On the Complexity of Equilibrium Computation in First-Price Auctions*

Aris Filos-Ratsikas
University of Edinburgh, United Kingdom
Aris.Filos-Ratsikas@ed.ac.uk

Yiannis Giannakopoulos†
FAU Erlangen-Nürnberg, Germany
yiannis.giannakopoulos@fau.de

Alexandros Hollender‡
University of Oxford, United Kingdom
alexandros.hollender@cs.ox.ac.uk

Philip Lazos§
IOHK
philip.lazos@iohk.io

Diogo Poças¶
LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal
dmpocas@fc.ul.pt

Abstract

We consider the problem of computing a (pure) Bayes-Nash equilibrium in the first-price auction with continuous value distributions and discrete bidding space. We prove that when bidders have independent *subjective* prior beliefs about the value distributions of the other bidders, computing an ε-equilibrium of the auction is PPAD-complete, and computing an *exact* equilibrium is FIXP-complete. We also provide an efficient algorithm for solving a special case of the problem, for a fixed number of bidders and available bids.

1 Introduction

Auctions are prime examples of economic environments in which the element of strategic behavior is prevalent. The associated theory can be traced back to as early as the 1960s and the seminal work of Vickrey [1961]. Over the years, auction theory and mechanism design have produced some of the most celebrated results in economics, as can be evidenced, e.g., by the relevant 1996, 2007 and 2020 Nobel Prizes.[1] Among the plethora of auction formats that this rich literature has proposed, some stand out, such as the second-price auction of Vickrey [1961] or the revenue-maximizing auction of Myerson [1981].

Arguably, though, the most fundamental auction format is that of the *first-price auction*, in which the highest bidder wins and is charged an amount equal to her bid. Compared to its counterparts mentioned above, the first-price auction does not enjoy the same desirable incentive properties: participants may

*A preliminary version of this paper appeared in EC'21 [Filos-Ratsikas et al., 2021].

†Part of this work was done while the author was a member of the Operations Research group at Technical University of Munich, School of Management, supported by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF).

‡Supported by an EPSRC doctoral studentship (Reference 1892947)

§Partially supported by the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets" and MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets".

¶Supported by FCT via LASIGE Research Unit, ref. UIDB/00408/2020.

[1] For the official Nobel Prize announcements see here, here and here.

have an incentive to misreport their true bids. At the same time, however, the first-price auction is very natural and simple to describe, implement and participate in, making it very suitable for a range of important applications. As a matter of fact, several online ad exchanges, including Google Ad Manager, have adopted this auction format for selling their ads, which has been coined "the first-price movement" (see, e.g., [Digiday.com, 2019; Paes Leme et al., 2020]).

There has been a large body of work studying incentives and bidding behavior in first-price auctions, dating back to the original paper of Vickrey [1961]. In particular, the literature has studied the equilibria of the auction in an incomplete information setting where the bidders have only probabilistic prior beliefs (or simply *priors*) about the values of other bidders, via the lens of Bayesian game theory [Harsanyi, 1967] (see also [Myerson, 1997; Hartline, 2012]). Several different scenarios of interest have been analyzed; see, e.g., [Griesmer et al., 1967; Riley and Samuelson, 1981; Plum, 1992; Marshall et al., 1994; Lebrun, 1996, 1999; Maskin and Riley, 2000; Lizzeri and Persico, 2000; Athey, 2001; Reny and Zamir, 2004; Chawla and Hartline, 2013; Bergemann et al., 2017]. It is no exaggeration to say that understanding the Bayes-Nash equilibria of the first-price auction has historically been one of the most important questions of auction theory.

The aforementioned literature has been primarily concerned with identifying conditions under which (pure Bayes-Nash) equilibria are guaranteed to exist. Among those, the seminal paper of Athey [2001] has been pivotal in establishing the existence of equilibria for fairly general settings with continuous priors. A natural follow-up question posed explicitly by Athey [2001], which was also very much present in earlier works, is whether these equilibria can also be "found"; in the context of the related literature, this is usually interpreted as coming up with closed-form solutions that describe them.

One of the most significant contributions of computer science to the field of game theory is to formalize and systematically study this notion of "finding" or "computing" equilibria in games. Roughly speaking, an equilibrium can be efficiently computed if it can be found using a limited number of standard operations that can be performed by a computer, where "limited" here typically means a number which is a polynomial function of the size of the input parameters.[2] In perhaps the most important result in computational game theory, Daskalakis et al. [2009] proved that in all likelihood, Nash equilibria of general games cannot always be computed efficiently. In particular, they proved that the problem of computing a Nash equilibrium is complete for the class PPAD [Papadimitriou, 1994], which is widely believed to include problems that are computationally hard to solve.

In this paper, we study the complexity of computing an equilibrium of the first-price auction, in settings with continuous priors and discrete bids. We offer the following main result.

Informal Theorem 1. *Computing a (pure, Bayes-Nash) equilibrium of a first-price auction with continuous subjective priors and discrete bids is* PPAD*-complete.*

This result can be interpreted intuitively as justification of why research in economics has only had limited success in providing closed forms or characterizations for the equilibria of the first-price auction. In addition, we consider it to be a quite valuable addition to the literature of total search problems [Megiddo and Papadimitriou, 1991], as it concerns the computation of equilibria of one of the most fundamental games in auction theory.

1.1 Discussion and Further Results

Below, we provide a more in-depth discussion of our main result and its assumptions, as well as some other related results that we obtain along the way.

Continuous Priors, Discrete Bids. Informal Theorem 1 applies to the case where the bidders' beliefs about the values of other bidders are continuous distributions, whereas the bidding space is a discrete set. The former assumption is standard in auction theory (see, e.g., [Myerson, 1997, Sec. 3.11] or [Krishna,

[2] We remark that contrary to earlier works in economics, Athey's interpretation of "finding" an equilibrium was very much of a computational nature.

2009]). From a technical standpoint, this also guarantees the existence of equilibria [Athey, 2001].[3] The assumption of the discrete bidding space is clearly motivated by any real-world scenario, in which the bids will be increments of some minimum monetary amount, e.g., 1 dollar or 1 cent, depending on the application. This setting has in fact been studied in several works for first-price auctions in particular (see, e.g., [Chwe, 1989; Athey, 2001; Escamocher et al., 2009; Cai et al., 2010; Rasooly and Gavidia-Calderon, 2020]).

Subjective Priors. In Informal Theorem 1 we assume that the priors are subjective, meaning that two different bidders might have different beliefs about the values of some other bidder. In the auction theory literature, it is often assumed that a "universal" prior exists, which is common knowledge among all players; this is known as the *independent private values* model. Indeed, such common priors are quite convenient in settings where there is an aggregate objective that needs to be optimized in expectation (e.g., the social welfare or the seller's revenue), since they can be used by the designer to tune the parameters of the auction in a way that works best for the optimization goal at hand; this is the case, e.g., for Myerson's revenue-maximizing auction [Myerson, 1981].

From our perspective however, where the goal is to study the players' incentives and compute an equilibrium, we believe it is natural to make the more general assumption that priors are still independent, but subjective: this is enough for the bidders to come up with their best responses. As a matter of fact, Harsanyi's original paper [Harsanyi, 1967], as well as classic textbooks in economics (e.g., [Myerson, 1997; Jehle and Reny, 2001]) introduce Bayesian games directly in the context of subjective beliefs.[4] Similar notions of subjective priors and "subjective equilibria" have also been studied rather extensively for general Bayesian games in economics [Hahn, 1973; Fudenberg and Levine, 1986; Battigalli and Guaitoli, 1997; Battigalli et al., 1992; Kalai and Lehrer, 1993, 1995; Rubinstein and Wolinsky, 1994] and computer science [Witkowski and Parkes, 2012; Frongillo and Witkowski, 2016].

The subjective priors assumption is necessary for our PPAD-hardness result, but we would of course be very interested in settling the complexity for the case of common priors as well. In fact, as we explain in Section 7, we consider this to be one of the most important open problems in computational game theory. Thus, besides being of standalone interest, one can also see our result for subjective priors as an important first step in the quest of answering this question. We remark that our PPAD-membership result obviously applies to common priors, as this is just a special case of subjective beliefs.

Approximate Equilibria. While Informal Theorem 1 states the PPAD-completeness of computing an equilibrium of the first-price auction, the formal statement is in fact about ε-equilibria, i.e., stable states in which bidders do not wish to unilaterally deviate unless they are better off by some small positive quantity ε. As we explain in Section 2, this is very much necessary: there are examples where the equilibrium is *irrational*, and therefore cannot be computed exactly in many standard models of computation. As a matter of fact, this is a common theme in most papers in equilibrium computation; see, e.g., [Daskalakis et al., 2009; Chen et al., 2009] or the survey of Goldberg [2011] for a related discussion.

Of course, the focus on ε-equilibria is only relevant for the membership result in PPAD; the computational hardness result for approximate equilibria is clearly stronger. In fact, we show that under some standard assumptions (see Section 2), the problem is PPAD-hard even when ε is allowed to be a (sufficiently small) constant, independent of the input parameters. This is the strongest type of PPAD-hardness one could hope for. For the computation of *exact* equilibria, Etessami and Yannakakis [2010] defined the computational class FIXP. At a high level, this class contains problems that can be stated as computations of (possibly irrational) fixed points of functions defined by means of algebraic circuits (see [Yannakakis, 2009]). We complement our main result about ε-equilibria with the following analogous result on exact ones:

[3] It is important to note here that in some versions of the problem, even *mixed* Bayes-Nash equilibria are not guaranteed to exist; see, e.g., [Lebrun, 1996].

[4] These works also usually provide discussions on "consistency" conditions, e.g., see [Harsanyi, 1967] and [Myerson, 1997, Sec. 2.8]. See also a related discussion in Section 7 of our work.

Informal Theorem 2. *Computing an* exact *(pure, Bayes-Nash) equilibrium of a first-price auction with continuous subjective priors and discrete bids is* FIXP*-complete.*

One way to interpret a FIXP-completeness result in the standard computational (Turing) model is in terms of *strong* vs *weak* approximations. A weak approximation is an ε-equilibrium as defined above and is captured by our PPAD-completeness result. A strong approximation is a set of strategies represented by rational numbers, which are "ε-close" to an exact equilibrium (in terms of the max norm), and is captured by our FIXP-completeness result. We remark that this is completely analogous to the computation of Nash equilibria in general games, see [Etessami and Yannakakis, 2010; Garg et al., 2016a] for a more in-depth discussion.

The Meaning of PPAD-completeness. As we mentioned earlier, a PPAD-hardness result is interpreted as an indication that the problem cannot be solved in polynomial time. In particular, it is as hard as finding Nash equilibria in general games [Daskalakis et al., 2009; Chen et al., 2009; Mehta, 2018; Rubinstein, 2018], market equilibria in Arrow-Debreu markets [Vazirani and Yannakakis, 2011; Chen et al., 2017] or solutions to fixed point theorems [Papadimitriou, 1994; Goldberg and Hollender, 2021]. Additionally, PPAD has been shown to be hard under various cryptographic assumptions (e.g., see [Bitansky et al., 2015; Garg et al., 2016b; Choudhuri et al., 2019; Rosen et al., 2021]), meaning that solving a PPAD-hard problem would "break" those assumptions as well. On the other hand, an "in PPAD" result can be interpreted as the existence of an (inefficient) algorithm that uses a path-following argument to reach a solution.

An Efficient Algorithm. Besides our main PPAD- and FIXP-completeness results, we identify a special case of the problem which can be solved efficiently, namely when the number of bidders and the size of the bidding space are constant, and the value distributions are "sufficiently smooth", in the sense that they are given by piecewise polynomial functions. To this end, we have the following theorem.

Informal Theorem 3. *A (pure, Bayes-Nash) equilibrium of the first-price auction can be computed in polynomial time when there is a constant number of bidders, a constant-size bidding space, and continuous (subjective) priors which are piecewise polynomial functions.*

Informal Theorem 3 complements our PPAD- and FIXP-hardness results rather tightly, as our reductions use a constant bidding space and very simple, piecewise constant distributions, but a large number of bidders.

1.2 Related Work

As we mentioned earlier, there is a significant amount of work in economic theory on the equilibria of the first-price auction [Griesmer et al., 1967; Riley and Samuelson, 1981; Plum, 1992; Marshall et al., 1994; Lebrun, 1996, 1999; Maskin and Riley, 2000; Lizzeri and Persico, 2000; Athey, 2001; Reny and Zamir, 2004; Bergemann et al., 2017; Cheng, 2006]. Among those, the most relevant work to us is that of Athey [2001], who established the existence of pure Bayes-Nash equilibria in games with discontinuous payoffs which satisfy the *single crossing property* of Milgrom and Shannon [1994], of which the first-price auction is a special case. Athey's proof applies to both discrete and continuous bidding spaces, and in fact the latter is established through the former, via a limit argument similar in spirit to [Lebrun, 1996; Maskin and Riley, 2000].

