The Price of Anarchy in Games of Incomplete Information*

Tim Roughgarden†

April 18, 2014

Abstract

We define smooth games of incomplete information. We prove an "extension theorem" for such games: price of anarchy bounds for pure Nash equilibria for all induced full-information games extend automatically, without quantitative degradation, to all mixed-strategy Bayes-Nash equilibria with respect to a product prior distribution over players' preferences. We also note that, for Bayes-Nash equilibria in games with correlated player preferences, there is no general extension theorem for smooth games.

We give several applications of our definition and extension theorem. First, we show that many games of incomplete information for which the price of anarchy has been studied are smooth in our sense. Our extension theorem unifies much of the known work on the price of anarchy in games of incomplete information. Second, we use our extension theorem to prove new bounds on the price of anarchy of Bayes-Nash equilibria in routing games with incomplete information.

1 Introduction

Every student of game theory learns early and often that *equilibria are inefficient.* Such inefficiency is ubiquitous, and is present in many real-world situations and for many different reasons. For example: Prisoner's Dilemma-type scenarios; uninternalized negative externalities in the tragedy of the commons and in games with congestion effects; uninternalized positive externalities with a public good or with network effects; a failure to coordinate in team games; and so on.

Research over the past fifteen years has provided an encouraging counterpoint to this widespread equilibrium inefficiency: in a number of interesting application domains, game-theoretic equilibria *provably approximate* the optimal outcome. That is, the *price of anarchy* — the worst-case ratio between the objective function value of an equilibrium and of an optimal outcome — is close to 1 in many interesting games.

The price of anarchy was first studied in network models (see [31, Chapters 17-21] for an overview), but the list of applications studied now spans the gamut from health care [23] to basketball [37]. Essentially all initial work on the price of anarchy studied *full-information* games, where all players' payoffs are common knowledge. Now that the study of equilibrium inefficiency has grown in scope and considers strategically interesting auctions and mechanisms — we give several

*A preliminary version of this paper appeared in the Proceedings of the 13th EC conference, June 2012.

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305.
This research was supported in part by NSF Awards CCF-1016885 and CCF-1215965, an AFOSR MURI grant, and an ONR PECASE Award. Email: `tim@cs.stanford.edu`.

concrete examples in Section 2 — there is presently a well-motivated focus on the price of anarchy in games of *incomplete information* [21], where players are uncertain about each others' payoffs. The goal of this paper is to develop a useful general theory for bounding the price of anarchy in such games.

1.1 Executive Summary: Price of Anarchy Bounds for Bayes-Nash Equilibria via Extension Theorems (i.e., Without the Pain)

Pure-strategy Nash equilibria — where each player deterministically picks a single action — are often easy to reason about. Or at least, they are easier to analyze than their more general cousins like mixed-strategy Nash equilibria (where players can randomize) and Bayes-Nash equilibria (where players don't even know with certainty what game they're playing in).

For this reason, the price of anarchy of a game is often analyzed, at least initially, only for the game's pure-strategy Nash equilibria. But as much as he or she might want to, the conscientious researcher cannot stop there. Performance guarantees for more general classes of equilibria are crucial for several reasons: pure-strategy Nash equilibria do not always exist (like in "Matching Pennies"); they can be intractable to compute, even when they are guaranteed to exist [14]; and even when efficiently computable by a centralized algorithm, they can elude natural learning dynamics [38]. Finally, a fundamental assumption behind the Nash equilibrium concept is that all players' preferences are common knowledge, and this assumption is violated in most auction and mechanism design contexts, where participants have private information.

Many researchers dutifully extended their (or their predecessors') price of anarchy bounds beyond pure-strategy Nash equilibria to more general concepts. Early on, researchers emphasized full-information equilibrium concepts that extend Nash equilibria (see [7, 8, 10, 19, 24, 27, 35, 42] for a number of examples); as Section 2 discusses, recent work has focused on Bayes-Nash equilibria in games of incomplete information.

Extending price of anarchy bounds beyond pure Nash equilibria is an extremely well motivated activity, but it is also potentially dispiriting, for two reasons. The first is that the analysis generally becomes more complex, with one or more unruly probability distributions obfuscating the core argument. The second is that enlarging the set of permissible equilibria can only degrade the price of anarchy, which is a worst-case measure. Thus the work can be difficult, and the news can only be bad.

Can we obtain price of anarchy bounds for equilibrium concepts more general than pure-strategy Nash equilibria without doing any additional work? Ideal would be an *extension theorem* that could be used in the following "black-box" way (Figure 1): (1) prove a bound on the price of anarchy of pure-strategy Nash equilibria of a game; (2) invoke the extension theorem to conclude immediately that the exact same approximation bound applies to some more general equilibrium concept. Such an extension theorem would dodge both potential problems with generalizing price of anarchy bounds beyond pure Nash equilibria — no extra work, and no loss in the approximation guarantee.

Since there are plenty of games in which (say) the worst mixed-strategy Nash equilibrium is worse than the worst pure-strategy Nash equilibrium (like "Chicken"), there is no universally applicable extension theorem of the above type. The next-best thing would be an extension theorem that applies under some conditions — perhaps on the game, or perhaps *on the method of proof used* to bound the price of anarchy of pure Nash equilibria. If such an extension theorem existed, it would reduce proving price of anarchy bounds for general equilibrium concepts to proving such bounds *in a prescribed way* for pure-strategy Nash equilibria.

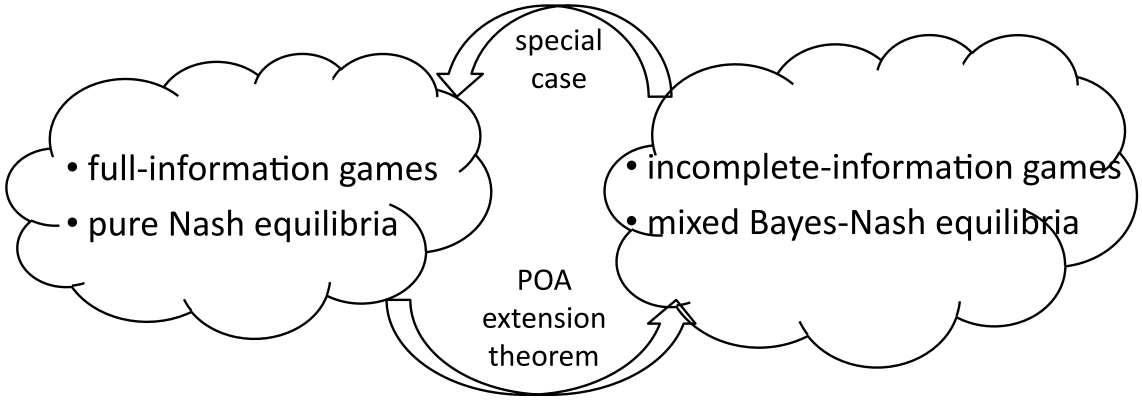

Figure 1: How to use an extension theorem. The first step is to bound the price of anarchy for the special case of pure Nash equilibria of full-information games (subject to conditions detailed in the main text). The second step is to apply an extension theorem to obtain the same approximation guarantee for mixed Bayes-Nash equilibria of incomplete-information games.

The first example of such an extension theorem was given in [35], for full-information games.[1] The key concept in [35] is that of a *smooth* game. We give an intuitive explanation here and a formal definition in Section 2.4. Conceptually, a full-information game is smooth if the objective function value of every pure-strategy Nash equilibrium $\mathbf{a}$ can be bounded using the following minimal recipe:

1. Choose an optimal outcome $\mathbf{a}^*$ of the game.
2. Invoke the Nash equilibrium hypothesis once per player, to derive that each player i's payoff in the Nash equilibrium $\mathbf{a}$ is at least as high as if it played a_i^* instead. *Do not use the Nash equilibrium hypothesis again in the rest of the proof.*
3. Use the inequalities of the previous step, possibly in conjunction with other properties of the game's payoffs, to prove that the objective function value of $\mathbf{a}$ is at least some fraction of that of $\mathbf{a}^*$.

Many interesting price of anarchy bounds follow from "smoothness proofs" of this type. The main extension theorem in [35] states that every price of anarchy bound proved in this way — seemingly only for pure Nash equilibria — automatically extends to every mixed-strategy Nash equilibrium, correlated equilibrium [2], and coarse correlated equilibrium [20, 28] of the game.

