t and $t + r$ highest bid evaluated at x are both 1. If $B(b^j \cdot z; x) > t + r$ then the conditional CDFs evaluated at x are both 0. Thus the conditional CDFs differ only when $B(b^j \cdot z; x) \in [t, t + r]$ and they differ by at most 1. Hence:

$$|F_t(x) - F_{t+r}(x)| \leq \Pr[B(b^j \cdot z; x) \in [t, t + r]] \quad (23)$$

By Lemma 17, if $t \geq \frac{4B^2r^4}{\epsilon^2\delta(1-\delta)}$, then for all $x \in [t, t + r]$, $\Pr[B(b^j \cdot z; x) = x] \leq \frac{\epsilon}{Br}$, and by the union bound: $|F_t(x) - F_{t+r}(x)| \leq \frac{\epsilon}{B}$.

Last observe that:

$$\begin{aligned} & \mathbb{E} [\theta_t(b^j \cdot z) - \theta_{t+r}(b^j \cdot z)] \\ & = \int_0^B 1 - F_t(x) dx - \int_0^B 1 - F_{t+r}(x) dx \\ & = \int_0^B F_{t+r}(x) - F_t(x) dx \leq \epsilon \end{aligned}$$

■

Lemma 23. If $k_j^n \geq \frac{16 \cdot m^2(B+H)^2r^8}{\epsilon^2\delta(1-\delta)} + r$ then for any valuation v_i and for any bid profile sequence b^n :

$$\left\| u_i^{n,\delta}(b^n; v_i) - U_i^{n,\delta}(b^n; v_i) \right\| = \epsilon \quad (24)$$

Proof. By same reasoning as in Lemma 11, the difference in utilities is upper bounded by the following quantity:

$$\begin{aligned} & r \cdot \left(\sum_{j \in [m]} (B + H) \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \right. \\ & \quad \left. + \left| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right| \right) \end{aligned}$$

and:

$$\begin{aligned} & \Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \\ & \leq \sum_{t=1}^r \sum_{q=k_j^n+1-r}^{k_j^n+1} \Pr[B(b^j \cdot z; b_i^{j,t}) = q] \\ & \leq r^2 \max_{x \in [0, B], q \in [k_j^n+1-r, k_j^n+1]} \Pr[B(b^j \cdot z; x) = q] \end{aligned}$$

By the previous lemmas, if $k_j^n \geq \frac{16 \cdot m^2 (B+H)^2 r^8}{\epsilon^2 \delta (1-\delta)} + r$ then:

$$\Pr [x_i^j(b^j \cdot z) \neq X_i^j(b^j \cdot z)] \leq r^2 \frac{\epsilon}{2mr^3(B+H)}$$

$$\left| \mathbb{E} \left[\theta_{k_j^n+1}(b_{-i}^j \cdot z_{-i}) - \theta_{k_j^n+1}(b^j \cdot z) \right] \right| \leq \frac{\epsilon}{2 \cdot r \cdot m}$$

which subsequently gives that the utility difference is at most ϵ . ■

B Greedy Combinatorial Auctions

This section is dedicated to the proof of Theorem 8. We begin by reviewing the proof sketch from Section 5.

Our approach will be to define a notion of approximate utility, then establish that this approximation satisfies the properties of a $(1, d)$ -smooth approximation in the large. To define the approximate utility, consider $\theta^k(T_i, b_{-i})$, which is the critical value for set T_i if agent i were not present. Write $X_i^k(b) = \max\{\ell : b_{i,\ell} > \theta^k(T_i, b_{-i})\}$. That is, $X_i^k(b)$ is the number of bids made by agent i that are strictly greater than $\theta^k(T_i, b_{-i})$. Then the approximate utility is:

$$U_i^{n,k}(b; v_i) = v_i(X_i^k(b)) - X_i^k(b) \cdot \theta^k(T_i, b_{-i}).$$

This is the utility of the original game, not taking into account the effect of player i 's bid upon the critical value of T_i . We denote by u_i^n and U_i^n the expected utility and approximate utility, respectively, in expectation over the distribution of k .

(1, d)-Smoothness of Approximate Utility. We will first show that the approximate utility U_i^n satisfies the conditions of being a $(1, d)$ -smooth approximation to the critical greedy auction, in the large. We do this in two steps. We first show that U_i^n satisfies the smoothness condition with respect to the critical greedy auction, then show that it approximates the utility of the original game.