To the best of our knowledge, there are only a few prior works on the computational complexity of equilibria in first-price auctions. Escamocher et al. [2009] study the problem of computing equilibria when *both* the priors and the bidding space are discrete. In that case, it is not hard to construct counter-examples that show that pure equilibria may not exist, and therefore they are concerned with the question of *deciding* their existence. Their results do not provide a conclusive answer (i.e., neither NP-hardness nor polynomial-time solvability is proven), except for the very special case of two bidders with bi-valued distributions. Wang et al. [2020] very recently studied the equilibrium computation problem in settings with *discrete priors* and *continuous bids* (in a sense, the opposite of what we do

here), and under the *Vickrey tie-breaking rule* for deciding the winner of the auction in case of a tie. According to this rule, ties are resolved by running an auxiliary second-price (Vickrey) auction among the potential winners of the first-price auction; effectively this allocates the item to the bidder with highest true valuation. This tie-breaking rule was introduced by Maskin and Riley [2000] primarily as a technical tool in proving their existence results for the *uniform tie-breaking rule*, where ties are broken uniformly at random among the bidders with the highest bid. Our results are proven for the uniform tie-breaking rule, which is the standard rule in the literature of the problem [Lebrun, 1996; Maskin and Riley, 2000; Athey, 2001; Krishna, 2009].

Finally, we remark that while we consider an equilibrium computation setting, our results are markedly different from other works on such problems, e.g., [Daskalakis et al., 2009]. This is because our paper concerns a much more specific and structured game, and crucially, a game which is *Bayesian*, which is not the case for most prior work. Conceptually closer to our work is the paper by Cai and Papadimitriou [2014] who study the complexity of Bayesian *combinatorial* auctions, a more complicated auction format which typically involves multiple items for sale and more complex agent valuations over subsets of items. The complexity of *general* Bayesian games (beyond auctions) has been studied in the literature, primarily resulting in NP-hardness results for several cases of interest, e.g., see [Gottlob et al., 2007; Conitzer and Sandholm, 2008].

2 Model and Notation

In a (Bayesian) *first-price auction (FPA)*, there is a set $N = \{1, 2, \dots, n\}$ of *bidders* (or *players*) and one item for sale. Each player i submits a *bid* $b_i \in B$, where the *bidding space* $B \subseteq [0, 1]$ is a finite set. We will also make the standard assumption (often referred to as the "null bid" in the literature) that $0 \in B$, which can be interpreted as the option of the bidders to not participate in the auction (see, e.g., [Maskin and Riley, 2000; Athey, 2001]).

The item is allocated to the player with the highest bid, who is charged a payment equal to her bid. If there are multiple players submitting the same highest bid, the winner is determined based on the *uniform tie-breaking* rule. Formally, for a *bid profile* $\boldsymbol{b} = (b_1, \dots, b_n)$, the *ex-post utility* of player i with true value v_i is given by

$$\tilde{u}_i(\boldsymbol{b}; v_i) \equiv \begin{cases} \frac{1}{|W(\boldsymbol{b})|}(v_i - b_i), & \text{if } \; i \in W(\boldsymbol{b}), \\ 0, & \text{otherwise}, \end{cases} \qquad \text{where } \; W(\boldsymbol{b}) = \operatorname*{argmax}_{j \in N} b_j \tag{1}$$

For each pair of players $i, j \in N$, $i \neq j$, there is a continuous value distribution $F_{i,j}$ over $[0, 1]$; we call this distribution the *prior* of bidder i over the values of bidder j. The *subjective belief* of player i for the values $\boldsymbol{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ of the other bidders is then given by the product distribution $\boldsymbol{F}_{-i} \equiv \times_{j \neq i} F_{i,j}$. In other words, from the perspective of bidder i, the values v_j for $j \neq i$ are drawn *independently* from distributions $F_{i,j}$. Notice that the special case where $F_{i,j} = F_{i',j}$ for all $j \in N$ and $i, i' \in N \setminus \{j\}$ corresponds to the classic *independent private values* model of auction theory, where the value of each bidder is drawn (independently of the others) from a single distribution. More formally, simplifying the notation by using F_j instead of $F_{i,j}$, $\boldsymbol{v}$ is drawn from the *common prior* distribution $\boldsymbol{F} = \times_{j \in N} F_j$. Obviously, while our hardness results rely on the fact that priors are subjective, all of our positive results trivially extend to the case of common priors as well.

The FPA described above naturally induces a game in which each bidder i selects her bid based on her own (true) value v_i, and her beliefs $\boldsymbol{F}_{-i}$. A *strategy* of bidder i is a function $\beta_i : [0, 1] \to B$ mapping values to bids. Given a strategy profile $\boldsymbol{\beta}_{-i}$ of the other players, the (interim) *utility* of player i with true value v_i when bidding $b \in B$ is

$$u_i(b, \boldsymbol{\beta}_{-i}; v_i) \equiv \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}} \left[\tilde{u}_i(b, \boldsymbol{\beta}_{-i}(\boldsymbol{v}_{-i}); v_i) \right],$$

where $\boldsymbol{\beta}_{-i}(\boldsymbol{v}_{-i})$ is a shorthand for $(\beta_1(v_1), \dots, \beta_{i-1}(v_{i-1}), \beta_{i+1}(v_{i+1}), \cdots, \beta_n(v_n))$. Intuitively, the player calculates her (expected) utility by drawing a value v_j for each bidder $j \neq i$ from her corresponding

subjective prior distribution $F_{i,j}$, and then using the strategy "rules" $\boldsymbol{\beta}_{-i}$ of the others to map their values to actual bids in B.

We are interested in "stable" states of the FPA, i.e., strategy profiles from which no bidder would like to unilaterally deviate to a different strategy. Formally, we have the following definition.

Definition 1 (ε-Bayes-Nash equilibrium of the FPA)**.** Let $\varepsilon \geq 0$. A strategy profile $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_n)$ is a (pure, interim) ε-Bayes-Nash equilibrium (ε-BNE) of the FPA if for any bidder $i \in N$ and any value $v_i \in [0, 1]$,

$$u_i(\beta_i(v_i), \boldsymbol{\beta}_{-i}; v_i) \geq u_i(b, \boldsymbol{\beta}_{-i}; v_i) - \varepsilon \qquad \text{for all } \; b \in B.$$

Given a fixed strategy profile $\boldsymbol{\beta}_{-i}$ of the other bidders, we will denote the set of *ε-best responses* of player i by

$$BR_i^{\varepsilon}(\boldsymbol{\beta}_{-i}) = \left\{ \beta_i \;\middle|\; u_i(\beta_i(v_i), \boldsymbol{\beta}_{-i}; v_i) \geq \max_{b \in B} u_i(b, \boldsymbol{\beta}_{-i}; v_i) - \varepsilon \quad \text{for all } \; v_i \in [0, 1] \right\}$$

Using this, the condition in Definition 1 can be equivalently written as $\beta_i \in BR_i^{\varepsilon}(\boldsymbol{\beta}_{-i})$ for all players i. For the special case of $\varepsilon = 0$, i.e. *exact* best-responses, we will drop the ε superscript.

Notice that, in Definition 1 we define a relaxed equilibrium concept, in which the bidder does not want to change to a different strategy unless it increases her utility by an additive factor larger than ε; obviously, when $\varepsilon = 0$ we recover the standard definition of the (exact) Bayes-Nash equilibrium.

No Overbidding. As part of our model, we will make the assumption that bidders will never submit a bid b_i which is higher than their valuation v_i. This is a standard assumption in the literature of the first-price auction [Maskin and Riley, 2000, 2003; Lebrun, 2006; Escamocher et al., 2009; Wang et al., 2020] and auctions in general [Caragiannis et al., 2015; Lucier and Borodin, 2010; Bhawalkar and Roughgarden, 2011; Feldman et al., 2020; Christodoulou et al., 2016; Paes Leme and Tardos, 2010]. The rationale behind it stems from the fact that, given the format of the utilities in the FPA (see (1)), it is arguably unreasonable to overbid, as bidding 0 will *always* result in at least the same utility. In game-theoretic terms, the overbidding strategy is *weakly dominated* by bidding 0, which can be interpreted as abstaining from the auction. These strategies are typically excluded from consideration to rule out unnatural equilibria (see [Feldman et al., 2020] for a discussion).

We are now ready to formally define our computational problem of finding an equilibrium of the FPA:

ε-Bayes-Nash Equilibrium in the First-Price Auction (ε-BNE-FPA)

Input:

- a set of bidders $N = \{1, 2, \ldots, n\}$;
- a finite bidding space $B \subseteq [0, 1]$;
- for each pair of bidders $i, j \in N$, a continuous value distribution $F_{i,j}$ over $[0, 1]$.

Output: An ε-Bayes-Nash equilibrium $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_n)$.

We will use the term exact-BNE-FPA instead of 0-BNE-FPA to denote the computational problem of finding an exact Bayes-Nash equilibrium of the auction. Some remarks related to the definition above are in order.

The Input Model for the Distributions. We have intentionally vaguely stated that the distributions $F_{i,j}$ should be provided as input to the problem, but we have not specified exactly how. Our positive results hold even when the functions $F_{i,j}$ are fairly general, and can be concisely and efficiently represented in a form that is appropriate for computation. In the interest of clarity, we omit the technical details here, and we refer the reader to Appendix A where we provide all the details of the input model. For the negative results on the other hand, we use fairly simple distributions $F_{i,j}$ – this

only makes our results stronger. In particular, we use *piecewise-constant* density functions, which can be represented by the endpoints and the value for each interval.

Explicit Bidding Space. We assume that the bidding space is explicitly given as part of the input. This assumption is required in Section 3 in order to show that we can compute best-responses efficiently. Even in the mildest of settings where the bidding space is given implicitly, computing best-responses turns out to be computationally and information-theoretically hard. We show this in Appendix B.

Equilibrium Representation. Besides the representation of the input, the output of our computational problem, i.e., the equilibrium of the FPA, should also be represented in some concise and efficient way. Following the standard literature of the problem, we will consider equilibria for which the strategy $\beta_i(v_i)$ of each bidder is a non-decreasing function of her value v_i (e.g., see [Athey, 2001; Maskin and Riley, 2000; Reny and Zamir, 2004] and [Krishna, 2009, Appendix G]) for which the existence of an equilibrium is always guaranteed [Athey, 2001]. These equilibria are in a sense the only "natural" ones, as, similar to the case of overbidding (see earlier discussion), any bidder's strategy is weakly dominated by a non-decreasing strategy.

Based on this, there is a straightforward and computationally efficient way of representing the best response of each player, as a step function with a finite set of "jump points", corresponding to the values at which the bidder "jumps" from one bid to the next [Athey, 2001]. Formally, we define

$$\alpha_i(b) = \sup\{v \mid \beta_i(v) \leq b\}. \tag{2}$$

Intuitively, $\alpha_i(b)$ is the largest value for which player i would bid b or lower. With a slight abuse of notation, we can write $\alpha_i = \beta_i^{-1}$, that is, α_i can be interpreted as an *inverse bidding* strategy. In that way, we can also rework β_i from α_i, as $\beta_i(v) = b$, where $v \in (\alpha_i(b^-), \alpha_i(b)]$ for any $b \in B$. Here we let b^- denote the previous bid, i.e., the largest $b' \in B$ with $b' < b$. Finally, to be able to handle the corner cases in a unified way, we set $a_i(b^-) = 0$ when $b = 0$. Notice also, that $\alpha_i(b) = 1$ when $b = \max B$.

In particular, this implies that bidding strategies are left-continuous (which is without loss of generality given our value distributions), as shown in Figure 1.

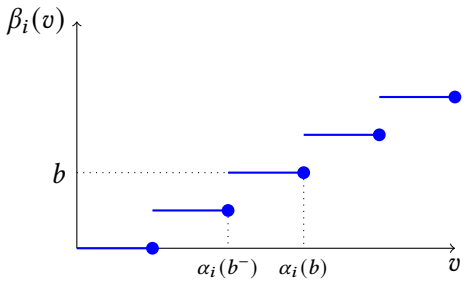

Figure 1: A monotone bidding strategy $\beta_i(\cdot)$ can be succinctly represented by its jump points, $\alpha_i(b)$ for $b \in B$.

Irrational Equilibria. As discussed in our introduction, for our PPAD-completeness result, we will be looking for an ε-approximate equilibrium, rather than an exact one. Of course, this only makes our hardness results even stronger; but besides that, it is actually very much necessary for our membership result in PPAD as well. In particular, as demonstrated by the example below, the FPA may have *only irrational* equilibria, even when all input parameters are rational numbers.

Example 1. Consider a FPA with $n = 3$ bidders and common priors, whose values are independently and identically distributed according to the uniform distribution on $[0, 1]$; that is, $F_i(x) = x$ for $i = 1, 2, 3$. Let the bidding space be $B = \{0, 1/2\}$. Clearly, this auction can be represented with piecewise-constant density functions (with a single piece) and with a finite number of rational quantities. It can be verified that the auction has a unique equilibrium, where a bidder bids 0 iff her valuation is below $\frac{-1+\sqrt{5}}{2} \approx 0.618$; therefore, the unique equilibrium is irrational. We provide the detailed derivation in Appendix C.

The appropriate setting for studying the computation of *exact* equilibria is the class FIXP of Etessami and Yannakakis [2010]. In Sections 4.2 and 5 we show that the problem of exact equilibrium computation of the FPA is FIXP-complete.

Further Notation. We conclude the section with the following terminology which will be useful in multiple sections of our paper. For $t_1 < t_2$, we will let $\mathrm{T}_{[t_1,t_2]}$ denote the *truncation* of a value x to $[t_1, t_2]$, i.e., $\mathrm{T}_{[t_1,t_2]}(x) = \max\{t_1, \min\{t_2, x\}\}$. Furthermore, for $k \in \mathbb{N}$ we sometimes use $[k]$ to denote $\{1, 2, \ldots, k\}$.