This paper presents a general extension theorem for games of incomplete information, where players' private preferences are drawn independently from prior distributions that are common knowledge. This extension theorem reduces, in a "black-box" fashion, the task of proving price of anarchy bounds for mixed-strategy Bayes-Nash equilibria to that of proving such bounds in a prescribed way for pure-strategy Nash equilibria in every induced game of full information (after conditioning on all players' preferences). With this extension theorem, one can prove equilibrium guarantees for games of incomplete information without ever leaving the safe confines of full-information games.

[1] See [9, 29, 36, 40] for subsequent refinements.

We conclude this section with an overview of the main points of this paper.[2]

1. We define smooth games of incomplete information. The definition is slightly stronger than requiring that every induced full-information game is smooth.

2. We prove an extension theorem for smooth games of incomplete information: price of anarchy bounds for pure Nash equilibria for all induced full-information games extend automatically to all mixed-strategy Bayes-Nash equilibria with respect to a product prior distribution over players' preferences.

3. We show that many games of incomplete information for which the price of anarchy has been studied are smooth in our sense. Thus our extension theorem unifies much of the known work on the price of anarchy in games of incomplete information.

4. We use our extension theorem to prove new bounds on the price of anarchy of Bayes-Nash equilibria in routing games with incomplete information.

5. We note that for Bayes-Nash equilibria in games with correlated player preferences, there is no general extension theorem for smooth games.

1.2 Organization of Paper

Section 2 reviews games of incomplete information, smooth full-information games, and several motivating examples in auctions and routing games. Section 3 defines smooth games of incomplete information and proves our extension theorems. Section 4 recovers known results for auctions and proves new results for routing games with incomplete information as special cases of our abstract framework. Section 5 offers conclusions.

2 Preliminaries and Examples

Section 2.1 defines games of incomplete information, Bayes-Nash equilibria, and the price of anarchy. This section can be skipped by the expert. Sections 2.2 and 2.3 describe in detail several games of incomplete information from the domains of mechanism design and routing, respectively. These concrete examples are useful to keep in mind throughout the abstract development in Section 3. Section 2.4 reviews smooth games and extension theorems for full-information games; this section provides context but can be skipped without much loss.

2.1 The Price of Anarchy in Games of Incomplete Information

In a game of incomplete information, there are n players. Player i has a *type space* $\mathcal{T}_i$ and an *action space* $\mathcal{A}_i$. We write $\mathcal{T} = \mathcal{T}_1 \times \cdots \times \mathcal{T}_n$ and $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$. We assume that the type vector $\mathbf{t}$ is drawn from a distribution $\mathbf{F}$ that is common knowledge. The distribution $\mathbf{F}$ may or may not be a product distribution — that is, players' types may or may not be stochastically independent. The *payoff* $u_i(t_i; \mathbf{a})$ of player i is determined by its type and by the actions $\mathbf{a}$ chosen by all of the players. For example, in a first-price auction, the actions (bids) determine whether or not a given player wins, and the price if it does win; its value for winning is given by its type.

[2] Some of our results were also obtained, subsequently but independently, by Syrgkanis [39]. There are also results in [39], developed further in [40], for some pay-as-bid mechanisms [22, 25] that are not considered here.

The point of the machinery above is to model situations where each player is uncertain about what the other players want, and is therefore also uncertain about what they will do. For example, suppose you are participating in a first-price auction. How should you bid? The answer depends on your beliefs about what the others' are bidding, which depends both on their types (i.e., valuations) and also on their bidding strategies (i.e., the player's bid given its valuation). When discussing equilibria, we assume that each player knows the others' bidding strategies, but is uncertain about their types.

In more detail, a *strategy* σ_i for player i is a function from types $\mathcal{T}_i$ to probability distributions over $\mathcal{A}_i$, with the semantics "when my type is t_i I will play the mixed strategy $\sigma_i(t_i)$". A strategy is *pure* if, for each type t_i, $\sigma_i(t_i)$ is a point mass on one action. A strategy profile σ is a *Bayes-Nash equilibrium* if, for every player i, type $t_i \in \mathcal{T}_i$, and action $a_i' \in \mathcal{A}_i$,

$$\mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{F}_{-i}{}^{(t_i)}}\left[\mathbf{E}_{\mathbf{a} \sim \sigma(\mathbf{t})}[u_i(t_i; \mathbf{a})]\right] \geq \mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{F}_{-i}{}^{(t_i)}}\left[\mathbf{E}_{\mathbf{a}_{-i} \sim \sigma_{-i}(\mathbf{t}_{-i})}\left[u_i(t_i; (a_i', \mathbf{a}_{-i}))\right]\right], \tag{1}$$

where $\mathbf{F}_{-i}{}^{(t_i)}$ denotes the distribution induced by $\mathbf{F}$ on $\mathcal{T}_{-i}$ after conditioning on t_i. Inequality (1) simply says that every (risk-neutral) player always plays a best response given all of the available information to it — the facts that its type is t_i, that other players' types are consequently distributed according to $\mathbf{F}_{-i}{}^{(t_i)}$, and that the other players are using the strategies σ_{-i}. If the distribution $\mathbf{F}$ is a point mass, so that there is no uncertainty about players' types, then the game is equivalent to a full-information game, and Bayes-Nash equilibria are simply Nash equilibria. In this sense, fixing a type vector $\mathbf{t}$ induces a full-information game.

Our motivating applications require a superficially more general model in which a player's action set $\mathcal{A}_i$ depends on its type t_i. We say that such actions are *feasible* for t_i. One canonical motivation for making some actions infeasible in a type-dependent way is to disallow "bluffing strategies" in second-price-type auctions. Another is in routing games, where the type-dependent origin and destination of a player determines which paths it can use. Infeasible strategies can be modeled by setting the player's payoff $u_i(t_i; \mathbf{a})$ to negative infinity whenever a_i is infeasible for t_i. We always assume that a player has at least one feasible action, no matter what its type is.

We now define the price of anarchy of a game of incomplete information. Let $W(\mathbf{t}; \mathbf{a})$ denote a non-negative objective function defined on the outcomes of the game (for each type profile), such as the sum of players' payoffs. Let $OPT(\mathbf{t})$ denote a profile of actions feasible for $\mathbf{t}$ that optimizes the objective function $W(\mathbf{t}; \mathbf{a})$ over all such profiles. The *price of anarchy* of the game is the worst-case, over the Bayes-Nash equilibria σ of the game, of the expected objective function value of a Bayes-Nash equilibrium and of an optimal outcome:

$$\frac{\mathbf{E}_{\mathbf{t} \sim \mathbf{F}}\left[\mathbf{E}_{\mathbf{a} \sim \sigma(\mathbf{t})}[W(\mathbf{t}; \mathbf{a})]\right]}{\mathbf{E}_{\mathbf{t} \sim \mathbf{F}}[W(\mathbf{t}; OPT(\mathbf{t}))]}. \tag{2}$$

We are particularly interested in bounds on the price of anarchy that are independent of the distribution $\mathbf{F}$ over types. To make this formal, by an (incomplete-information) *game structure*, we mean all of the ingredients of a game of incomplete information, save for the distribution $\mathbf{F}$. By definition, the *independent POA (iPOA)* of a game structure is the worst-case POA of a game of incomplete information induced by a product distribution F. The *correlated POA (cPOA)* is the worst-case POA induced by an arbitrary distribution F. Obviously, the cPOA can only be worse than the iPOA.

Product distributions include fixed type vectors as degenerate special cases, so the iPOA can only be worse than the worst-case full-information POA corresponding to the game structure. Put another way, the best-case scenario for bounding the iPOA is to extend a POA bound that applies to every induced full-information game.