Lemma 24. *For each n , U_i^n satisfies the $(1, d)$ -smoothness property with respect to the greedy critical price auction.*

Proof. Fix valuation profile v , and let $x^{*,k}$ denote the welfare-optimal allocation for supply k . We will consider the utility of agent i when declaring his true valuation v_i . We have

$$\begin{aligned} U_i^n(v_i, b_{-i}; v_i) &= \mathbb{E}_k \left[v_i(X_i^k(v_i, b_{-i})) \right. \\ &\quad \left. - X_i^k(v_i, b_{-i}) \cdot \theta^k(T_i, b_{-i}) \right] \\ &= \mathbb{E}_k \left[\sum_{\ell=1}^r (v_i(\ell) - v_i(\ell-1) - \theta^k(T_i, b_{-i}))^+ \right] \\ &\geq \mathbb{E}_k \left[\sum_{\ell=1}^r (v_i(\ell) - v_i(\ell-1) - \theta^k(T_i, b))^+ \right] \\ &\geq \mathbb{E}_k \left[v_i(x_i^{*,k}) - x_i^{*,k} \cdot \theta^k(T_i, b) \right]. \end{aligned}$$

Taking a sum over all i and applying linearity of expectation, we have

$$\sum_i U_i^n(v_i, b_{-i}; v_i) \geq \text{OPT}^n(v) - \mathbb{E}_k \left[\sum_i x_i^{*,k} \cdot \theta^k(T_i, b) \right].$$

Since $\theta^{k(n)}(T_i, b) = \max_{j \in T_i} \theta_j^k(b) \leq \sum_{j \in T_i} \theta_j^k(b)$, we have

$$\begin{aligned} \sum_i U_i^n(v_i, b_{-i}; v_i) &\geq \text{OPT}^n(v) - \mathbb{E}_k \left[\sum_j \theta_j^k(b) \sum_{i: T_i \ni j} x_i^{*,k} \right] \\ &\geq \text{OPT}^n(v) - \mathbb{E}_k \left[\sum_j \theta_j^k(b) \cdot k_j \right] \\ &\geq \text{OPT}^n(v) - d \cdot \mathbb{E}_k \left[\sum_i x_i^k(b) \cdot \theta^k(T_i, b) \right] \\ &= \text{OPT}^n(v) - d \cdot \mathcal{R}^n(b) \end{aligned}$$

as required, where in the last inequality we made use of the fact that $\theta^k(T_i, b) \geq \frac{1}{d} \sum_{j \in T_i} \theta_j^k(b)$, plus the fact that $\theta_j^{k(n)}(b) = 0$ if not all copies of item j are allocated in $x^k(b)$. \blacksquare

Approximation. Now we show that U_i^n approximates u_i^n as n grows large.

Lemma 25. *For any valuation v_i and for any bid profile sequence b^n :*

$$\lim_{n \rightarrow \infty} \|u_i^n(b^n; v_i) - U_i^n(b^n; v_i)\| = 0 \quad (25)$$

Proof. This proof closely follows the proof of Lemma 11. Our goal is to find an upper bound on $\|u_i^n(b^n; v_i) - U_i^n(b^n; v_i)\|$. Applying the triangle inequality to the definition of u_i^n and U_i^n , we have

$$\begin{aligned} &\|u_i^n(b^n; v_i) - U_i^n(b^n; v_i)\| \\ &\leq \|\mathbb{E}_k[v_i(x_i^k(b)) - v_i(X_i^k(b))]\| \\ &\quad + \|\mathbb{E}_k[x_i^k(b) \cdot \theta^k(T_i, b) - X_i^k(b) \cdot \theta^k(T_i, b_{-i})]\| \end{aligned}$$

We'll bound separately each of the two terms on the right hand side. The first can be bounded by

$$\|\mathbb{E}_k[v_i(x_i^k(b)) - v_i(X_i^k(b))]\| \leq H \cdot \Pr[x_i^k(b) \neq X_i^k(b)]$$