2.1 Outline

In Section 3 we provide a useful characterization of BNE and then show how to compute the best responses in polynomial time. In Section 4, first we provide a new existence proof via Brouwer's fixed point theorem, and then proceed to prove the membership of the equilibrium computation problems in PPAD and FIXP. In Section 5 we show the computational hardness for these classes. In Section 6 we present an efficient algorithm for a natural special case. We conclude with some interesting future directions in Section 7.

3 Equilibrium Characterization and Best Response Computation

In this section we begin by presenting a useful characterization of ε-BNE that is crucial for many parts of the paper. Then, we show how best-responses of bidders can be checked and computed in polynomial time. We remark that the reductions that we will construct in Section 4 to show the PPAD-membership and the FIXP-membership of the problem do not technically require the computation of the whole best-response function, but rather only the probabilities of winning the item given the bidder's bid and the bidding strategies of the other bidders. However, the best-response computation is interesting in its own right, and that is why we present this here.

Characterization. The following lemma essentially states that an ε-BNE is characterized by the behavior of the bidding function at the jump points. Recall that for any bid b, we let b^- denote the previous bid, and we use the convention $\alpha_i(b^-) = 0$ when $b = 0$. Notice that $\alpha_i(b) = 1$ when $b = \max B$. Furthermore, when $\alpha_i(b^-) = \alpha_i(b)$, the corresponding strategy β_i does not use bid b, but instead jumps from the previous bid directly to the following one. We will call all bids where this does not happen, i.e, we have $\alpha_i(b^-) < \alpha_i(b)$, *non-degenerate* for bidder i.

Lemma 3.1 (Characterization of ε-BNE)**.** *Fix an $\varepsilon \geq 0$. A strategy profile $\boldsymbol{\beta}$ is an ε-BNE of the FPA, if and only if, for every bidder i and every non-degenerate bid b,*

$$u_i(b, \boldsymbol{\beta}_{-i}; \alpha_i(b^-)) \geq u_i(b', \boldsymbol{\beta}_{-i}; \alpha_i(b^-)) - \varepsilon \qquad \text{for all } \; b' < b \tag{3}$$

and

$$u_i(b, \boldsymbol{\beta}_{-i}; \alpha_i(b)) \geq u_i(b', \boldsymbol{\beta}_{-i}; \alpha_i(b)) - \varepsilon \qquad \text{for all } \; b' > b. \tag{4}$$

The H-functions. Before proving this characterization, we introduce some useful notation. We use the term $H_i(b, \boldsymbol{\beta}_{-i})$ to denote the (perceived) probability that bidder i wins the item with bid b, when the other bidders use bids according to the bidding strategy $\boldsymbol{\beta}_{-i}$, i.e.,

$$H_i(b, \boldsymbol{\beta}_{-i}) = \Pr\left[\text{bidder } i \text{ wins} | b, \boldsymbol{\beta}_{-i}\right]$$

The utility can easily be expressed in terms of this function, namely $u_i(b, \boldsymbol{\beta}_{-i}; v_i) = (v_i - b) \cdot H_i(b, \boldsymbol{\beta}_{-i})$.

Proof of Lemma 3.1. ($\Rightarrow$): Fix a bidder i and a bid b with $\alpha_i(b^-) < \alpha_i(b)$. Since bidder i bids b inside the non-empty interval $(\alpha_i(b^-), \alpha_i(b)]$, and $\boldsymbol{\beta}$ is an ε-BNE, we get that $u_i(b, \boldsymbol{\beta}_{-i}; v_i) \geq u_i(b', \boldsymbol{\beta}_{-i}; v_i) - \varepsilon$ for every $v_i \in (\alpha_i(b^-), \alpha_i(b)]$ and $b' \neq b$. Since the utilities are continuous functions on v_i, the inequalities must also hold at the interval endpoints.

($\Leftarrow$): Suppose (3, 4) hold. Take any bidder i and any valuation v_i, and let $(\alpha_i(b^-), \alpha_i(b)]$ be the interval containing v_i. Notice that the utilities $u_i(b, \boldsymbol{\beta}_{-i}; v_i)$, $u_i(b', \boldsymbol{\beta}_{-i}; v_i)$ are linear functions on v_i, with slopes given by $H_i(b, \boldsymbol{\beta}_{-i})$, $H_i(b', \boldsymbol{\beta}_{-i})$ respectively. For $b' < b$, we know that $H_i(b', \boldsymbol{\beta}_{-i}) \leq H_i(b, \boldsymbol{\beta}_{-i})$ and $u_i(b, \boldsymbol{\beta}_{-i}; v) \geq u_i(b', \boldsymbol{\beta}_{-i}; v) - \varepsilon$ holds at $v = \alpha_i(b^-)$; therefore it must hold also at $v = v_i$. Similarly for $b' > b$, we know that $H_i(b', \boldsymbol{\beta}_{-i}) \geq H_i(b, \boldsymbol{\beta}_{-i})$ and $u_i(b, \boldsymbol{\beta}_{-i}; v) \geq u_i(b', \boldsymbol{\beta}_{-i}; v) - \varepsilon$ holds at $v = \alpha_i(b)$; therefore it must hold also at $v = v_i$. We thus conclude that $\boldsymbol{\beta}$ is an ε-BNE. □

We now consider the basic computational problems of checking and computing best-responses of bidders. We assume throughout that bidding strategies provided in the input are given via rational quantities corresponding to the jump points $\alpha_j(b)$, as defined in Section 2. The first step to be able to check or compute best-responses is the efficient computation of the H-functions defined above.

Computation of the H-functions. Recall that $H_i(b, \boldsymbol{\beta}_{-i}) = \Pr[\text{bidder } i \text{ wins} | b, \boldsymbol{\beta}_{-i}]$. This probability clearly depends on bidder i's prior on the other bidders' distributions, as well as on whether b is the highest bid, and if it is, how many other highest bids there are in the auction, in case of a tie. While the form of the functions H_i can be devised analytically, the expression involves exponentially many terms in the number of bidders n; therefore it is not obvious that it can be computed efficiently. The following lemma states that this is in fact possible.

Lemma 3.2. *Given a bidder i, a bid b and bidding strategies $\boldsymbol{\beta}_{-i}$ of the other bidders, the probability $H_i(b, \boldsymbol{\beta}_{-i})$ of bidder i winning the item can be computed in polynomial time.*

Proof. For ease of notation, we present the proof for bidder $i = n$. The cases for the other bidders are analogous and can be handled, e.g., via an appropriate relabeling. The probability that bidder n wins (given her bid and the bidding strategies of the other bidders) can be written as

$$H_n(b, \boldsymbol{\beta}_{-n}) = \sum_{k=0}^{n-1} \frac{1}{k+1} T(b, n-1, k), \tag{5}$$

where, for $0 \leq k \leq \ell \leq n-1$, we use $T(b, \ell, k)$ to denote the probability that *exactly* k out of the first ℓ bidders bid exactly b, and the remaining $\ell - k$ bidders all bid below b; in other words, for the special case where $\ell = n-1$ in the above expression, $T(b, n-1, k)$ is the probability of the highest bid being b, with $k+1$ bidders (including bidder n) being tied for the highest bid. Next, for a given bidder j, let

$$G_{j,b^-} = F_{n,j}(\alpha_j(b^-)) = \Pr\left[\beta_j(v_j) < b\right], \quad g_{j,b} = F_{n,j}(\alpha_j(b)) - G_{jb^-} = \Pr\left[\beta_j(v_j) = b\right]$$

denote the (perceived from the perspective of bidder n) probabilities that bidder j bids below b, and exactly b, respectively. Note that G_{j,b^-} and $g_{j,b}$ can be efficiently computed with access to $F_{n,j}$ and $\boldsymbol{\alpha}_{-n}$. Moreover, one could write

$$T(b, n-1, k) = \sum_{\substack{S \subseteq [n-1] \\ |S| = k}} \prod_{j \in S} g_{j,b} \cdot \prod_{j \notin S} G_{j,b^-}. \tag{6}$$

Notice that (6) does not yield an efficient way of computing the probabilities, as the number of summands can be exponential in n. To bypass this obstacle, we observe that, more generally, the probabilities $T(b, \ell, k)$ can be computed from G_{ℓ,b^-} and $g_{\ell,b}$ via dynamic programming, by conditioning on bidder ℓ's bid, in the following way:

$$\begin{aligned}
T(b, 0, 0) &= 1; \\
T(b, \ell, k) &= 0, && \text{for } k > \ell; \\
T(b, \ell+1, 0) &= T(b, \ell, 0) G_{\ell+1,b^-}; \\
T(b, \ell+1, k+1) &= T(b, \ell, k) g_{\ell+1,b} + T(b, \ell, k+1) G_{\ell+1,b^-}; && \text{for } k \leq \ell.
\end{aligned}$$

Thus, all values of $T(b, n-1, k)$, for $k = 0, \ldots, n-1$, can be computed with a total number of $O(n^2)$ recursive calls, so that $H_n(b, \boldsymbol{\beta}_{-n})$ can be computed in polynomial time. □

Lemma 3.2 implies that the utilities in (3, 4) of the characterization (Lemma 3.1) can be computed in polynomial time. Since there are $O(n|B|^2)$ inequalities to check in Lemma 3.1, we immediately conclude the following.

Corollary 3.3. *Given $\varepsilon \geq 0$, and a strategy profile $\boldsymbol{\beta}$ in a first-price auction with subjective priors, one can determine in polynomial time if $\boldsymbol{\beta}$ constitutes an ε-BNE.*

Using Lemma 3.2, we can now also efficiently compute best-responses, and, in fact, even *exact* best-responses (i.e., ε-best-responses for $\varepsilon = 0$).

Theorem 3.4. *In a first-price auction with subjective priors, the bidders' best-responses can be computed in polynomial time.*

Proof. Given a bidder i and the vector of bidding strategies $\boldsymbol{\beta}_{-i}$, one can compute in polynomial time the probabilities $H_i(b, \boldsymbol{\beta}_{-i})$ for each bid $b \in B$ using Lemma 3.2. Now recall that the utility of bidder i, when having a valuation of v_i and bidding b, is given by $u_i(b, \boldsymbol{\beta}_{-i}; v_i) = (v_i - b) \cdot H_i(b, \boldsymbol{\beta}_{-i})$, which is a linear function on v_i having slope $H_i(b, \boldsymbol{\beta}_{-i})$. Thus, maximizing the utility amounts to taking the maximum (or *upper envelope*) of $|B|$ linear functions; the result is a piecewise linear function whose jump points can be efficiently computed by solving linear equations. In particular, given bids $b < b'$, we can compute $\alpha = \tilde{\alpha}_i(b, b')$ as the solution of $u_i(b, \boldsymbol{\beta}_{-i}; \alpha) = u_i(b', \boldsymbol{\beta}_{-i}; \alpha)$, that is,

$$\tilde{\alpha}_i(b, b') = \begin{cases} \frac{b' H_i(b', \boldsymbol{\beta}_{-i}) - b H_i(b, \boldsymbol{\beta}_{-i})}{H_i(b', \boldsymbol{\beta}_{-i}) - H_i(b, \boldsymbol{\beta}_{-i})} & \text{if} \quad H_i(b', \boldsymbol{\beta}_{-i}) \neq H_i(b, \boldsymbol{\beta}_{-i}), \\ +\infty & \text{otherwise.} \end{cases}$$

Intuitively, $\tilde{\alpha}_i(b, b')$ is the jump point corresponding to bidding b versus bidding b': bidder i achieves higher utility by bidding b iff $v_i < \tilde{\alpha}_i(b, b')$. Now the highest value for which bidder i (weakly) prefers bidding b versus any other higher bid is $\min_{b' > b} \tilde{\alpha}_i(b, b')$; if at this valuation, bidding b also achieves higher utility than bidding any other lower bid, then $\min_{b' > b} \tilde{\alpha}_i(b, b')$ is indeed one of the desired jump points. Otherwise, b is a degenerate bid, in the sense that there is no valuation for which b is an optimal response. Therefore, the jump points introduced in (2) are given by $\alpha_i(b) = \max_{b' \leq b} \min_{b'' > b'} \tilde{\alpha}_i(b', b'')$.[5] Clearly then, the $\alpha_i(b)$ can be found in polynomial time. □

4 Existence and Membership in PPAD and FIXP

The existence of equilibria in our setting can essentially be established by adapting a proof by Athey [2001], which relies on Kakutani's fixed point theorem. Unfortunately, proofs that are based on this fixed point theorem cannot easily be turned into membership results for computational classes such as PPAD and FIXP. This is especially true for FIXP which is essentially defined as the class of all problems that can be solved by finding a Brouwer fixed point. In order to circumvent this obstacle we present a new proof that uses Brouwer's fixed point theorem. In this section, we first present this proof, and then utilize it to prove membership of our problems of interest in PPAD and FIXP.