2.2 Motivating Examples from Mechanism Design

Mechanisms for allocating goods or resources furnish relevant and technically interesting examples of games of incomplete information. We consider mechanisms that comprise an *allocation rule* $\mathbf{x}$ and *payment rule* $\mathbf{p}$, which map bid vectors to allocation vectors (i.e., who gets what) and payment vectors (i.e., who pays what), respectively. The type of a player i is a *valuation* v_i, which specifies the player's value for each allocation that it might receive. The action of a player is a bid b_i. The (quasi-linear) payoff of a player is determined by its type and the computed allocation and payment: $u_i(v_i;\mathbf{b}) = v_i(x_i(\mathbf{b})) - p_i(\mathbf{b})$. Such a description of valuation spaces, feasible bid spaces, an allocation rule, and a payment rule defines a game structure in the sense above — once a distribution over types is specified, we have all of the ingredients of a game of incomplete information. The most commonly studied objective in such settings is welfare-maximization, where the *welfare* of an allocation is the sum of players' values for what they were allocated. Thus, with the notation above, a bid vector $\mathbf{b}$ yields an allocation with welfare $\sum_{i=1}^{n} v_i(x_i(\mathbf{b}))$.

We next discuss three well-studied examples; the list is illustrative, not exhaustive. All concern welfare-maximization mechanism design problems, but the specifics vary widely. In addition to reinforcing the concepts and notation above, these examples serve as interesting and diverse special cases of the general theory of POA bounds developed in Section 3.

Example 2.1 (The Generalized Second Price Auction) In the standard single-shot sponsored search auction model, there are k slots with associated click-through rates $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$. The private type of player i is its valuation v_i per click. The feasible action (or bid) space for player i with type v_i is $[0, v_i]$.[3]

The payoff to player i when it is assigned slot j with a total payment of p is $v_i \cdot \alpha_j - p$.

In the Generalized Second Price (GSP) auction, the allocation rule $\mathbf{x}$ assigns the ith highest bidder to the ith highest slot, for each $i = 1, 2, \ldots, k$. The payment rule $\mathbf{p}$ charges the bidder in the ith slot the $(i+1)$th highest bid $b_{(i+1)}$ per click, for an overall payment of $\alpha_i b_{(i+1)}$.

The GSP auction was first proposed as a model of sponsored search auctions in [13, 41]. The POA in the full-information and incomplete information versions of the GSP auction was first analyzed in [32] and [26], respectively. Currently, the best lower bound known for the cPOA in this model is 0.342 — that is, for every (joint) distribution of player valuations, the expected welfare of every Bayes-Nash equilibrium of the GSP auction is at least 0.342 times the expected maximum-possible welfare [9]. Better bounds are known for various full-information equilibrium concepts [9].

Example 2.2 (Simultaneous Second-Price Auctions) In a combinatorial auction, there are m goods for sale. The private type of player i is a valuation function v_i that specifies its value

[3] Some kind of restriction on bidding is necessary for non-trivial price of anarchy guarantees for mechanisms with the "critical-bid" payment rules that we consider. Even in the simple Vickrey auction, there are Nash equilibria with arbitrarily bad welfare. (Consider two bidders with known valuations 1 and 0, who bid 0 and 1, respectively.) Many authors have, by necessity, made and discussed such "no overbidding" or "conservative bidding" assumptions; see [12, 25, 32] for further details. More generally, such assumptions can be parameterized by the largest factor by which bids can exceed valuations [5]; all of the price of anarchy guarantees we discuss degrade gracefully with this parameter.

for each subset of the goods. With *item bidding*, the action space of each player is much smaller than its type space, and is a subset of $\mathbb{R}_+^m$, with one bid per good. An action $b_{i1}, \ldots, b_{im}$ is feasible for the type v_i if the player does not overbid on any bundle of goods: $\sum_{j \in S} b_{ij} \leq v_i(S)$ for every bundle S of goods.

The standard allocation rule $\mathbf{x}$ with item bidding is to assign independently each good j to the highest bidder $\text{argmax}_i\, b_{ij}$ for it, breaking ties according to some fixed rule. In this paper, we consider the second-price payment rule $\mathbf{p}$ that charges the winner of bundle S the price $\sum_{j \in S} b_{(2)j}$, where $b_{(2)j}$ denotes the second-highest bid on good j.

Simultaneous second-price auctions were first studied in [12], where bidder valuation functions were required to be submodular.[4] They proved that the iPOA is precisely $\frac{1}{2}$. They did not consider the cPOA, which was later shown to be inverse polynomial in m [5]. More general classes of valuations [5, 15] and other payment rules [6, 22, 40] have also been considered.

Example 2.3 (Greedy Combinatorial Auctions) We again consider combinatorial auctions, but with a full bid space. That is, a player i with valuation v_i submits one bid $b_i(S)$ for each subset S of the goods, subject only to the constraint that $b_i(S) \leq v_i(S)$ for every subset S of goods.

A *greedy* allocation rule works as follows. At each step it irrevocably allocates a bundle to a single player, with each player considered exactly once. At each step, player-bundle pairs are ranked according to a function that depends only on the player, the bundle, the player's value for that bundle, and the assignments made thus far. The highest-ranked pair (subject to feasibility) determines the next player and the bundle assigned to it. The corresponding *critical bid* payment rule charges each player the minimum bid at which it would continue to receive the same bundle from the allocation rule.

Greedy combinatorial auctions were first considered in [25], where it was proved that if the greedy allocation rule is a $\frac{1}{c}$-approximation algorithm for the underlying welfare maximization problem, then the iPOA is at least $\frac{1}{c+1}$, and that this bound is tight in general. The cPOA has not been studied in this model.

2.3 Motivating Examples from Routing Games

Our theory is also relevant for games of incomplete information that arise naturally outside of mechanism design. We give two examples related to *(atomic) selfish routing games* [34].[5] Traditionally, a (weighted) selfish routing game is a game of full information. There are n players, and each picks a path in a network $G = (V, E)$. Specifically, each player i has a weight w_i, an origin $o_i \in V$, a destination $d_i \in V$, and its actions $\mathcal{A}_i$ are the o_i-d_i paths of G. In routing games, it is convenient to use costs, which everyone wants to minimize, instead of payoffs. Each edge $e \in E$ has a cost function $\ell_e : \mathbb{R}^+ \to \mathbb{R}^+$ that specifies the per-unit-weight cost incurred by players on the edge e, as a function of the total weight of the players that choose paths that include e. The overall cost incurred by player i is then additive over the edges in its path a_i: $w_i \cdot \sum_{e \in a_i} \ell_e(f_e)$,

[4] A set function $f : 2^U \to \mathbb{R}$ is *submodular* if for every $S \subseteq T$ and $j \in U \setminus T$, $f(T \cup \{j\}) - f(T) \leq f(S \cup \{j\}) - f(S)$. This is a set-theoretic notion of "diminishing returns". The results in [12] are proved more generally for "XOS valuations;" we discuss only submodular valuations to keep the exposition simple.

[5] We could equally well consider the more general but abstract class of *congestion games* [33].

where $f_e = \sum_{j \,:\, e \in a_j} w_j$. The price of anarchy in such (full-information) games is thoroughly understood [1, 3, 4, 10, 11, 35]. For example, when every edge cost function is affine, the worst-case POA is 2.5 with unit-weight players and $(1+\sqrt{5})/2 \approx 2.618$ with arbitrary-weight players [3, 10, 11].

We propose two simple and new models to address potential uncertainty about players' origin-destination pairs and weights, respectively. There are also other ways of incorporating uncertainty into selfish routing models [16, 17, 18, 30]. Our goal here is to provide a simple demonstration of the relevance of our general framework to application domains beyond mechanism design.

Example 2.4 (Routing Games with Uncertain Origin-Destination Pairs) The game structure is as follows. There are n players, each with unit weight. The network $G = (V, E)$ and edge cost functions are publicly known. The private type of a player i is its origin-destination pair. The feasible actions of a player are the unsplittable unit flows from its origin to its destination.

Example 2.5 (Routing Games with Unknown Weights) Here, there are n players, each with a publicly known origin-destination pair. The network $G = (V, E)$ and edge cost functions are also known. The private type of a player i is its weight w_i. The feasible actions of a player are the unsplittable flows of w_i units from its origin to its destination.

Neither the iPOA nor the cPOA have been considered previously in either model.

2.4 Smooth Full-Information Games

For the purposes of completeness and comparison, we review the definition and interpretation of smooth full-information games from [35].