For the second term, we have

$$\begin{aligned} &\|\mathbb{E}_k[x_i^k(b) \cdot \theta^k(T_i, b) - X_i^k(b) \cdot \theta^k(T_i, b_{-i})]\| \\ &\leq \|\mathbb{E}_k[x_i^k(b)(\theta^k(T_i, b) - \theta^k(T_i, b_{-i}))]\| \\ &\quad + \|\mathbb{E}_k[(X_i^k(b) - x_i^k(b))\theta^k(T_i, b)]\| \\ &\leq r \cdot \|\mathbb{E}_k[\theta^k(T_i, b) - \theta^k(T_i, b_{-i})]\| + H \cdot r \cdot \Pr[x_i^k(b) \neq X_i^k(b)] \end{aligned}$$

Given these bounds, it suffices to show that, for all $\epsilon > 0$, there exists an $n(\epsilon)$ such that, for all $n > n(\epsilon)$, we have

$$\Pr[x_i^k(b) \neq X_i^k(b)] < \epsilon$$

and

$$\|\mathbb{E}_k[\theta^k(T_i, b) - \theta^k(T_i, b_{-i})]\| < \epsilon.$$

We will complete the proof by establishing these bounds in separate lemmas.

Lemma 26. *For all $\epsilon > 0$ there exists $n(\epsilon)$ such that for all $n > n(\epsilon)$, and all i and b , $\Pr[x_i^k(b) \neq X_i^k(b)] < \epsilon$.*

Proof. By the union bound, we have

$$\Pr[x_i^k(b) \neq X_i^k(b)] \leq \sum_j \sum_{\ell=1}^r \Pr[\theta_j^k(b) > b_{i,\ell} \geq \theta_j^k(b_{-i})].$$

It therefore suffices to bound $\Pr[\theta_j^k(b) > b_{i,\ell} \geq \theta_j^k(b_{-i})]$. Fix the quantities of all items but j , and suppose there are infinitely many units of item j . Among the marginal bids in b , consider the winning bids for sets containing j ; let $(z_1 \geq z_2 \geq \dots)$ be those bids in decreasing order. Note then that $\theta_j^k(b) = z_{k_j+1}$, and $\theta_j^k(b_{-i}) \geq z_{k_j+r+1}$ (as agent i is allocated at most r copies of item j).

Let ℓ be the unique index such that $z_\ell > b_{i,\ell} \geq z_{\ell+1}$. We then have

$$\Pr[\theta_j^k(b) > b_{i,\ell} \geq \theta_j^k(b_{-i})] \leq \Pr[\ell + 1 \leq k_j \leq \ell + r].$$

The union bound combined with the definition of supply uncertainty implies that, for sufficiently large n , this probability is at most $\epsilon \cdot r$. Taking an appropriate choice of ϵ completes the proof. \blacksquare

Lemma 27. *For all $\epsilon > 0$ there exists $n(\epsilon)$ such that, for all $n > n(\epsilon)$, and for all i, j , and b ,*

$$|\mathbb{E}_k [\theta_j^k(b)] - \mathbb{E}_k [\theta_j^k(b_{-i})]| < \epsilon. \quad (26)$$

Proof. Define values $(z_1 \geq z_2 \geq \dots)$ as in Lemma 26. Recalling that $\theta_j^k(b) = z_{k_j+1}$ and $\theta_j^k(b_{-i}) \geq z_{k_j+r+1}$, we have

$$|\mathbb{E}_k [\theta_j^k(b)] - \mathbb{E}_k [\theta_j^k(b_{-i})]| \leq \sum_{\ell \geq 1} (z_\ell - z_{\ell+r}) \cdot \Pr[k_j = \ell].$$

Supply uncertainty implies that, for sufficiently large n , $\Pr[k_j = \ell] < \epsilon$ for all ℓ and hence

$$\begin{aligned} |\mathbb{E}_k [\theta_j^k(b)] - \mathbb{E}_k [\theta_j^k(b_{-i})]| &< \sum_{\ell \geq 1} (z_\ell - z_{\ell+r}) \cdot \epsilon \\ &\leq \sum_{\ell=1}^r z_\ell \cdot \epsilon < rH\epsilon. \end{aligned}$$

Taking an appropriate choice of ϵ therefore completes the proof. \blacksquare

Applying Lemma 26 and Lemma 27 then completes the proof of Lemma 25. \blacksquare