4.1 Existence of Equilibria via Brouwer's Fixed Point Theorem

Theorem 4.1. *Every first-price auction with continuous subjective priors and finite bidding space admits a monotone non-decreasing and non-overbidding pure Bayes-Nash equilibrium.*

Proof. Let $N = \{1, 2, \ldots, n\}$ be the set of bidders, $F_{i,j}$ the continuous subjective priors, and $0 = b_0, b_1, \ldots, b_m$ be the ordered list of bids, i.e., the elements of $B \subseteq [0, 1]$. Recall that a monotone non-decreasing strategy $\beta_i : [0, 1] \to B$ can be represented by its jump points $\alpha_i(b)$. Let

$$\mathcal{D} = \{\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_n) \in ([0, 1]^m)^n \mid \forall i \in N, j \in [m] : \alpha_i(b_{j-2}) \leq \alpha_i(b_{j-1}) \wedge b_j \leq \alpha_i(b_{j-1})\}$$

[5] The maximization over $b' \leq b$ serves to exclude degenerate cases, e.g. if $b' < b < b''$ but $\tilde{\alpha}_i(b, b'') < \tilde{\alpha}_i(b', b'') < \tilde{\alpha}_i(b, b')$.

where we use the convention $\alpha_i(b_{-1}) := 0$ to keep the notation simple. The domain $\mathcal{D}$ is the set of all monotone non-decreasing non-overbidding strategy profiles, represented by their jump points. Note that $\mathcal{D}$ is compact and convex.

In what follows we slightly abuse notation by replacing the strategy profile $\boldsymbol{\beta}$ by its representation $\boldsymbol{\alpha}$ in some terms. Recall the functions $H_i(b, \boldsymbol{\alpha}_{-i})$ defined in Section 3, which represent the probability that bidder i wins the auction, if they bid b. By inspecting the proof of Lemma 3.2, it is easy to see that the quantities G_{jb^-} and g_{jb} are continuous with respect to $\boldsymbol{\alpha}_{-i}$, since the distributions are continuous. As a result, the terms $T(b, n-1, j)$ are also continuous in $\boldsymbol{\alpha}_{-i}$ (by (6)), which implies that $H_i(b, \boldsymbol{\beta}_{-i})$ is also continuous in $\boldsymbol{\alpha}_{-i}$. Since the utility functions can be written as $u_i(b, \boldsymbol{\alpha}_{-i}; v_i) = (v_i - b) \cdot H_i(b, \boldsymbol{\alpha}_{-i})$, it follows that the functions $(\boldsymbol{\alpha}_{-i}, v_i) \mapsto u_i(b, \boldsymbol{\alpha}_{-i}; v_i)$ are continuous.

We now construct a function $G : \mathcal{D} \to \mathcal{D}$. For any bidder $i \in N$ and any $j \in [m]$, define the continuous function $\Delta^i_j : \mathcal{D} \to \mathbb{R}$ by

$$\Delta^i_j(\boldsymbol{\alpha}) = u_i(b_{j-1}, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})) - \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})).$$

The intuition behind the construction of $\Delta^i_j(\boldsymbol{\alpha})$ is as follows. We consider the point $\alpha_i(b_{j-1})$ (the last point where bidder i currently bids b_{j-1}) and compare how beneficial it is for bidder i to bid b_{j-1} at this point, compared to using a larger bid. The sign of $\Delta^i_j(\boldsymbol{\alpha})$ essentially encodes the result of this comparison. As a result, it encodes whether bidder i: (a) is happy with the current value of $\alpha_i(b_{j-1})$ (i.e., $\Delta^i_j(\boldsymbol{\alpha}) = 0$), (b) would like to increase $\alpha_i(b_{j-1})$ (i.e., $\Delta^i_j(\boldsymbol{\alpha}) > 0$), or (c) would like to decrease $\alpha_i(b_{j-1})$ (i.e., $\Delta^i_j(\boldsymbol{\alpha}) < 0$).

Now, for any $\boldsymbol{\alpha} \in \mathcal{D}$, let $G(\boldsymbol{\alpha}) = \boldsymbol{\alpha}'$, where for all $i \in N$ and $j = 1, 2, \ldots, m$ (consecutively and in that order)

$$\alpha'_i(b_{j-1}) = \mathrm{T}_{[\max\{b_j, \alpha'_i(b_{j-2})\}, 1]}(\alpha_i(b_{j-1}) + \Delta^i_j(\boldsymbol{\alpha})). \tag{7}$$

Note in particular that this is well-defined, since $\alpha'_i(b_{j-2})$ is defined before $\alpha'_i(b_{j-1})$. The truncation operator immediately ensures that $\boldsymbol{\alpha}' \in \mathcal{D}$. Since G is also clearly continuous, and $\mathcal{D}$ is compact and convex, it follows by Brouwer's fixed point theorem that there exists a $\boldsymbol{\alpha} \in \mathcal{D}$ with $G(\boldsymbol{\alpha}) = \boldsymbol{\alpha}$. It remains to prove that $\boldsymbol{\alpha}$ corresponds to an equilibrium of the auction.

Consider some bidder $i \in N$. We will show that α_i is a best-response to $\boldsymbol{\alpha}_{-i}$ using the characterization of Lemma 3.1. Consider any non-empty interval of non-empty interior $[\alpha_i(b_{j-1}), \alpha_i(b_j)]$, for some $j \in \{0, 1, \ldots, m\}$, where we use the convention that $\alpha_i(b_{-1}) = 0$, and recall that $\alpha_i(b_m) = 1$.

- First, we show that $u_i(b_j, \boldsymbol{\alpha}_{-i}; \alpha_i(b_j)) \geq \max_{\ell > j} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_j))$. Clearly, for $j = m$ this holds trivially. For $j < m$, this can immediately be rephrased as showing $\Delta^i_{j+1}(\boldsymbol{\alpha}) \geq 0$. Now, note that by assumption we have $\alpha_i(b_j) > \alpha_i(b_{j-1})$. Thus, since $\alpha_i(b_j)$ remains fixed under G, it must be that $\alpha_i(b_j) = b_{j+1}$ or $\Delta^i_{j+1}(\boldsymbol{\alpha}) \geq 0$. However, if $\alpha_i(b_j) = b_{j+1}$, then it also trivially holds that $\Delta^i_{j+1}(\boldsymbol{\alpha}) \geq 0$.

- Next, we show that $u_i(b_j, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})) \geq \max_{\ell < j} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1}))$. Again, this holds trivially for $j = 0$, so we now consider $j > 0$. By the first bullet above, it holds that

  $$u_i(b_j, \boldsymbol{\alpha}_{-i}; \alpha_i(b_j)) = \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_j)).$$

  By the monotonicity of the H-functions, we can simply replace $\alpha_i(b_j)$ by $\alpha_i(b_{j-1})$ in the equation above. Indeed, since by definition of the H-functions we have that $H_i(b, \boldsymbol{\alpha}_{-i}) \leq H_i(b', \boldsymbol{\alpha}_{-i})$ for any two bids $b \leq b'$, it follows in particular that the function $v \mapsto u_i(b, \boldsymbol{\alpha}_{-i}; v) - u_i(b', \boldsymbol{\alpha}_{-i}; v)$ is monotonically non-increasing. To see this, it suffices to write the utilities in terms of the H-functions, i.e., $u_i(b, \boldsymbol{\alpha}_{-i}; v) = (v - b) \cdot H_i(b, \boldsymbol{\alpha}_{-i})$, and similarly for b'. As a result, we obtain that

  $$u_i(b_j, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})) = \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})).$$

On the other hand, since $\alpha_i(b_{j-1}) < \alpha_i(b_j)$, it follows in particular that $\alpha_i(b_k) < 1$ for all $k < j$. As a result, since $\alpha_i(b_k)$ remains fixed under G, it must be that $\Delta^i_{k+1}(\boldsymbol{\alpha}) \leq 0$ for all $k < j$, i.e.,

$$u_i(b_k, \boldsymbol{\alpha}_{-i}; \alpha_i(b_k)) \leq \max_{\ell \geq k+1} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_k))$$

which by monotonicity of the H-functions, as explained above, continues to hold if we replace $\alpha_i(b_k)$ by $\alpha_i(b_{j-1})$, i.e., for all $k < j$ we have

$$u_i(b_k, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})) \leq \max_{\ell \geq k+1} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})).$$

As a result it follows by induction that for all $k < j$

$$u_i(b_k, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})) \leq \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})) = u_i(b_j, \boldsymbol{\alpha}_{-i}; \alpha_i(b_{j-1})).$$

By Lemma 3.1, it immediately follows that α_i is a best-response to $\boldsymbol{\alpha}_{-i}$. Since this holds for all bidders $i \in N$, $\boldsymbol{\alpha}$ is an equilibrium. □

4.2 FIXP Membership

In order to study the exact equilibrium problem for the first-price auction in the context of FIXP, we consider the model where the distributions $F_{i,j}$ are given by algebraic circuits using the operations $\{+, -, \times, /, \max, \min, \sqrt[k]{\cdot}\}$ and rational constants, as is usual in this setting [Etessami and Yannakakis, 2010]. We show that the proof of existence in the previous section can be turned into a reduction.

Theorem 4.2. *The problem* EXACT-BNE-FPA *lies in* FIXP.

Proof. Clearly, the domain $\mathcal{D}$ of the function $G : \mathcal{D} \to \mathcal{D}$ from the proof of Theorem 4.1 can be represented by a set of linear inequalities that can be constructed in polynomial time in n, m and the representation length of B. Thus, it remains to show that we can construct in polynomial time an algebraic circuit that computes G.

We now describe how to construct a circuit for G that only uses operations $\{+, -, \times, /, \max, \min, \sqrt[k]{\cdot}\}$ and rational constants. First of all, note that probabilities of the form $\Pr_{v_j \sim F_{i,j}}[\beta_j(v_j) \leq b] = F_{i,j}(\alpha_j(b))$ can easily be computed by the circuit, since the (cumulative) distribution functions $F_{i,j}$ are provided as algebraic circuits, and $\boldsymbol{\alpha}$ is the input to the circuit. It follows that the quantities G_{jb^-} and g_{jb} defined in the proof of Lemma 3.2 can also be computed by the circuit. As a result, we can use the dynamic programming procedure described in the proof of Lemma 3.2, to compute the terms $T(b, n-1, j)$ by only using a polynomial number of operations. Note in particular, that the dynamic programming assignment rules can all be implemented using the available set of operations. With the terms $T(b, n-1, j)$ in hand, we can then easily compute the terms $H_i(b, \boldsymbol{\alpha}_{-i})$ for all $b \in B$, and thus evaluate the utility function $u_i(b, \boldsymbol{\alpha}_{-i}; v_i) = (v_i - b) \cdot H_i(b, \boldsymbol{\alpha}_{-i})$ at any given $v_i \in [0, 1]$. Finally, using the utility functions and the max operation we can now compute the terms $\Delta^i_j(\boldsymbol{\alpha})$ from the proof of Theorem 4.1, and then using $+$, max, min and the constant 1 we can output $\boldsymbol{\alpha}' = G(\boldsymbol{\alpha})$ by noting that

$$\alpha'_i(b_{j-1}) = \max\{\max\{b_j, \alpha'_i(b_{j-2})\}, \min\{1, \alpha_i(b_{j-1}) + \Delta^i_j(\boldsymbol{\alpha})\}\}.$$ □

4.3 PPAD Membership

In order to study the approximate equilibrium problem for the first-price auction in the context of PPAD, we consider a model where the distributions $F_{i,j}$ are *polynomially-computable*, i.e., can be evaluated in polynomial time by a Turing machine.[6] In order to guarantee that an approximate equilibrium with

[6] Note that a function represented as an algebraic circuit (as in the previous section on FIXP) is not necessarily polynomially-computable, e.g., because the circuit can use "repeated squaring" to construct numbers with exponential bit complexity. Conversely, a function that is polynomially-computable cannot necessarily be represented as an algebraic circuit, because a Turing machine is not restricted to using arithmetic gates. We note that these two different models for representing functions are standard for FIXP and PPAD respectively.

polynomial bit complexity exists, we also assume that the distributions are *polynomially continuous*. For a formal definition of these two standard properties in the context of PPAD, see Appendix A. In this section we show that in this model, the problem of computing an ε-BNE lies in the class PPAD. We begin by observing that the polynomial-continuity of the distribution functions $F_{i,j}$ implies that the utility functions are also polynomially-continuous. This is proved in Appendix D.

Lemma 4.3. *If the distributions $F_{i,j}$ are polynomially-continuous, then so are the utility functions $\boldsymbol{\alpha} \mapsto u_i(b, \boldsymbol{\alpha}_{-i}; v_i)$. In more detail, given $\varepsilon > 0$, we can in polynomial time compute $\delta > 0$ such that for all $i \in N$, $b \in B$ and $v_i \in [0, 1]$*

$$\|\boldsymbol{\alpha} - \boldsymbol{\alpha}'\|_\infty \leq \delta \implies |u_i(b, \boldsymbol{\alpha}_{-i}; v_i) - u_i(b, \boldsymbol{\alpha}'_{-i}; v_i)| \leq \varepsilon.$$

In particular, δ can be represented using a polynomial number of bits.

We are now ready to state the main result of this section.

Theorem 4.4. *The problem* ε-BNE-FPA *lies in* PPAD.

Proof. We show that the existence proof of Theorem 4.1 can be turned into a polynomial-time many-one reduction to the problem of computing an approximate Brouwer fixed point of a polynomially-computable and polynomially-continuous function over a bounded polytope given by linear inequalities, known to lie in PPAD [Etessami and Yannakakis, 2010, Proposition 2].

Since the distributions $F_{i,j}$ are polynomially-computable, and by the arguments provided in the proof of Theorem 4.2 (including the dynamic programming procedure from Lemma 3.2), it immediately follows that G is polynomially-computable. The polynomial-continuity of G also immediately follows from the polynomial-continuity of the utility functions (Lemma 4.3). Thus, the problem of computing an approximate fixed point of G lies in PPAD.