Definition 2.6 ([35]) A game $(\mathcal{A}, \mathbf{u})$ is (λ, μ)-*smooth* with respect to an outcome $\mathbf{a}^*$ and a maximization objective $W : \mathcal{A} \to \mathbb{R}_+$ if

$$\sum_{i=1}^{n} u_i(a_i^*, \mathbf{a}_{-i}) \geq \lambda \cdot W(\mathbf{a}^*) - \mu \cdot W(\mathbf{a})$$

for every outcome $\mathbf{a}$.

There is an analogous definition for minimization objectives [35].[6]

Smooth games correspond to proofs in a prescribed format, outlined in Section 1.1, that bound the price of anarchy of pure-strategy Nash equilibria. To see this, suppose that the objective function W is *payoff-dominating*, meaning that it is always at least as large the sum of players' payoffs. In auction contexts, W is generally the welfare of the outcome. In the common case where the net payments from the buyers to the seller are non-negative, the welfare objective is payoff-dominating. Then, if a game is (λ, μ)-smooth with respect to an outcome $\mathbf{a}^*$, every pure-strategy Nash equilibrium $\mathbf{a}$ has objective function value at least $\lambda/(1+\mu)$ times that of $\mathbf{a}^*$. Precisely, apply payoff dominance, the Nash equilibrium assumption (once per player i, with the hypothetical deviation a_i^*), and smoothness to derive

$$W(\mathbf{a}) \geq \sum_{i=1}^{n} u_i(\mathbf{a}) \geq \sum_{i=1}^{n} u_i(a_i^*, \mathbf{a}_{-i}) \geq \lambda \cdot W(\mathbf{a}^*) - \mu \cdot W(\mathbf{a}), \tag{3}$$

[6] Very roughly, the condition in Definition 2.6 asserts that player payoffs following a "one-dimensional perturbation" (i.e., in $(a_i^*, \mathbf{a}_{-i})$) can be related to the initial outcome $\mathbf{a}$ and the "perturbation directions" $\mathbf{a}^*$. This interpretation is the motivation for the term "smooth."

and then rearrange terms. Many of the known price of anarchy bounds for classes of full-information games are smoothness proofs in this sense.

The main extension theorem in [35] states that, for a payoff-dominating objective and a (λ, μ)-smooth game, the approximation guarantee of $\lambda/(1+\mu)$ extends automatically to the expected objective function value of mixed-strategy Nash equilibria, correlated equilibria, and coarse correlated equilibria.[7]

3 Smooth Games of Incomplete Information

This section defines smooth games of incomplete information and proves our main extension theorem. This section is necessarily abstract; Section 4 instantiates these concepts for each of Examples 2.1–2.5.

3.1 The Definitions

There are two analogous definitions of smooth games of incomplete information, one for maximization objectives (like welfare in an auction) and one for minimization objectives (like the total delay in a routing game). We emphasize that, in a game of incomplete information, the objective function value depends on the actions taken *and* on players' types. For example, the welfare of an allocation in an auction depends on what the players' valuations are. Thus, while Definition 2.6 is parameterized by a single action $\mathbf{a}^*$ (canonically, an optimal action profile), the definitions below are parameterized by a *choice function* $\mathbf{c}^*$ that chooses a feasible action profile for each type profile (canonically, an action profile that is optimal for the given type profile).

We begin with the maximization version of smooth games.

Definition 3.1 (Smooth Games — Maximization Version) Let $\Gamma = (\mathcal{T}, \mathcal{A}, \mathbf{u})$ denote a game structure and $W : \mathcal{T} \times \mathcal{A} \to \mathbb{R}_+$ a maximization objective function. The structure Γ is (λ, μ)-*smooth* with respect to the choice function $\mathbf{c}^* : \mathcal{T} \to \mathcal{A}$ if

$$\sum_{i=1}^{n} u_i(t_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) \geq \lambda \cdot W(\mathbf{t}; \mathbf{c}^*(\mathbf{t})) - \mu \cdot W(\mathbf{s}; \mathbf{a}) \tag{4}$$

for every type vector $\mathbf{t}$, every type vector $\mathbf{s}$, and every outcome $\mathbf{a}$ feasible for $\mathbf{s}$.

Remark 3.2 (Discussion of Definition 3.1) As one would hope, Definition 3.1 specializes to Definition 2.6 in the special case where each player has only one possible type (i.e., in a full-information game). An alternative, more permissive definition would be to call a game of incomplete information (λ, μ)-smooth with respect to $\mathbf{c}^*$ whenever every full-information game induced by a type vector $\mathbf{t}$ is (λ, μ)-smooth (according to Definition 2.6) with respect to $\mathbf{c}^*(\mathbf{t})$. This alternative definition corresponds to requiring (4) only when $\mathbf{s} = \mathbf{t}$, rather than for all $\mathbf{s}$. The more stringent requirements of Definition 3.1 appear necessary for the most general extension theorem (Theorem 3.5). We do not know any interesting examples that meet the weaker condition but fail to satisfy Definition 3.1.

Does Definition 3.1 hold, for reasonable values of λ and μ, in any interesting classes of games? Fortunately, as with Definition 2.6, the canonical method by which one bounds the price of anarchy

[7] Formal definitions of these equilibrium concepts are not needed in this paper.

of pure-strategy Nash equilibria in every induced full-information game — by following an analogue of the three-step approach outlined in Section 1.1 — typically verifies Definition 3.1. Section 4 gives several concrete examples; Example 3.4 provides another.

Modifying Definition 3.1 for minimization objective functions is straightforward.

Definition 3.3 (Smooth Games — Minimization Version) Let $\Gamma = (\mathcal{T}, \mathcal{A}, \ell)$ denote a game structure and $L : \mathcal{T} \times \mathcal{A} \to \mathbb{R}_+$ a minimization objective function. The structure Γ is (λ, μ)-*smooth* with respect to the choice function $\mathbf{c}^* : \mathcal{T} \to \mathcal{A}$ if

$$\sum_{i=1}^{n} \ell_i(t_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) \leq \lambda \cdot L(\mathbf{t}; \mathbf{c}^*(\mathbf{t})) + \mu \cdot L(\mathbf{s}; \mathbf{a})$$

for every type vector $\mathbf{t}$, every type vector $\mathbf{s}$, and every outcome $\mathbf{a}$ feasible for $\mathbf{s}$.

While examples are mostly relegated to Section 4, we pause here for a relatively simple one, to increase the intuition for and plausible utility of the definitions above. The following argument is from Lucier and Paes Leme [26], rephrased in our terminology.

Example 3.4 (The GSP Auction Is a Smooth Game [26]) Recall the sponsored search auction model of Example 2.1, with k slots with known click-through rates $\alpha_1 \geq \cdots \geq \alpha_k$. It is convenient to think of there being n slots, with $\alpha_i = 0$ for $i \in \{k+1, \ldots, n\}$. The natural objective function is welfare maximization. Define the choice function $\mathbf{c}^*$ by $c_i^*(\mathbf{v}) = \frac{v_i}{2}$ for every i; observe that bidders are ranked by valuation with the bid vector $\mathbf{c}^*(\mathbf{v})$.

We verify Definition 3.1 with respect to $\mathbf{c}^*$ with the parameters $\lambda = \frac{1}{2}$ and $\mu = 1$. Fix a type vector $\mathbf{t}$, meaning a per-click valuation v_i for each player i, and an outcome $\mathbf{a}$, meaning a bid $b_i \in [0, v_i]$ for each player i. Since the auction is anonymous, we can rename the players so that $v_1 \geq \cdots \geq v_n$. Let $\kappa(i)$ denote the name of the player with the ith highest bid in $\mathbf{b}$. We claim that

$$u_i(v_i; (c_i^*(\mathbf{v}), \mathbf{a}_{-i})) \geq \tfrac{1}{2}\alpha_i v_i - \alpha_i b_{\kappa(i)} \tag{5}$$

for every player i. To see why, fix i and suppose that player i receives a slot $j \leq i$ in $(c_i^*(\mathbf{v}), \mathbf{a}_{-i})$. Since click-through rates are nonincreasing and the player's price per click is at most its bid $c_i^*(\mathbf{v}) = \frac{1}{2}v_i$, its utility is at least $\alpha_j(v_i - c_i^*(\mathbf{v})) \geq \frac{1}{2}\alpha_i v_i$. If player i is not assigned such a slot in $(c_i^*(\mathbf{v}), \mathbf{a}_{-i})$, then $b_{\kappa(i)} \geq c_i^*(\mathbf{v}) = \frac{1}{2}v_i$ and the right-hand side of (5) is non-positive, so inequality (5) holds.