Given $\varepsilon > 0$, by Lemma 4.3 we can compute $\delta > 0$ so that for all $i \in N$, $b \in B$ and $v_i \in [0, 1]$

$$\|\boldsymbol{\alpha} - \boldsymbol{\alpha}'\|_\infty \leq \delta \implies |u_i(b, \boldsymbol{\alpha}_{-i}; v_i) - u_i(b, \boldsymbol{\alpha}'_{-i}; v_i)| \leq \frac{\varepsilon}{16m}.$$

Now consider any δ-approximate fixed point of G, i.e., $\boldsymbol{\alpha} \in \mathcal{D}$ such that $\|G(\boldsymbol{\alpha}) - \boldsymbol{\alpha}\|_\infty \leq \delta$. Let $\boldsymbol{\alpha}' = G(\boldsymbol{\alpha})$. We prove that $\boldsymbol{\alpha}'$ is an ε-approximate equilibrium of the first-price auction. This shows that ε-BNE-FPA reduces to the Brouwer fixed point computation problem, and thus lies in PPAD.

Since $\boldsymbol{\alpha}' = G(\boldsymbol{\alpha})$ and $\|G(\boldsymbol{\alpha}) - \boldsymbol{\alpha}\|_\infty \leq \delta$, it holds that $\|\boldsymbol{\alpha} - \boldsymbol{\alpha}'\|_\infty \leq \delta$ and thus

$$|u_i(b, \boldsymbol{\alpha}_{-i}; v_i) - u_i(b, \boldsymbol{\alpha}'_{-i}; v_i)| \leq \frac{\varepsilon}{16m} \tag{8}$$

for all $i \in N$, $b \in B$ and $v_i \in [0, 1]$. In particular, we also have that $|\Delta^i_j(\boldsymbol{\alpha}) - \Delta^i_j(\boldsymbol{\alpha}')| \leq 2(\varepsilon/16m + \delta) \leq \varepsilon/4m$ (since the utility functions are also 1-Lipschitz with respect to v_i). Note that here we assumed without loss of generality that $\delta \leq \varepsilon/16m$.

Fix some bidder $i \in N$. Consider any non-empty interval $[\alpha'_i(b_{j-1}), \alpha'_i(b_j)]$ for some $j \in \{0, 1, \ldots, m\}$, where we use the convention that $\alpha'_i(b_{-1}) = 0$, and recall that $\alpha'_i(b_m) = 1$.

- First, we show that $u_i(b_j, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_j)) \geq \max_{\ell > j} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_j)) - \varepsilon/2$. Clearly, for $j = m$ this holds trivially. For $j < m$, this can immediately be rephrased as showing $\Delta^i_{j+1}(\boldsymbol{\alpha}') \geq -\varepsilon/2$. By (8), it suffices to show that $\Delta^i_{j+1}(\boldsymbol{\alpha}) \geq -\varepsilon/2 + \varepsilon/4m$. But if $\Delta^i_{j+1}(\boldsymbol{\alpha}) < -\varepsilon/2 + \varepsilon/4m \leq -\varepsilon/16m \leq -\delta$, then by construction of G, since $|\alpha_i(b_j) - \alpha'_i(b_j)| \leq \delta$, it must be that $\alpha'_i(b_j) = b_{j+1}$ or $\alpha'_i(b_j) = \alpha'_i(b_{j-1})$. In the former case, it trivially holds that $\Delta^i_{j+1}(\boldsymbol{\alpha}') \geq 0 \geq -\varepsilon/2$. The latter case is impossible, since we assumed that $\alpha'_i(b_{j-1}) < \alpha'_i(b_j)$.

- Next, we show that $u_i(b_j, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) \geq \max_{\ell < j} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) - \varepsilon$. Again, this holds trivially for $j = 0$, so we now consider $j > 0$. By the first bullet above, it holds that

$$u_i(b_j, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_j)) \geq \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_j)) - \varepsilon/2.$$

As a result, by the monotonicity of the H-functions, as explained in the proof of Theorem 4.1, this continues to hold if we replace $\alpha'_i(b_j)$ by $\alpha'_i(b_{j-1})$, i.e.,

$$u_i(b_j, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) \geq \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) - \varepsilon/2. \tag{9}$$

On the other hand, since $\alpha'_i(b_{j-1}) < \alpha'_i(b_j)$, it follows in particular that $\alpha'_i(b_k) < 1$ for all $k < j$. As a result, by construction of G, and since $|\alpha_i(b_k) - \alpha'_i(b_k)| \leq \delta$, it must be that $\Delta^i_{k+1}(\boldsymbol{\alpha}) \leq \delta$ for all $k < j$. By (8) it follows that $\Delta^i_{k+1}(\boldsymbol{\alpha}') \leq \delta + \varepsilon/4m \leq \varepsilon/2m$ for all $k < j$, which yields

$$u_i(b_k, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_k)) \leq \max_{\ell \geq k+1} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_k)) + \frac{\varepsilon}{2m}$$

which by monotonicity of the H-functions, as explained in the proof of Theorem 4.1, continues to hold if we replace $\alpha'_i(b_k)$ by $\alpha'_i(b_{j-1})$, i.e., for all $k < j$ we have

$$u_i(b_k, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) \leq \max_{\ell \geq k+1} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) + \frac{\varepsilon}{2m}.$$

As a result it follows by induction that for all $k < j$

$$u_i(b_k, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) \leq \max_{\ell \geq j} u_i(b_\ell, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) + (j-k)\frac{\varepsilon}{2m}$$

which together with (9) yields that for all $k < j$

$$u_i(b_k, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) \leq u_i(b_j, \boldsymbol{\alpha}'_{-i}; \alpha'_i(b_{j-1})) + m\frac{\varepsilon}{2m} + \frac{\varepsilon}{2}.$$

By Lemma 3.1, it immediately follows that α'_i is an ε-best-response to $\boldsymbol{\alpha}'_{-i}$. Since this holds for all bidders $i \in N$, $\boldsymbol{\alpha}'$ is an ε-equilibrium. □

5 Computational Hardness

In this section we prove computational hardness results for the problem of computing an equilibrium of a first-price auction with subjective priors. Namely, we show that computing an ε-BNE is PPAD-hard, while computing an exact BNE is FIXP-hard. Our computational hardness results are particularly robust, because they hold even if we apply all of the following restrictions:

- the bidding space is $B = \{0, 1/5, 2/5, 3/5, 4/5\}$,
- the value distributions $F_{i,j}$ are given by very simple piecewise constant density functions,
- ε is some sufficiently small *constant.* *(only relevant for ε-BNE)*

In particular, by a simple rescaling argument, the hardness results also hold when the bidding space consists of all monetary amounts that are increments of some fixed denomination (e.g., one cent) up to some number m.[7] For example, there exists a sufficiently small constant ε such that it is PPAD-hard to compute an ε-BNE when the bidding space is $B = \{0, 1/100, 2/100, \ldots, 99/100, 1, 101/100, \ldots, m - 1/100, m\}$.

Together with the corresponding membership results proved in the previous section (Theorems 4.2 and 4.4), we thus obtain the following two theorems, which are the main results of this paper.

Theorem 5.1. *There exists a constant $\varepsilon > 0$ such that the problem ε-BNE-FPA is PPAD-complete.*

[7] Note that m should be provided in the input in *unary* representation. This is necessary to ensure that the bidding space has polynomial size, thus allowing efficient computation of best-responses. See the discussion in Section 2 regarding our assumption of an explicit bidding space.

Theorem 5.2. *The problem* EXACT-BNE-FPA *is* FIXP*-complete.*

In the rest of this section, we present the proof of our hardness results. A nice feature of our proof is that we provide a *single* reduction to prove both PPAD- and FIXP-hardness. In more detail, we reduce from the so-called *Generalized Circuit problem*, which has been instrumental for proving PPAD-hardness results for Nash equilibrium computation problems [Daskalakis et al., 2009; Chen et al., 2009; Rubinstein, 2018]. In fact, we show that it suffices to consider significantly restricted versions of the Generalized Circuit problem when proving hardness results, and that an exact version of the problem can also be used to prove FIXP-hardness. Since we believe that these points may be of independent interest for future works, they are presented separately in Section 5.1. Our reduction from this problem to equilibrium computation in first-price auctions is then presented in Section 5.2.

5.1 The Generalized Circuit Problem

Generalized circuits, defined by Chen et al. [2009], can be viewed as a generalization of arithmetic circuits where we also allow *cycles*. This means that instead of representing a function, a generalized circuit represents a certain kind of constraint satisfaction problem. Indeed, the goal in the Generalized Circuit problem is to assign a value to each gate of the circuit such that all the gates are (approximately) satisfied. Importantly, gates are only allowed to take values in $[0, 1]$ and arithmetic operations are truncated accordingly. As a result, it can be shown that by Brouwer's fixed point theorem, there always exists an assignment of values that satisfies all the gates. However, computing even an approximate assignment is already PPAD-hard, i.e., essentially as hard as any Brouwer fixed point computation. We now provide some formal definitions.

Definition 2. A *generalized circuit*[8] with gate-types $\mathcal{G}$ is a list of gates $g_1, g_2, \ldots, g_{\mathfrak{m}}$. Every gate g_i is a 3-tuple $g_i = (G, j, k)$, where $G \in \mathcal{G}$ is the type of the gate, and $j, k \in [\mathfrak{m}] = \{1, \ldots, \mathfrak{m}\}$ are the indices of the input gates g_j, g_k (i, j, k distinct).

Before describing possible types of gates, we introduce some notation. Let $\mathrm{T} = \mathrm{T}_{[0,1]}$. Furthermore, we use the notation $x = y \pm \varepsilon$ to denote that $|x - y| \leq \varepsilon$.

Consider a generalized circuit $g_1, g_2, \ldots, g_{\mathfrak{m}}$ and an assignment $\mathbf{v} : [\mathfrak{m}] \to [0, 1]$ of values to its gates. We say that a gate is ε-satisfied by the assignment, if the constraint imposed by this gate is satisfied with error at most ε. The constraint that a gate $g_i = (G, j, k)$ must satisfy depends on its gate-type $G \in \mathcal{G}$, e.g.,

- if $G = G_1$, then $\mathbf{v}[g_i] = 1 \pm \varepsilon$ *(constant 1)*
- if $G = G_+$, then $\mathbf{v}[g_i] = \mathrm{T}(\mathbf{v}[g_j] + \mathbf{v}[g_k]) \pm \varepsilon$ *(addition)*
- if $G = G_-$, then $\mathbf{v}[g_i] = \mathrm{T}(\mathbf{v}[g_j] - \mathbf{v}[g_k]) \pm \varepsilon$ *(subtraction)*
- if $G = G_{1-}$, then $\mathbf{v}[g_i] = 1 - \mathbf{v}[g_j] \pm \varepsilon$ *(complement)*
- if $G = G_{\times 2}$, then $\mathbf{v}[g_i] = \mathrm{T}(2 \cdot \mathbf{v}[g_j]) \pm \varepsilon$ *(multiplication by 2)*
- if $G = G_\times$, then $\mathbf{v}[g_i] = \mathbf{v}[g_j] \cdot \mathbf{v}[g_k] \pm \varepsilon$ *(multiplication)*
- if $G = G_{(\cdot)^2}$, then $\mathbf{v}[g_i] = (\mathbf{v}[g_j])^2 \pm \varepsilon$ *(square)*

We are now ready to define the associated computational problem.

[8] Note that in the usual definition of generalized circuits, every gate also contains a rational parameter $\zeta \in [0, 1]$, which is used by some gate-types, e.g., a gate performing multiplication by the constant ζ. In our definition, gates do not contain this rational parameter, because, as we show in Propositions 5.3 and 5.4, these gate-types are actually not needed for the problems to be hard.

Definition 3. Let $\varepsilon > 0$. The problem ε-GCIRCUIT with gate-types $\mathcal{G}$ is defined as follows: given a generalized circuit $g_1, g_2, \ldots, g_{\mathfrak{m}}$ with gate-types $\mathcal{G}$, find an assignment $\mathbf{v} : [\mathfrak{m}] \to [0,1]$ to the gates such that they are all ε-satisfied.

Rubinstein [2018] proved that this problem is PPAD-complete for some sufficiently small constant $\varepsilon > 0$ and a relatively large set of gate-types $\mathcal{G}$. In Appendix E.1, we prove that the problem remains hard, even with a very restricted set of gate-types.

Proposition 5.3. *There exists a constant $\varepsilon > 0$ such that the problem ε-GCIRCUIT with gate-types $\mathcal{G} = \{G_+, G_{1-}\}$ is* PPAD*-complete. This continues to hold if we instead take $\mathcal{G} = \{G_1, G_-\}$.*

We can also define a problem EXACT-GCIRCUIT, where the goal is to find an assignment that *exactly* satisfies all constraints (i.e., with $\varepsilon = 0$). In Appendix E.2, we prove the following result.