Summing (5) over all players gives

$$\sum_{i=1}^{n} u_i(v_i; (c_i^*(\mathbf{v}), \mathbf{a}_{-i})) \geq \tfrac{1}{2}\sum_{i=1}^{n} \alpha_i v_i - \sum_{i=1}^{n} \alpha_i b_{\kappa(i)}.$$

The left-hand side is the same as that in Definition 3.1. Because bids in $\mathbf{c}^*(\mathbf{v})$ are ordered by valuation, the first summation $\sum_{i=1}^{n} \alpha_i v_i$ on the right-hand side equals $W(\mathbf{v}; \mathbf{c}^*(\mathbf{v}))$. The second summation on the right-hand side is at most $W(\mathbf{s}; \mathbf{b})$ for every type profile $\mathbf{s}$ for which $\mathbf{b}$ is feasible — that is, every valuation profile $\mathbf{v}'$ with $\mathbf{v}' \geq \mathbf{b}$ component-wise has welfare $W(\mathbf{v}'; \mathbf{b}) = \sum_{i=1}^{n} \alpha_i v'_{\kappa(i)} \geq \sum_{i=1}^{n} \alpha_i b_{\kappa(i)}$ under the bid profile $\mathbf{b}$. (Recall from Example 2.1 that overbidding is infeasible in this model.) Since the type and action profiles were arbitrary, Definition 3.1 holds with respect to $\mathbf{c}^*$ with the constants $\lambda = \frac{1}{2}$ and $\mu = 1$.

A key point in Example 3.4 is that the entire argument works only with a fixed type vector and pure action profiles; no randomization over types or over strategies is considered. The argument is in essence meant for the pure-strategy Nash equilibria of a full-information game, but it happens to meet additional criteria that enable the application of an extension theorem.

3.2 The Extension Theorems

This section states and proves our main extension theorem. By an *optimal* choice function (for a fixed objective function), we mean one that always chooses an action profile $\mathbf{c}^*(\mathbf{t})$ that is optimal for the types $\mathbf{t}$. Recall that an objective function W is payoff-dominating if it is at least the sum of the players' payoffs (like welfare in an auction). We now show that if a game structure is (λ, μ)-smooth with respect to an optimal choice function, then the price of anarchy of (mixed-strategy) Bayes-Nash equilibria is at least $\lambda/(1+\mu)$ in every game of incomplete information induced by a product distribution over players' types.

Theorem 3.5 (Extension Theorem - Maximization Version) *If a game structure $\Gamma = (\mathcal{T}, \mathcal{A}, \mathbf{u})$ is (λ, μ)-smooth with respect to an optimal choice function for a payoff-dominating maximization objective W, then the iPOA of Γ with respect to W is at least $\lambda/(1+\mu)$.*

Proof: Let Γ be (λ, μ)-smooth with respect to the optimal choice function $\mathbf{c}^*$. Let $\mathbf{F}$ be a product distribution on $\mathcal{T}$. Let σ be a Bayes-Nash equilibrium in the induced game of incomplete information. For every i and t_i, $\sigma_i(t_i)$ is feasible for t_i with probability 1.

Let $\hat{\sigma}_i(t_i)$ denote the following mixed-strategy deviation for player i when its type is t_i: sample $\mathbf{s}_{-i}^{(i)} \sim \mathbf{F}_{-i}$ and play the action $c_i^*(t_i, \mathbf{s}_{-i}^{(i)})$.[8] Importantly, because $\mathbf{F}$ is a product distribution, when sampling $\mathbf{s}_{-i}^{(i)}$ it makes no difference whether or not we condition on i's type t_i — the conditional distribution is simply the product of the (unconditional) marginals of F for the players other than i.

For the first phase of our derivation, we use the fact that W is payoff-dominating, linearity of expectation, the fact that σ is a Bayes-Nash equilibrium, and the definition of the $\hat{\sigma}_i$'s to derive the following lower bound on the expected objective function value of the equilibrium:

$$\begin{aligned}
\mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[W(\mathbf{t};\mathbf{a})]\right] &\geq \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}\left[\sum_{i=1}^{n} u_i(t_i;\mathbf{a})\right]\right] \\
&= \sum_{i=1}^{n} \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[u_i(t_i;\mathbf{a})]\right] \\
&\geq \sum_{i=1}^{n} \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\hat{a}_i\sim\hat{\sigma}_i(t_i),\mathbf{a}\sim\sigma(\mathbf{t})}[u_i(t_i;(\hat{a}_i,\mathbf{a}_{-i}))]\right] \\
&= \sum_{i=1}^{n} \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{s}_{-i}^{(i)}\sim\mathbf{F}_{-i},\mathbf{a}\sim\sigma(\mathbf{t}_{-i})}\left[u_i(t_i;(c_i^*(t_i,\mathbf{s}_{-i}^{(i)}),\mathbf{a}_{-i}))\right]\right]. \qquad (6)
\end{aligned}$$

The second phase of the derivation leans on the stochastic independence of players' types. Since the distributions of the $\mathbf{s}_{-i}^{(i)}$'s are projections of a common product distribution $\mathbf{F}$, we can use

[8] The intuition behind the deviation $\hat{\sigma}$ is as follows. Ideally, we would like to consider a deviation by i from its equilibrium strategy to its strategy in an optimal solution. Unfortunately, the optimal solution is a function of the other players' types $\mathbf{t}_{-i}$, which are unknown to i. The closest player i can come to this ideal deviation on its own is to simulate the random types of the other players and then play the corresponding hypothetically optimal action.

linearity of expectation to write

$$\sum_{i=1}^{n} \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{s}_{-i}^{(i)}\sim\mathbf{F}_{-i},\mathbf{a}\sim\sigma_{-i}(\mathbf{t})}\left[u_i(t_i;(c_i^*(t_i,\mathbf{s}_{-i}^{(i)}),\mathbf{a}_{-i}))\right]\right] = \sum_{i=1}^{n} \mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[u_i(t_i;(c_i^*(t_i,\mathbf{s}_{-i}),\mathbf{a}_{-i}))]\right]; \tag{7}$$

that is, we can sample $\mathbf{s}$ once "up front" and use its projections $\mathbf{s}_{-i}$, rather than having each player i sample its own independent copy $\mathbf{s}_{-i}^{(i)}$.

Next, for each player i and type t_i, the random variables $\mathbf{t}_{-i}$ and $\mathbf{s}_{-i}$ are independent and identically distributed, so

$$\mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[u_i(t_i;(c_i^*(t_i,\mathbf{s}_{-i}),\mathbf{a}_{-i}))]\right] = \mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{s})}[u_i(t_i;(c_i^*(t_i,\mathbf{t}_{-i}),\mathbf{a}_{-i}))]\right], \tag{8}$$

where we are also using that $u_i(t_i;(c_i^*(t_i,\mathbf{s}_{-i}),\mathbf{a}_{-i}))$ is independent of a_i.

The third and final phase of the derivation uses the smoothness assumption. After combining (6)–(8) with linearity of expectation, we use the fact that the game is (λ,μ)-smooth with respect to $\mathbf{c}^*$ to obtain

$$\begin{aligned}\mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[W(\mathbf{t};\mathbf{a})]\right] &\geq \mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{s})}\left[\sum_{i=1}^{n} u_i(t_i;(c_i^*(t_i,\mathbf{t}_{-i}),\mathbf{a}_{-i}))\right]\right] \\ &\geq \mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{s})}[\lambda\cdot W(\mathbf{t};\mathbf{c}^*(\mathbf{t}))-\mu\cdot W(\mathbf{s};\mathbf{a})]\right],\end{aligned}$$

where in applying Definition 3.1 we use the fact that $\mathbf{a}\sim\sigma(\mathbf{s})$ is feasible for $\mathbf{s}$ with probability 1. To wrap things up, we note that the term

$$\mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{s})}[W(\mathbf{t};\mathbf{c}^*(\mathbf{t}))]\right] = \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}[W(\mathbf{t};\mathbf{c}^*(\mathbf{t}))]$$

equals the expected optimal objective function value (since $\mathbf{c}^*$ is an optimal choice function), and the term

$$\mathbf{E}_{\mathbf{t},\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{s})}[W(\mathbf{s};\mathbf{a})]\right] = \mathbf{E}_{\mathbf{s}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{s})}[W(\mathbf{s};\mathbf{a})]\right]$$

equals the expected objective function value of the Bayes-Nash equilibrium σ (since $\mathbf{t},\mathbf{s}$ are identically distributed). Rearranging terms shows that the expected objective function value of σ is at least a $\lambda/(1+\mu)$ fraction of that of the maximum possible. Since $\mathbf{F}$ was an arbitrary product distribution and σ was an arbitrary Bayes-Nash equilibrium with respect to it, the proof is complete. ■

Remark 3.6 (Discussion of Theorem 3.5) The proof of Theorem 3.5 is somewhat messy, but this is to be expected given the three probability distributions — over types, equilibrium strategies, and deviation strategies — that have to be carefully managed. The proof also has some subtle steps that use the independence of players' types, but this is also to be expected, since the extension theorem is generally false for the cPOA (see [5] and Example 4.6). Finally, the proof generalizes a sequence of analogous arguments for specific games, beginning with the price of anarchy bound for Bayes-Nash equilibria in simultaneous second-price auctions (Example 2.2) given in [12].

An analogous proof establishes an extension theorem for smooth games with respect to cost-dominated minimization objectives.

Theorem 3.7 (Extension Theorem - Minimization Version) *If a game structure $\Gamma = (\mathcal{T},\mathcal{A},\ell)$ is (λ,μ)-smooth with respect to an optimal choice function for a cost-dominated minimization objective L, then the iPOA of Γ with respect to L is at most $\lambda/(1-\mu)$.*

4 Applications

This section shows that each of the game structures in Examples 2.1–2.5 is smooth with reasonable constants λ and μ. The first three examples recover known results in diverse mechanism design settings in a unified and modular way. The second two examples provide new price of anarchy bounds for routing games with incomplete information.

4.1 The Generalized Second Price Auction

Recall the game structure defined by the Generalized Second Price auction for sponsored search (Example 2.1). Types correspond to valuations-per-click and actions to bids-per-click. In Example 3.4, we showed that this game structure is $(\frac{1}{2}, 1)$-smooth with respect to the (payoff-dominating) welfare objective function and the choice function $\mathbf{c}^*$ defined by $c_i^*(\mathbf{v}) = \frac{v_i}{2}$ for every i. Since bidders are ranked by valuation under the bid profile $\mathbf{c}^*(\mathbf{v})$ and slot click-through rates are nonincreasing, an easy exchange argument shows that $\mathbf{c}^*$ is an optimal choice function. Applying Theorem 3.5 immediately implies the following.

Theorem 4.1 ([26]) *For every Generalized Second Price auction setting and product distribution over players' valuations, the expected welfare of every (mixed-strategy) Bayes-Nash equilibrium is at least $\frac{1}{4}$ times the expected maximum welfare.*

The lower bound of Theorem 4.1 was recently improved in [9] to 0.342 via more sophisticated (smoothness) arguments.

4.2 An Extension Theorem for Separable Choice Functions and Correlated Types

As noted by Lucier and Paes Leme [26], a much stronger version of Theorem 4.1 also holds. The choice function $\mathbf{c}^*$ used to prove Theorem 4.1 has the remarkable property that it is *separable*, meaning that it can be written as $(c_1^*, \dots, c_n^*)$, where c_i^* is a function from $\mathcal{T}_i$ to $\mathcal{A}_i$. Precisely, in the proof of Theorem 4.1, $c_i^*(\mathbf{v}) = \frac{1}{2}v_i$, independent of $\mathbf{v}_{-i}$.

The following extension theorem shows that, whenever the hypotheses of Theorem 3.5 are satisfied with a separable choice function, the conclusion holds even with *correlated* player types (i.e., for the cPOA). This extension theorem is implicit in Lucier and Paes Leme [26] and explicit in Caragiannis et al. [9]; we include the proof here for completeness.

Theorem 4.2 (Extension Theorem - Maximization Version) *If a game structure $\Gamma = (\mathcal{T}, \mathcal{A}, \mathbf{u})$ is (λ, μ)-smooth with respect to an optimal separable choice function for a payoff-dominating maximization objective W, then the cPOA of Γ with respect to W is at least $\lambda/(1 + \mu)$.*

Proof: Let Γ be (λ, μ)-smooth with respect to the optimal separable choice function $\mathbf{c}^* = (c_1^*, \dots, c_n^*)$, where c_i^* maps $\mathcal{T}_i$ to $\mathcal{A}_i$. Taking $\mathbf{s} = \mathbf{t}$ in Definition 3.1 implies that every full-information game induced by a type vector $\mathbf{t}$ is (λ, μ)-smooth with respect to $\mathbf{c}^*(\mathbf{t})$ in the sense of Definition 3.1; see also Remark 3.2. We only need this weaker version of smoothness in the following proof.

Let $\mathbf{F}$ be an arbitrary distribution on $\mathcal{T}$ and σ a Bayes-Nash equilibrium in the induced game of incomplete information. We have

$$\begin{aligned}
\mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[W(\mathbf{t};\mathbf{a})]\right] &\geq \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}\left[\sum_{i=1}^{n} u_i(t_i;\mathbf{a})\right]\right] \\
&= \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\sum_{i=1}^{n} \mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[u_i(t_i;\mathbf{a})]\right] \\
&\geq \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\sum_{i=1}^{n} \mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[u_i(t_i;(c_i^*(t_i),\mathbf{a}_{-i}))]\right] \\
&= \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}\left[\sum_{i=1}^{n} u_i(t_i;(c_i^*(t_i),\mathbf{a}_{-i}))\right]\right] \\
&\geq \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[\lambda\cdot W(\mathbf{t};\mathbf{c}^*(\mathbf{t}))-\mu\cdot W(\mathbf{t};\mathbf{a})]\right] \\
&= \lambda\cdot\mathbf{E}_{\mathbf{t}\sim\mathbf{F}}[W(\mathbf{t};\mathbf{c}^*(\mathbf{t}))]-\mu\cdot\mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\left[\mathbf{E}_{\mathbf{a}\sim\sigma(\mathbf{t})}[W(\mathbf{t};\mathbf{a})]\right],
\end{aligned}$$

where the first inequality follows from payoff dominance, the second from the fact that σ is a Bayes-Nash equilibrium (using the well-defined hypothetical deviation $c_i^*(t_i)$ for player i), and the third from the fact that every induced full-information game is (λ,μ)-smooth respect to $\mathbf{c}^*(\mathbf{t})$. All equalities follow from the linearity of expectation. Since $\mathbf{c}^*$ is an optimal choice function, rearranging terms shows that the expected objective function value of σ is at least a $\lambda/(1+\mu)$ fraction of that of the maximum possible, completing the proof. ∎

The optimal choice functions used in the remaining four examples are not separable, so Theorem 4.2 does not apply to them.

4.3 Simultaneous Second-Price Auctions

Recall the setting of Example 2.2, where there are m goods, types correspond to submodular valuation functions, and actions correspond to bid vectors (with one bid per good). Feasible bid vectors are those that don't overbid on any bundle. Each good is allocated independently, to the highest bidder for it, at a price equal to the second-highest bid for the good.

We isolate a few key inequalities in [12] and show how they imply that this game structure is (1,1)-smooth for an optimal choice function and the welfare objective function. The choice function $\mathbf{c}^*$ is defined as follows. For a given type vector $\mathbf{t}$ — a submodular valuation function v_i for each player i — let $(S_1^*,\ldots,S_n^*)$ denote an allocation maximizing the welfare $\sum_{i=1}^n v_i(S_i)$ over all feasible allocations $(S_1,\ldots,S_n)$. Now consider a player i. Assume by relabeling that S_i^* contains the goods $1,2,\ldots,d$ for some d. Set $b_{ij}^* = v_i(\{1,2,\ldots,j\}) - v_i(\{1,2,\ldots,j-1\})$ for $j=1,2,\ldots,d$ and $b_{ij}^*=0$ for $j>d$. Since v_i is submodular, this bid is feasible for v_i: $v_i(T)\geq\sum_{j\in T} b_{ij}^*$ for every bundle T. We define $c_i^*(\mathbf{t})$ to be this bid vector b_i^*. It is easy to see that $\mathbf{c}^*$ is an optimal choice function: for every type vector $\mathbf{t}$, every player i bids a positive amount on the goods it receives in the optimal allocation for $\mathbf{t}$, and all other players bid zero on these goods.