Proposition 5.4. *The problem* EXACT-GCIRCUIT *with gate-types $\mathcal{G} = \{G_{1-}, G_+, G_{(\cdot)^2}\}$ is* FIXP*-complete. This continues to hold if we instead take $\mathcal{G} = \{G_{1-}, G_{\times 2}, G_\times\}$.*

5.2 Reduction to BNE-FPA

In this section, we present a reduction that achieves the following: given a generalized circuit, it constructs (in polynomial time) an instance of the first-price auction problem, such that for all $\varepsilon \in [0, 1/10^5]$, from any ε-BNE we can extract an 500ε-satisfying assignment for the generalized circuit. Furthermore, this "extraction" of the assignment from an ε-BNE can be done efficiently and, in fact, using a simple so-called separable linear transformation. This ensures that in the case $\varepsilon = 0$, we obtain a so-called SL-reduction from EXACT-GCIRCUIT, which yields the FIXP-hardness result [Etessami and Yannakakis, 2010]. If we let $\tilde{\varepsilon} > 0$ be a constant such that $\tilde{\varepsilon}$-GCIRCUIT is PPAD-hard, then for $\varepsilon = \min\{1/10^5, \tilde{\varepsilon}/500\}$ the reduction is a valid polynomial-time many-one reduction, which yields the PPAD-hardness result.

An obstacle to obtaining the desired reduction is that it is unclear how to simulate a G_+-gate or a $G_\times$-gate. As a result, we reduce from the GCIRCUIT problem with gate-types $\mathcal{G} = \{G_{\times 2}, G_{1-}, G_\phi\}$, where $\phi : [0,1]^2 \to [0,1]$, $(x, y) \mapsto \frac{1}{4}(x+1)(y+1)$. This means that a gate $g_i = (G_\phi, j, k)$ enforces the constraint $\mathbf{v}[g_i] = \phi(\mathbf{v}[g_j], \mathbf{v}[g_k]) \pm \varepsilon$. In Appendix E.3 we prove that this set of gate-types is sufficient for our desired hardness results.

Lemma 5.5. *Let $\mathcal{G} = \{G_{\times 2}, G_{1-}, G_\phi\}$. There exists a constant $\tilde{\varepsilon} > 0$ such that the problem $\tilde{\varepsilon}$-GCIRCUIT with gate-types $\mathcal{G}$ is* PPAD*-complete. Furthermore,* EXACT-GCIRCUIT *with gate-types $\mathcal{G}$ is* FIXP*-complete.*

The reduction. We begin with a high-level description of the reduction. Consider a generalized circuit $g_1, g_2, \ldots, g_{\mathfrak{m}}$ with gate-types $\mathcal{G} = \{G_{\times 2}, G_{1-}, G_\phi\}$. We construct a first-price auction with bidding space $B = \{0, 1/5, 2/5, 3/5, 4/5\}$ and a set of bidders $N = \{1, 2, \ldots, n\}$ where $n = 10\mathfrak{m}$. For every $i \in [\mathfrak{m}]$, bidder i will "correspond" to gate g_i, in the sense that, in any ε-BNE $\boldsymbol{\beta}$, the position of the second jump point of β_i, i.e., $\alpha_i(1/5)$ will encode the value $\mathbf{v}[g_i]$ that we will assign to gate g_i. Thus, we will refer to the bidders $1, 2, \ldots, \mathfrak{m}$ as *gate-bidders*. The rest of the bidders will be used as intermediate steps to enforce the desired constraints on the strategies of the gate-bidders. Accordingly, we will refer to them as *auxiliary-bidders*. Note that for every gate-bidder, there are 9 auxiliary-bidders available (if needed). For convenience, we describe the construction with the value space $[0, 5]$ instead of $[0, 1]$. This is without consequence, since this re-scaling of the instance simply means that we have to replace ε by 5ε at the end. Note that as a result of the re-scaling, the bidding space is now simply $B = \{0, 1, 2, 3, 4\}$.

Valid strategies and encoded value. Let $\boldsymbol{\beta}$ be any ε-BNE of the auction. A bidder $i \in N$ is said to be *valid*, if $\alpha_i(0) \in [1, 1 + 1/2]$, $\alpha_i(1) \in [2 + 1/3 - 2\varepsilon, 2 + 2/3 + 2\varepsilon]$, $\alpha_i(2) \in [3 + 1/2, 5]$ and $\alpha_i(3) = 5$. The bidder i is *almost-valid*, if the condition on $\alpha_i(1)$ is relaxed to $\alpha_i(1) \in [2, 3]$. For every bidder $i \in N$, we define the value encoded by bidder i according to $\boldsymbol{\beta}$, as

$$\mathbf{v}_{\boldsymbol{\beta}}[i] = \begin{cases} \mathrm{T}_{[0,1]}(3(\alpha_i(1) - 2 - 1/3)) & \text{if } i \text{ is valid,} \\ \text{null} & \text{otherwise.} \end{cases}$$

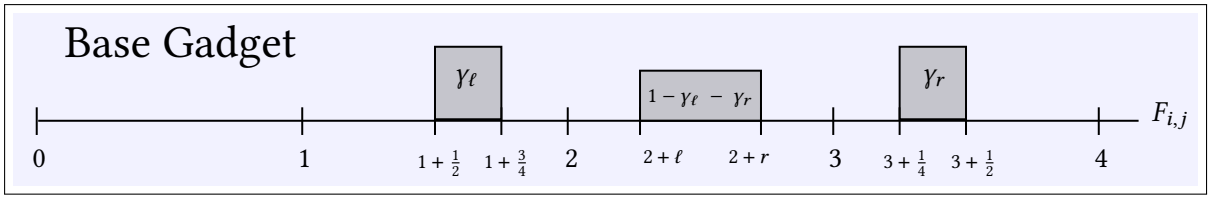

Figure 2: An illustration of the base gadget. The density of $F_{i,j}$ is depicted. When $\gamma_\ell = \gamma_r = 1/3$, $\ell = 1/3$ and $r = 2/3$ we obtain a *standard* base gadget, which essentially (approximately) "copies" the value $\mathbf{v}[i]$ of the input bidder i to the value $\mathbf{v}[j]$ of the output bidder j.

Note that we always have $\mathbf{v}_{\boldsymbol{\beta}}[i] \in [0,1] \cup \{\text{null}\}$. In the rest of the proof, we drop the subscript $\boldsymbol{\beta}$, since it is understood from the context. Our construction will ensure that for all $i \in [\mathfrak{m}]$, bidder i is valid and as a result $\mathbf{v}[i] \in [0,1]$. Furthermore, letting $\mathbf{v}[g_i] := \mathbf{v}[i]$ will yield an 100ε-satisfying assignment of the generalized circuit.

Gadgets. The rest of the proof describes the construction of the distribution functions $F_{i,j}$. We begin by constructing some *unary* gadgets. A unary gadget has a single "input" bidder $j \in N$ and an output bidder $i \in N \setminus \{j\}$. The goal of such a gadget is to establish a constraint on β_i that depends on β_j, but not on the strategy of any other bidder. This is achieved by setting $F_{i,k}$ for all $k \in N \setminus \{i,j\}$, such that its (piecewise constant) density function has a single piece of volume 1 lying in $[0,1]$. As a result, because of the no-overbidding assumption, bidder i will believe that all bidders $k \in N \setminus \{i,j\}$ bid 0 with probability 1. The behavior of the gadget is then determined by the precise construction of $F_{i,j}$.

Base Gadget. The base gadget is a unary gadget with input bidder j and output bidder i that has four parameters $\gamma_\ell, \gamma_r, \ell, r \in [0,1]$ with $\gamma_\ell + \gamma_r < 1$ and $r - \ell > 0$. The piecewise constant density function of $F_{i,j}$ is defined as follows. There is a piece of volume γ_ℓ in the interval $[1+1/2, 1+3/4]$, a piece of volume $1-\gamma_\ell-\gamma_r$ in $[2+\ell, 2+r]$, and finally a piece of volume γ_r in $[3+1/4, 3+1/2]$. See Figure 2 for an illustration.

When the parameters are $(\gamma_\ell, \gamma_r, \ell, r) = (1/3, 1/3, 1/3, 2/3)$, we call this the *standard* base gadget. It will immediately follow from Claim 1 below that if the input bidder j of the standard base gadget is valid, then so is the output bidder i, and furthermore $\mathbf{v}[i] = \mathbf{v}[j] \pm 6\varepsilon$. In other words, this gadget can be used to copy the value encoded by one bidder onto some other bidder.

Claim 1. *Let $\gamma_\ell, \gamma_r, \ell, r \in [0,1]$ with $\gamma_\ell, \gamma_r \geq 1/20$, $\gamma_\ell + \gamma_r < 1$ and $\ell < r$. Consider a base gadget with input bidder j and output bidder i, and parameters $(\gamma_\ell, \gamma_r, \ell, r)$. It holds that:*

- *If the input bidder j is almost-valid, then the output bidder i is also almost-valid.*
- *If $\gamma_\ell, \gamma_r \geq 1/3$ and j is almost-valid, then i is valid and*

$$\mathbf{v}[i] = (3\gamma_\ell - 1) + 3(1-\gamma_\ell-\gamma_r)\frac{T_{[2+\ell,2+r]}(\alpha_j(1)) - (2+\ell)}{r-\ell} \pm 6\varepsilon.$$

Proof. We begin by obtaining some equations that will be useful for various proofs in this section. Consider any unary gadget with input bidder j and output bidder i. To simplify notation, for $b \in \{0,1,2,3,4\}$, let p_b be the probability that bidder j bids b, as perceived by bidder i. Formally,

$$p_b := \Pr_{v_j \sim F_{i,j}}[\beta_j(v_j) = b] = \begin{cases} F_{i,j}(\alpha_j(b)) - F_{i,j}(\alpha_j(b-1)) & \text{if } b \in \{1,2,3,4\} \\ F_{i,j}(\alpha_j(0)) & \text{if } b = 0. \end{cases}$$

Recall the quantity $H_i(b, \boldsymbol{\beta}_{-i})$ defined in Section 3, which represents the probability that bidder i wins the auction if she bids b, and the other bidders act according to $\boldsymbol{\beta}_{-i}$. We drop $\boldsymbol{\beta}_{-i}$ from the notation, since it is clear from the context. Going back to our unary gadget, it is easy to see that $H_i(0) = p_0/n$, $H_i(1) = p_0 + p_1/2$, $H_i(2) = p_0 + p_1 + p_2/2$, $H_i(3) = p_0 + p_1 + p_2 + p_3/2$ and $H_i(4) = p_0 + p_1 + p_2 + p_3 + p_4/2$.

Here we used the fact that, by construction of $F_{i,k}$, bidder i perceives that all other bidders $k \in N \setminus \{i, j\}$ bid 0 with probability 1.

Now, by Lemma 3.1 the first jump point $\alpha_i(0)$ of β_i must necessarily satisfy $u_i(0, \boldsymbol{\beta}_{-i}; \alpha_i(0)) \geq u_i(1, \boldsymbol{\beta}_{-i}; \alpha_i(0)) - \varepsilon$ (because the interval $(0, \alpha_i(0))$ is non-empty by the non-overbidding assumption). We can rewrite this as $H_i(0) \cdot (\alpha_i(0) - 0) \geq H_i(1) \cdot (\alpha_i(0) - 1) - \varepsilon$, which yields

$$\alpha_i(0) \leq \frac{H_i(1) + \varepsilon}{H_i(1) - H_i(0)} = 1 + \frac{H_i(0) + \varepsilon}{H_i(1) - H_i(0)} = 1 + \frac{p_0/n + \varepsilon}{p_0(n-1)/n + p_1/2} \tag{10}$$

where the fraction is interpreted as $+\infty$ when $p_0 + p_1 = 0$. Similarly, by Lemma 3.1, the fourth jump point must satisfy $u_i(4, \boldsymbol{\beta}_{-i}; \alpha_i(3)) \geq u_i(3, \boldsymbol{\beta}_{-i}; \alpha_i(3)) - \varepsilon$, unless $\alpha_i(3) = 5$. Rewriting this as $H_i(4) \cdot (\alpha_i(3) - 4) \geq H_i(3) \cdot (\alpha_i(3) - 3) - \varepsilon$, we obtain that $\alpha_i(3) = 5$ or

$$\alpha_i(3) \geq \frac{4H_i(4) - 3H_i(3) - \varepsilon}{H_i(4) - H_i(3)} = 4 + \frac{H_i(3) - \varepsilon}{H_i(4) - H_i(3)} = 4 + \frac{p_0 + p_1 + p_2 + p_3/2 - \varepsilon}{p_3/2 + p_4/2}. \tag{11}$$

Again, by Lemma 3.1, the third jump point must satisfy $u_i(3, \boldsymbol{\beta}_{-i}; \alpha_i(2)) \geq u_i(2, \boldsymbol{\beta}_{-i}; \alpha_i(2)) - \varepsilon$, unless $\alpha_i(2) = \alpha_i(3)$, and it must satisfy $u_i(2, \boldsymbol{\beta}_{-i}; \alpha_i(2)) \geq u_i(3, \boldsymbol{\beta}_{-i}; \alpha_i(2)) - \varepsilon$, unless $\alpha_i(2) = \alpha_i(1)$. Thus it follows that

$$\alpha_i(2) = \mathrm{T}_{[\alpha_i(1), \alpha_i(3)]}\left(\frac{3H_i(3) - 2H_i(2) \pm \varepsilon}{H_i(3) - H_i(2)}\right) = \mathrm{T}_{[\alpha_i(1), \alpha_i(3)]}\left(3 + \frac{2p_0 + 2p_1 + p_2 \pm 2\varepsilon}{p_2 + p_3}\right). \tag{12}$$

Finally, by Lemma 3.1, the second jump point must satisfy $u_i(2, \boldsymbol{\beta}_{-i}; \alpha_i(1)) \geq u_i(1, \boldsymbol{\beta}_{-i}; \alpha_i(1)) - \varepsilon$, unless $\alpha_i(1) = \alpha_i(2)$, and it must satisfy $u_i(1, \boldsymbol{\beta}_{-i}; \alpha_i(1)) \geq u_i(2, \boldsymbol{\beta}_{-i}; \alpha_i(1)) - \varepsilon$, unless $\alpha_i(1) = \alpha_i(0)$. As a result, it must be that

$$\alpha_i(1) = \mathrm{T}_{[\alpha_i(0), \alpha_i(2)]}\left(\frac{2H_i(2) - H_i(1) \pm \varepsilon}{H_i(2) - H_i(1)}\right) = \mathrm{T}_{[\alpha_i(0), \alpha_i(2)]}\left(2 + \frac{2p_0 + p_1 \pm 2\varepsilon}{p_1 + p_2}\right). \tag{13}$$

We are now ready to prove Claim 1. Consider a base gadget with input bidder j, output bidder i and parameters $(\gamma_\ell, \gamma_r, \ell, r)$, such that $\gamma_\ell, \gamma_r \geq 1/20$, $\gamma_\ell + \gamma_r < 1$ and $\ell < r$. Let p_b denote the probability that bidder j bids b, as perceived by bidder i.