Now we prove smoothness. Fix a type vector $\mathbf{t}$ (i.e., submodular valuations $\mathbf{v}$). Fix an action vector $\mathbf{a}$ — a bid b_{ij} by each player i for each good j. Suppose player i bids on the goods S_i^* in $c_i^*(\mathbf{t})$

and wins the goods $T \subseteq S_i^*$. Using the definition of $c_i^*(\mathbf{t})$ and a second-price auction, we have

$$\begin{aligned} u_i(v_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) &= v_i(T) - \sum_{j \in T} \max_{k \neq i} b_{kj} \\ &\geq \sum_{j \in T} b_{ij}^* - \sum_{j \in T} \max_{k \neq i} b_{kj} \\ &\geq \sum_{j \in S_i^*} b_{ij}^* - \sum_{j \in S_i^*} \max_{k \neq i} b_{kj} \\ &= v_i(S_i^*) - \sum_{j \in S_i^*} \max_{k \neq i} b_{kj} \\ &\geq v_i(S_i^*) - \sum_{j \in S_i^*} \max_{k} b_{kj}. \end{aligned}$$

Summing over the players and using the fact that the S_i^*'s are a partition of the goods, we have

$$\sum_{i=1}^{n} u_i(v_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) \geq \sum_{i=1}^{n} v_i(S_i^*) - \sum_{j=1}^{m} \max_k b_{kj}.$$

The left-hand side agrees with that in Definition 3.1. The first term on the right-hand side is, by definition of the S_i^*'s, the optimal welfare $W(\mathbf{t}; \mathbf{c}^*(\mathbf{t}))$ for $\mathbf{t}$. For the final term, let $(S_1, \ldots, S_m)$ denote the allocation induced by the bid vector $\mathbf{b}$, and use the fact that highest bidders win to rewrite the term as $\sum_{i=1}^{n} \sum_{j \in S_i} b_{ij}$. By the definition of feasible bids, this quantity is at most the welfare $\sum_{i=1}^{n} v_i'(S_i)$ for every type vector $\mathbf{s}$ (i.e., valuations $\mathbf{v}'$) for which the bids $\mathbf{b}$ are feasible.

Since the game structure is (1,1)-smooth with respect to the optimal choice function $\mathbf{c}^*$, Theorem 3.5 immediately gives the following.

Theorem 4.3 ([12]) *With simultaneous second-price auctions, for every product distribution over players' submodular valuations, the expected welfare of every (mixed-strategy) Bayes-Nash equilibrium is at least $\frac{1}{2}$ times the expected maximum welfare.*

4.4 Greedy Combinatorial Auctions

The greedy combinatorial auctions of Example 2.3 can be treated similarly to simultaneous second-price auctions. A few key observations and inequalities in [25] show that an auction derived from a greedy $\frac{1}{c}$-approximation algorithm is $(1, c)$-smooth with respect to a natural optimal choice function. Theorem 3.5 immediately gives the following.

Theorem 4.4 ([25]) *For every combinatorial auction with a $\frac{1}{c}$-approximate greedy allocation rule and a critical bid payment rule, and every product distribution over players' valuations, the expected welfare of every (mixed-strategy) Bayes-Nash equilibrium is at least $\frac{1}{c+1}$ times the expected maximum welfare.*

4.5 Routing Games with Uncertain Origin-Destination Pairs

We now turn to a different class of examples: routing games. The results in this section and the next are new.

We first recall the game structure introduced in Example 2.4. There is a fixed network $G = (V, E)$, and each edge $e \in E$ has a cost function ℓ_e. To keep the discussion simple, we first assume that every cost function is affine (and nonnegative and nondecreasing). Every player has a unit amount of traffic that it has to route on a single path. The private type t_i of a player i is its origin-destination pair (o_i, d_i); its feasible strategies are the o_i-d_i paths in G. The standard objective function is to minimize the sum of the players' costs. There is an obvious optimal choice function $\mathbf{c}^*$: given types $\mathbf{t}$, let $\mathbf{a}^*$ be the action profile that minimizes the sum of the players' costs over all feasible routings (given their o_i-d_i pairs), and set $c_i^*(\mathbf{t}) = a_i^*$.[9]

We now prove smoothness. One interesting difference between this proof and the previous three is that we argue edge-by-edge, rather than player-by-player. Fix a type vector $\mathbf{t}$ — that is, an (o_i, d_i) pair for each player i. The choice function $\mathbf{c}^*(\mathbf{t})$ corresponds to paths $P_1^*, \ldots, P_n^*$, where P_i^* is an o_i-d_i path in G. Fix an action vector $\mathbf{a}$ — a set of n paths $P_1, \ldots, P_n$ in G with origins $o'_1, \ldots, o'_n$ and $d'_1, \ldots, d'_n$. We emphasize that there need not be any relationship between the o_i-d_i pairs (which correspond to the types $\mathbf{t}$) and the o'_i-d'_i pairs (which correspond to some other types $\mathbf{s}$). In contrast to the preceding auction settings, there is a unique set of types $\mathbf{s}$ for which the actions $\mathbf{a}$ are feasible.

For an edge e, define f_e^* and f_e as the number of paths from each set that include e: $f_e^* = |\{i \in \{1, 2, \ldots, n\} : e \in P_i^*\}|$ and $f_e = |\{i \in \{1, 2, \ldots, n\} : e \in P_i\}|$. Since cost functions are nondecreasing, we can write

$$\sum_{i=1}^{n} \ell_i(t_i; (P_i^*, P_{-i})) \leq \sum_{e \in E} f_e^* \cdot \ell_e(f_e + 1). \tag{9}$$

Next we use the elementary fact that $y(z+1) \leq \frac{5}{3}y^2 + \frac{1}{3}z^2$ for all nonnegative integers y, z [11, Lemma 1]. Expanding the affine cost functions c_e, applying this inequality to each edge e (with $y = f_e^*$ and $z = f_e$), and rearranging we obtain

$$\sum_{i=1}^{n} \ell_i(t_i; (P_i^*, P_{-i})) \leq \frac{5}{3} \sum_{e \in E} \ell_e(f_e^*) f_e^* + \frac{1}{3} \sum_{e \in E} \ell_e(f_e) f_e. \tag{10}$$

The two sums on the right-hand side of (10) are precisely the objective function values of $\mathbf{c}^*(\mathbf{t})$ and of $\mathbf{a}$ — in the latter case, for the unique types $\mathbf{s}$ for which $\mathbf{a}$ is feasible. Since $\mathbf{t}$ and $\mathbf{a}$ were arbitrary, this proves that the game structure is $(\frac{5}{3}, \frac{1}{3})$-smooth with respect to the optimal choice function $\mathbf{c}^*$. Applying Theorem 3.7 immediately yields the following theorem.

Theorem 4.5 *For every unweighted selfish routing game with affine cost functions and product distribution over players' origin-destination pairs, the expected cost of every (mixed-strategy) Bayes-Nash equilibrium is at most $\frac{5}{2}$ times the expected minimum cost.*

The bound of $\frac{5}{2}$ is tight, even for pure Nash equilibria in the full-information model [3, 11].

Analogues of Theorem 4.5 hold for all classes of cost functions, with the approximation bound degrading with the "nonlinearity" of the cost functions in a well-understood way. Precisely, previous works [1, 35] identify, for every set $\mathcal{L}$ of cost functions, the coefficients λ and μ (above, $\frac{5}{3}$ and $\frac{1}{3}$) that minimize the price of anarchy upper bound $\lambda/(1-\mu)$ subject to the smoothness constraint (10).

[9] Routing games are in this sense simpler than the auction models studied earlier in the paper: rather than bidding to coax indirectly a desired allocation from a mechanisms, a player can choose directly the desired path.

These works also show matching lower bounds for every set $\mathcal{L}$, which even apply to pure Nash equilibria.

Example 4.6 (cPOA Lower Bound) The upper bound of 5/2 on the iPOA in Theorem 4.5 does not hold more generally for the cPOA. Consider the following game structure. For an arbitrarily large perfect square n, there are n players. There is a set S of $n + \sqrt{n}$ edges that each have the cost function $c(x) = x$. For each subset $A \subseteq S$ of size $\sqrt{n} + 1$, there is an origin o_A with zero-cost edges to the tails of the edges in A, and a destination d_A with zero-cost edges from the heads of the edges in A. For each player, the possible types correspond to these sets A. As discussed above, this game structure is $(\frac{5}{3}, \frac{1}{3})$-smooth in the sense of Definition 3.3.

We describe the (correlated) distribution $\mathbf{F}$ over types via a sampling algorithm. The first step chooses a random subset $B \subseteq S$ of size $\sqrt{n}$. The second step chooses a random bijection f from the players to $S \setminus B$. The type of player i is defined as the origin-destination pair o_{A_i}-d_{A_i}, where $A_i = B \cup \{f(i)\}$.

With probability 1 over the random choice of types, there is an outcome of the routing game in which players use disjoint paths (with player i using edge $f(i)$) and incur total cost n. By symmetry, the strategy profile in which every player always randomizes uniformly over its $\sqrt{n} + 1$ paths is a Bayes-Nash equilibrium. The expected cost of the path chosen by a player is $\approx \sqrt{n}$, and the expected total cost in this Bayes-Nash equilibrium is $\approx n^{3/2}$.

4.6 Routing Games with Unknown Weights

Routing games with uncertain weights (Example 2.5) can be treated in a similar way. Here, the induced full-information games are weighted selfish routing games. Weighted routing games are harder to analyze than their unweighted counterparts, but recent work has determined the worst-case price of anarchy for all classes of cost functions that satisfy mild closure conditions [1, 4]. The analogues of inequalities (9) and (10) are

$$\sum_{i=1}^{n} \ell_i(t_i; (P_i^*, P_{-i})) \leq \sum_{e \in E} f_e^* \cdot \ell_e(f_e + f_e^*) \leq \lambda \sum_{e \in E} \ell_e(f_e^*) f_e^* + \mu \sum_{e \in E} \ell_e(f_e) f_e, \tag{11}$$

where f_e and f_e^* denote the total weight of players that pick a path that contains e in an arbitrary outcome and an optimal outcome, respectively. Previous works [1, 4] identified, for every set of cost functions that satisfies mild conditions, the coefficients λ and μ that minimize $\lambda/(1-\mu)$ subject to (11). There are matching lower bounds, even for pure Nash equilibria. Since these upper bound proofs show more generally that routing games with type-dependent player weights are (λ, μ)-smooth in the sense of Definition 3.3, Theorem 3.7 extends these upper bounds to all Bayes-Nash equilibria with respect to a product prior distribution over players' weights. For example, we have the following.

Theorem 4.7 *For every selfish routing game with affine cost functions and product distribution over players' weights, the expected cost of every (mixed-strategy) Bayes-Nash equilibrium is at most* $\frac{1+\sqrt{5}}{2} \approx 2.618$ *times the expected minimum cost.*

The constant in Theorem 4.7 is slightly bigger than that in Theorem 4.5 because the price of anarchy in weighted routing games is larger than in unweighted games.

Finally, the bound in Theorem 4.7 continues to apply in the more general model in which players' weights and origin-destination pairs are both uncertain. The weight and origin-destination pair of a given player can be correlated, as long as the types of different players are independent.

5 Conclusions

This paper proposed a notion of smooth games of incomplete information and proved extension theorems for such games. This theory reduces the goal of proving approximation guarantees for mixed-strategy Bayes-Nash equilibria with respect to an arbitrary product prior distribution to the simpler task of proving such guarantees for pure-strategy Nash equilibria of full-information games. Several previous works on the price of anarchy of Bayes-Nash equilibria are, in hindsight, specific instantiations of this general approach. Ideally, researchers can focus their creativity on understanding the pure-strategy Nash equilibria of interesting new full-information models, relying on extension theorems to perform the dirty work needed for more general results.

There are also opportunities to tailor the theory developed here to important problem domains. For example, Syrgkanis [39] proves that in domains where (i) the set $\mathcal{A}_i$ of feasible actions is independent of the player's type t_i; and (ii) the objective function $W(;)$ equals the sum of players' payoffs (instead of only being payoff-dominating), there is an extension theorem for mixed-strategy Bayes-Nash equilibria assuming only that every induced full-information game is smooth (cf., Remark 3.2).

Motivated by mechanism design problems with quasi-linear player utilities, including some not discussed in Section 2.2, Syrgkanis and Tardos [40] propose several modifications of our smoothness definition and extension theorems. The simplest one replaces the term $W(\mathbf{s};\mathbf{a})$ on the right-hand side of (4) with the revenue of the auction when the actions (i.e., bids) are $\mathbf{a}$. This definition and a corresponding extension theorem lead to good iPOA bounds for certain pay-as-bid mechanisms [22, 25] without any no overbidding restrictions, and to composition theorems that extend smoothness guarantees for a single mechanism to a collection of mechanisms running simultaneously or sequentially.

Finally, Feldman et al. [15] recently extended the iPOA bound of $\frac{1}{2}$ for simultaneous second-price auctions (Example 2.2) to the setting of subadditive bidder valuations, by departing from the smoothness paradigm and reasoning directly about mixed-strategy Bayes-Nash equilibria. It would be interesting to understand better the power and limitations of their approach.

References

[1] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing*, 40(5):1211–1233, 2011.

[2] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[3] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.

[4] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, volume 2, pages 17–28, 2010.

[5] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 700–709, 2011.

[6] K. Bhawalkar and T. Roughgarden. Simultaneous single-item auctions. In *Proceedings of the 8th International Workshop on Internet and Network Economies (WINE)*, pages 338–350, 2012.

[7] A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. *Theory of Computing*, 6(1):179–199, 2010.

[8] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 373–382, 2008.

[9] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. *Journal of Economic Theory*, 2014. To appear.

[10] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, volume 3669 of *Lecture Notes in Computer Science*, pages 59–70, 2005.

[11] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.

[12] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 820–832, 2008.

[13] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[14] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

[15] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 201–210, 2013.

[16] M. Gairing. Malicious Bayesian congestion games. In *Proceedings of the 6th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 119–132, 2008.

[17] M. Gairing, B. Monien, and K. Tiemann. Selfish routing with incomplete information. *Theory of Computing Systems*, 42(1):91–130, 2008.

[18] C. Georgiou, T. Pavlides, and A. Philippou. Selfish routing in the presence of network uncertainty. *Parallel Processing Letters*, 19(1):141–157, 2009.

[19] M. X. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

[20] J. Hannan. Approximation to Bayes risk in repeated play. *Annals of Mathematical Studies*, 39:97–139, 1957.

[21] J. C. Harsanyi. Games with incomplete information played by 'Bayesian' players, parts I, II, and III. *Management Science*, 14:159–182, 320–334, and 486–502, 1967.

[22] A. Hassidim, H. Kaplan, M. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 295–296, 2011.

[23] V. A. Knight and P. R. Harper. The price of anarchy in health care. Submitted, 2011.

[24] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.

[25] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 537–553, 2010.

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

[27] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 183–194, 2004.

[28] H. Moulin and J. P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3/4):201–221, 1978.

[29] U. Nadav and T. Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *Proceedings of the 6th International Workshop on Internet and Network Economies (WINE)*, pages 319–326, 2010.

[30] E. Nikolova and N. E. Stier Moses. A mean-risk model for the traffic assignment problem with stochastic travel times. *Operations Research*, 2014. To appear.

[31] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[32] R. Paes Leme and É. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 735–744, 2010.

[33] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[34] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.

[35] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.

[36] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. *Journal of Economic Theory*, 2014. To appear.

[37] B. Skinner. The price of anarchy in basketball. *Journal of Quantitative Analysis in Sports*, 6(1), 2010. Article 3.

[38] A. Skopalik and B. Vöcking. Inapproximability of pure Nash equilibria. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 355–364, 2008.

[39] V. Syrgkanis. Bayesian games and the smoothness framework. arXiv.cs.GT:1203.5155v1, 2012.

[40] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013.

[41] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

[42] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.