Assume first that bidder j is almost-valid. Then, by the construction of $F_{i,j}$, we obtain that $p_0 = p_3 = p_4 = 0$, $p_1 \in [\gamma_\ell, 1 - \gamma_r]$ and $p_2 = 1 - p_1$. Using (10) we have that $\alpha_i(0) \leq 1 + \frac{\varepsilon}{p_1/2} \leq 1 + \frac{2\varepsilon}{\gamma_\ell} \leq 1 + 1/2$ since $\gamma_\ell \geq 4\varepsilon$. Using (11) we obtain that $\alpha_i(3) = 5$, since $p_3 = p_4 = 0$ and $1 - \varepsilon > 0$. (12) yields that $\alpha_i(2) \geq 3 + \frac{1 + p_1 - 2\varepsilon}{1 - p_1} \geq 3 + 1/2$, since $\varepsilon \leq 1/4$. Thus, in order to show that bidder i is almost-valid, it remains to prove that $\alpha_i(1) \in [2, 3]$. Using (13) we can write

$$\alpha_i(1) = \mathrm{T}_{[\alpha_i(0), \alpha_i(2)]}\left(2 + \frac{2p_0 + p_1 \pm 2\varepsilon}{p_1 + p_2}\right) = \mathrm{T}_{[\alpha_i(0), \alpha_i(2)]}\left(2 + p_1 \pm 2\varepsilon\right) = 2 + p_1 \pm 2\varepsilon$$

where we used the fact that $p_1 + 2\varepsilon \leq 1$, since $p_1 \leq 1 - \gamma_r$ and $\gamma_r \geq 2\varepsilon$. Note that this also yields that $\alpha_i(1) \leq 3$, while the bound $\alpha_i(1) \geq 2$ holds because $p_1 \geq \gamma_\ell$ and $\gamma_\ell \geq 2\varepsilon$ (or simply because of the no-overbidding assumption). As a result, bidder i is almost-valid.

Now consider the case where, in addition, $\gamma_\ell, \gamma_r \geq 1/3$. We can write

$$p_1 = \gamma_\ell + (1 - \gamma_\ell - \gamma_r)\frac{\mathrm{T}_{[2+\ell, 2+r]}(\alpha_j(1)) - (2 + \ell)}{r - \ell}.$$

In particular, it holds that $p_1 \in [1/3, 2/3]$. Since, as shown above, $\alpha_i(1) = 2 + p_1 \pm 2\varepsilon$, we immediately obtain that $\alpha_i(1) \in [2 + 1/3 - 2\varepsilon, 2 + 2/3 + 2\varepsilon]$, i.e., bidder i is valid. Furthermore, we can write

$$\mathbf{v}[i] = \mathrm{T}_{[0,1]}(3(\alpha_i(1) - 2 - 1/3)) = 3p_1 - 1 \pm 6\varepsilon = (3\gamma_\ell - 1) + 3(1 - \gamma_\ell - \gamma_r)\frac{\mathrm{T}_{[2+\ell, 2+r]}(\alpha_j(1)) - (2 + \ell)}{r - \ell} \pm 6\varepsilon$$

which proves the claim. □

Projection Gadget. The projection gadget with input bidder j and output bidder i, uses two additional auxiliary-bidders k and k', and consists of three uses of the standard base gadget. Concretely, the first standard base gadget has input j and output k, the second such gadget has input k and output k', and the third has input k' and output i. See Figure 3a for an illustration. As stated in the claim below, the projection gadget has the notable property that the output bidder i is *always* valid. This gadget will be used to ultimately ensure that all the gate-bidders are valid.

Claim 2. *The projection gadget with input bidder j and output bidder i ensures that:*

- *the output bidder i is valid, and*
- *if the input bidder j is valid, then $\mathbf{v}[i] = \mathbf{v}[j] \pm 18\varepsilon$.*

Proof. The second point follows immediately from Claim 1 applied to the standard base gate. Thus, it remains to show that the output bidder i is always valid. Consider the first standard base gadget, which has input bidder j and output bidder k. Let p_b denote the probability that bidder j bids b, as perceived by bidder k. Since the density function of $F_{k,j}$ has a block of volume $1/3$ lying in $[1 + 1/2, 1 + 3/4]$, and since we do not allow overbidding, it follows that $p_0 + p_1 \geq 1/3$. Using (10) this implies that

$$\alpha_k(0) \leq 1 + \frac{p_0/n + \varepsilon}{p_0(n-1)/n + p_1/2} \leq 1 + 6/n + 6\varepsilon \leq 1 + 1/2$$

since wlog $n \geq 24$ and $\varepsilon \leq 1/24$. Next, using (11) we immediately get that $\alpha_k(3) \geq 4$ since $\varepsilon < 1/3$ (or just by using the no-overbidding assumption). Then, (12) implies that

$$\alpha_k(2) = \mathrm{T}_{[\alpha_k(1),\alpha_k(3)]}\left(3 + \frac{2p_0 + 2p_1 + p_2 \pm 2\varepsilon}{p_2 + p_3}\right) \geq 4 - 2\varepsilon \geq 3 + 1/2$$

where we used $p_0 + p_1 \geq 1/3$, $p_2 + p_3 \leq 2/3$, and $\varepsilon \leq 1/4$. Finally, note that $\alpha_k(1) \geq 2$ by the no-overbidding assumption.

Next, consider the second standard base gadget, which has input bidder k and output bidder k'. Let p_b denote the probability that bidder k bids b, as perceived by bidder k'. From the construction of the density function of $F_{k',k}$ and the bounds obtained on the jump points of k in the first step, it follows that $p_0 = p_3 = p_4 = 0$ and $p_1 \geq 1/3$. Using Equations (10) to (12) similarly to above, we obtain that $\alpha_{k'}(0) \leq 1 + 1/2$, $\alpha_{k'}(2) \geq 3 + 1/2$ and $\alpha_{k'}(3) = 5$. As before, we have that $\alpha_{k'}(1) \geq 2$ by the no-overbidding assumption, and using (13) we also obtain that

$$\alpha_{k'}(1) = \mathrm{T}_{[\alpha_{k'}(0),\alpha_{k'}(2)]}\left(2 + \frac{2p_0 + p_1 \pm 2\varepsilon}{p_1 + p_2}\right) \leq 2 + 1 + 2\varepsilon \leq 3 + 1/4$$

since $\varepsilon \leq 1/8$.

Finally, consider the third and last standard base gadget, which has input bidder k' and output bidder i. Let p_b denote the probability that bidder k' bids b, as perceived by bidder i. From the construction of the density function of $F_{i,k'}$ and the bounds obtained on the jump points of k' in the previous step, it follows that $p_0 = p_3 = p_4 = 0$, $p_1 \geq 1/3$ and $p_2 \geq 1/3$. Again using Equations (10) to (12) as in the previous step, we obtain that $\alpha_i(0) \leq 1 + 1/2$, $\alpha_i(2) \geq 3 + 1/2$ and $\alpha_i(3) = 5$. Using (13) we have that

$$\alpha_i(1) = \mathrm{T}_{[\alpha_i(0),\alpha_i(2)]}\left(2 + \frac{2p_0 + p_1 \pm 2\varepsilon}{p_1 + p_2}\right) = 2 + p_1 \pm 2\varepsilon \in [2 + 1/3 - 2\varepsilon, 2 + 2/3 + 2\varepsilon]$$

and thus bidder i is indeed valid. □

$G_{\times 2}$ Gadget. The $G_{\times 2}$ gadget with input bidder j and output bidder i, uses an additional auxiliary-bidder k, and consists of one use of the base gadget and one use of the projection gadget. In more detail, the base gadget has input j, output k and parameters $(\gamma_\ell, \gamma_r, \ell, r) = (1/3, 1/3, 1/3, 1/2)$, while the projection gate has input k and output i. See Figure 3b for an illustration.

Claim 3. *The $G_{\times 2}$ gadget with input bidder j and output bidder i ensures that:*

- *the output bidder i is valid, and*
- *if the input bidder j is valid, then $\mathbf{v}[i] = \mathrm{T}(2 \cdot \mathbf{v}[j]) \pm 24\varepsilon$.*

Proof. The fact that bidder i is valid follows from our use of the projection gadget and the first bullet point in Claim 2. Now consider the case where bidder j is valid. Since $\gamma_\ell = \gamma_r = 1/3$, by Claim 1 we know that bidder k is also valid and it holds that

$$\mathbf{v}[k] = \frac{\mathrm{T}_{[2+\ell,2+r]}(\alpha_j(1)) - (2+\ell)}{r-\ell} \pm 6\varepsilon = \mathrm{T}_{[0,1]}\left(6\alpha_j(1) - 14\right) \pm 6\varepsilon = \mathrm{T}_{[0,1]}(2 \cdot \mathbf{v}[j]) \pm 6\varepsilon.$$

Since k is valid, we can use the second bullet point in Claim 2, which yields $\mathbf{v}[i] = \mathbf{v}[k] \pm 18\varepsilon = \mathrm{T}_{[0,1]}(2 \cdot \mathbf{v}[j]) \pm 24\varepsilon$. □

G_{1-} Gadget. The G_{1-} gadget with input bidder j and output bidder i uses three additional auxiliary-bidders k_1, k_2, k_3. First, a base gadget is used with input j, output k_1 and parameters $(\gamma_\ell, \gamma_r, \ell, r) = (1/6, 2/3, 1/3, 2/3)$. Next, the density function of F_{k_2,k_1} has a block of volume $2/3$ in $[1+1/2, 1+3/4]$, and a block of volume $1/3$ in $[4,5]$. Then, we use a base gadget with input k_2, output k_3 and parameters $(\gamma_\ell, \gamma_r, \ell, r) = (1/3, 1/3, 2/3, 5/6)$. Finally, we use a projection gadget with input k_3 and output i. See Figure 3c for an illustration.

The crucial idea behind the construction of this gadget is that the third jump point (instead of the second one) is used to encode information in some intermediate step. This allows us to simulate the non-monotone operation $x \mapsto 1 - x$.

Claim 4. *The G_{1-} gadget with input bidder j and output bidder i ensures that:*

- *the output bidder i is valid, and*
- *if the input bidder j is valid, then $\mathbf{v}[i] = 1 - \mathbf{v}[j] \pm 60\varepsilon$.*

Proof. First of all, note that i must be valid, because of the corresponding property of the projection gadget (Claim 2). Now consider the case where j is valid. By Claim 1 it follows that bidder k_1 is almost-valid, in particular $\alpha_{k_1}(3) = 5$ and $\alpha_{k_1}(1) \le 3$. Let p_b denote the probability that bidder j bids b, as perceived by bidder k_1. Since j is valid, we immediately obtain that $p_0 = p_3 = p_4 = 0$. Furthermore, by the construction of $F_{k_1,j}$, it is easy to see that $p_1 = 1/6 + (1 - 1/6 - 2/3)\mathbf{v}[j] = 1/6 + \mathbf{v}[j]/6$. Next, using (12) we can write

$$\begin{aligned}
\alpha_{k_1}(2) = \mathrm{T}_{[\alpha_{k_1}(1),\alpha_{k_1}(3)]}\left(3 + \frac{2p_0 + 2p_1 + p_2 \pm 2\varepsilon}{p_2 + p_3}\right) &= \mathrm{T}_{[\alpha_{k_1}(1),\alpha_{k_1}(3)]}\left(3 + \frac{1 + p_1 \pm 2\varepsilon}{1 - p_1}\right) \\
&= 3 + \frac{7/6 + \mathbf{v}[j]/6}{5/6 - \mathbf{v}[j]/6} \pm 3\varepsilon \\
&= 4 + \frac{2 + 2\mathbf{v}[j]}{5 - \mathbf{v}[j]} \pm 3\varepsilon.
\end{aligned}$$

Now consider bidder k_2. Let p_b denote the probability that bidder k_1 bids b, as perceived by bidder k_2. By construction of F_{k_2,k_1} and since k_1 is almost-valid, it is easy to see that $p_0 = p_4 = 0$, $p_1 = 2/3$ and $p_2 + p_3 = 1/3$. By the same arguments used in the proof of Claim 1 it follows that $\alpha_{k_2}(0) \le 1 + 1/2$. By using (11) we obtain $\alpha_{k_2}(3) \ge 4 + \frac{2/3+1/6-\varepsilon}{1/6} \ge 5$. Next, using (12) we obtain

$$\alpha_{k_2}(2) \ge \mathrm{T}_{[\alpha_{k_2}(1),\alpha_{k_2}(3)]}\left(3 + \frac{2p_0 + 2p_1 + p_2 \pm 2\varepsilon}{p_2 + p_3}\right) \ge \mathrm{T}_{[\alpha_{k_2}(1),5]}\left(3 + \frac{4/3 - 2\varepsilon}{1/3}\right) = 5.$$

Now observe that by construction of F_{k_2,k_1} and the expression obtained earlier for $\alpha_{k_1}(2)$

$$p_2 = \frac{\mathrm{T}_{[4,5]}(\alpha_{k_1}(2)) - 4}{3} = \frac{2 + 2\mathbf{v}[j]}{15 - 3\mathbf{v}[j]} \pm \varepsilon.$$

As a result, it follows that

$$\frac{2p_0+p_1}{p_1+p_2} = \frac{2/3}{2/3+\frac{2+2\mathbf{v}[j]}{15-3\mathbf{v}[j]}\pm\varepsilon} = \frac{2/3}{2/3+\frac{2+2\mathbf{v}[j]}{15-3\mathbf{v}[j]}} \pm 3\varepsilon = 5/6-\mathbf{v}[j]/6\pm 3\varepsilon$$

where we used $\varepsilon \le 2/15$. Finally, using (13) we obtain

$$\begin{aligned}\alpha_{k_2}(1) = \mathrm{T}_{[\alpha_{k_2}(0),\alpha_{k_2}(2)]}\left(2+\frac{2p_0+p_1\pm 2\varepsilon}{p_1+p_2}\right) &= \mathrm{T}_{[\alpha_{k_2}(0),\alpha_{k_2}(2)]}\left(2+5/6-\mathbf{v}[j]/6\pm 3\varepsilon\pm\frac{2\varepsilon}{p_1+p_2}\right)\\ &= 2+5/6-\mathbf{v}[j]/6\pm 6\varepsilon.\end{aligned}$$

Note in particular that bidder k_2 is almost-valid, since $\varepsilon \le 1/36$.

Since bidder k_2 is almost-valid, and we use a base gadget with $\gamma_\ell = \gamma_r = 1/3$ with input k_2 and output k_3, it follows by Claim 1 that bidder k_3 is valid and

$$\mathbf{v}[k_3] = \frac{\mathrm{T}_{[2+\ell,2+r]}(\alpha_{k_2}(1)) - (2+4/6)}{1/6} \pm 6\varepsilon = 1-\mathbf{v}[j]\pm 42\varepsilon.$$

Finally, the projection gadget with input k_3 and output i ensures that $\mathbf{v}[i] = \mathbf{v}[k_3]\pm 18\varepsilon = 1-\mathbf{v}[j]\pm 60\varepsilon$. □

G_ϕ Gadget. The G_ϕ gadget with input bidders j_1 and j_2 and output bidder i is a binary gadget with additional auxiliary-bidders k_1, k_2, k_3. First of all, for all $t \in N \setminus \{j_1, j_2, k_1\}$, we set $F_{k_1,t}$ to have density function with a single block of volume 1 in $[0,1]$. We set *both* F_{k_1,j_1} and F_{k_1,j_2} to be distributions as in our construction of the base gadget with parameters $(\gamma_\ell, \gamma_r, \ell, r) = (1/20, 8/20, 1/3, 2/3)$. The density function of F_{k_2,k_1} has a block of volume $1/2$ in $[1+1/2, 1+3/4]$, and a block of volume $1/2$ in $[3+1/2, 5]$. Next, we use a base gadget with input k_2, output k_3 and parameters $(\gamma_\ell, \gamma_r, \ell, r) = (1/3, 1/3(1+1/4), 104/200, 779/800)$. Finally, we use a G_{1-} gadget with input k_3 and output i. See Figure 4 for an illustration. We have the following claim.

Claim 5. *The G_ϕ gadget with input bidders j_1, j_2 and output bidder i ensures that:*

- *the output bidder i is valid, and*
- *if the input bidders j_1 and j_2 are valid, then*

$$\mathbf{v}[i] = \phi(\mathbf{v}[j_1], \mathbf{v}[j_2]) \pm 86\varepsilon = \frac{1}{4}(\mathbf{v}[j_1]+1)(\mathbf{v}[j_2]+1) \pm 86\varepsilon.$$

Proof. Bidder i is guaranteed to be valid, because it is the output bidder of the G_{1-} gadget (Claim 4). Now assume that j_1 and j_2 are valid. Let p_b denote the probability that bidder j_1 bids b, as perceived by bidder k_1. Similarly, let q_b denote the probability that bidder j_2 bids b, as perceived by bidder k_1. By construction of F_{k_1,j_1} and F_{k_1,j_2}, and because j_1 and j_2 are valid, we know that $p_0 = p_3 = p_4 = q_0 = q_3 = q_4 = 0$, $p_1, q_1 \ge 1/20$ and $p_2, q_2 \ge 8/20$. Recall that $H_{k_1}(b)$ is used to denote the probability that bidder k_1 wins if she bids b (from k_1's perspective). Thus we immediately obtain that $H_{k_1}(0) = 0$, $H_{k_1}(1) = p_1q_1/3$, $H_{k_1}(2) = p_1q_1 + p_2q_1/2 + p_1q_2/2 + p_2q_2/3 = 1/3 + (p_1+q_1)/6 + p_1q_1/3$ and $H_{k_1}(3) = H_{k_1}(4) = 1$. With this in hand, we now obtain (just as we did for Equations (10) to (13)):

$$\alpha_{k_1}(0) \le 1 + \frac{H_{k_1}(0)+\varepsilon}{H_{k_1}(1)-H_{k_1}(0)} = 1+\frac{\varepsilon}{p_1q_1/3} \le 1+1200\varepsilon \le 1+1/2$$

since $\varepsilon \le 1/2400$. Similarly, since $H_{k_1}(4) - H_{k_1}(3) = 0$ and $H_{k_1}(3) = 1 > \varepsilon$, we have that $\alpha_{k_1}(3) = 5$. We also have

$$\alpha_{k_1}(1) \le 2 + \frac{H_{k_1}(1)+\varepsilon}{H_{k_1}(2)-H_{k_1}(1)} \le 2+\frac{p_1q_1/3+\varepsilon}{1/3+(p_1+q_1)/6} \le 3$$

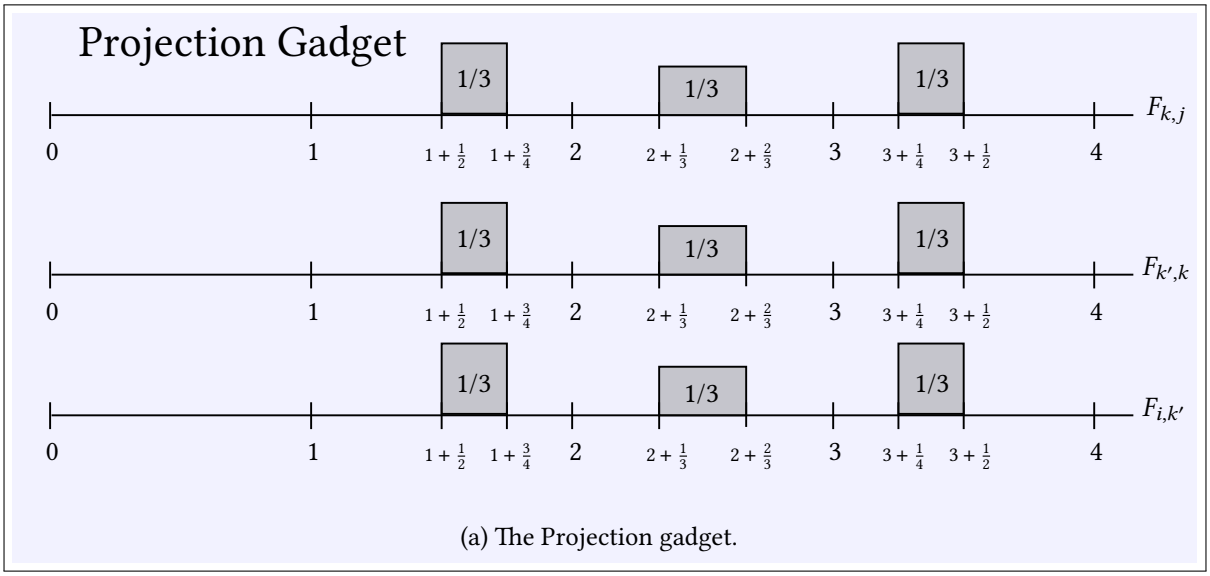

(a) The Projection gadget.

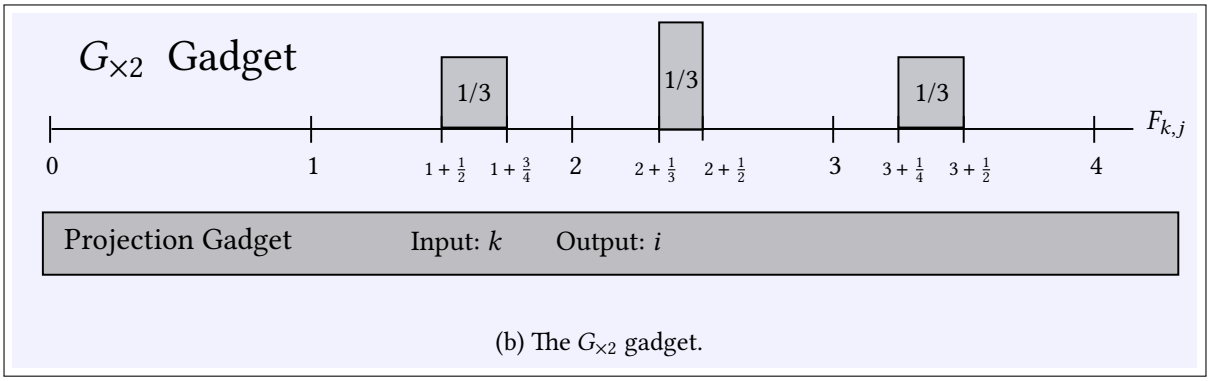

(b) The $G_{\times 2}$ gadget.

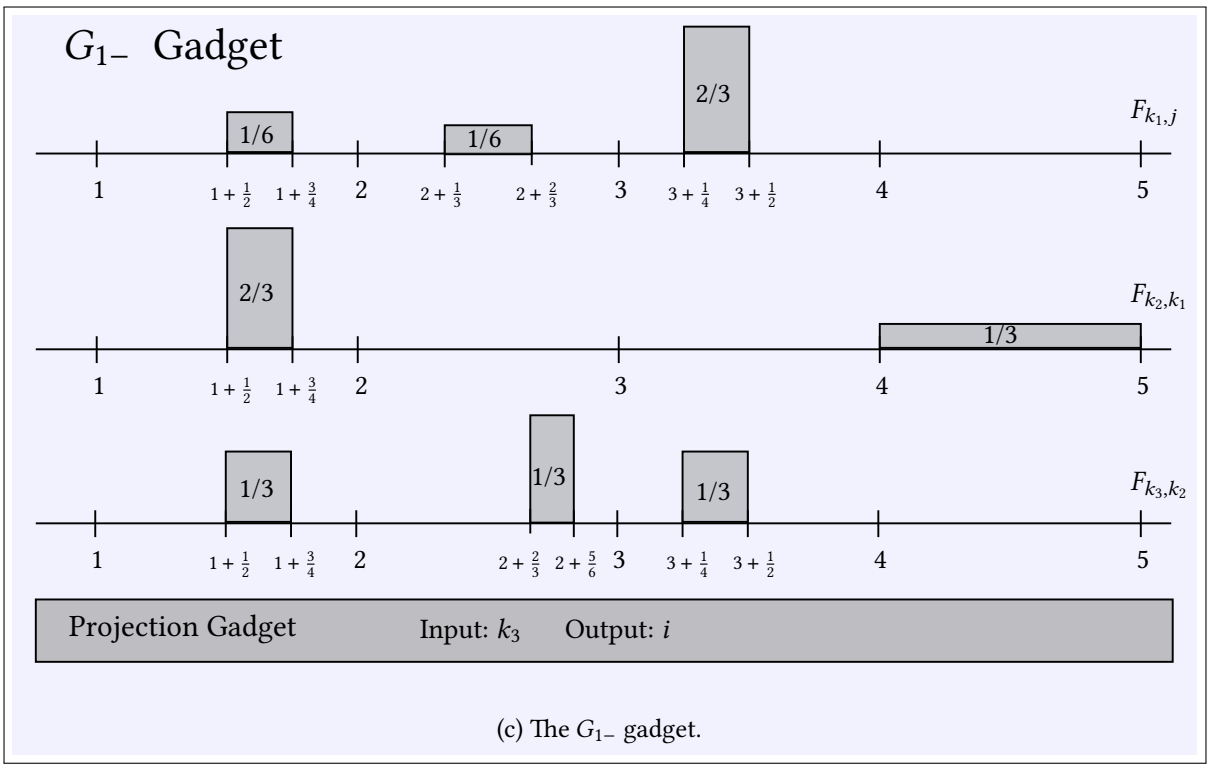

(c) The G_{1-} gadget.

Figure 3: The Projection, $G_{\times 2}$ and G_{1-} gadgets. The probability density functions of the corresponding subjective priors are shown.

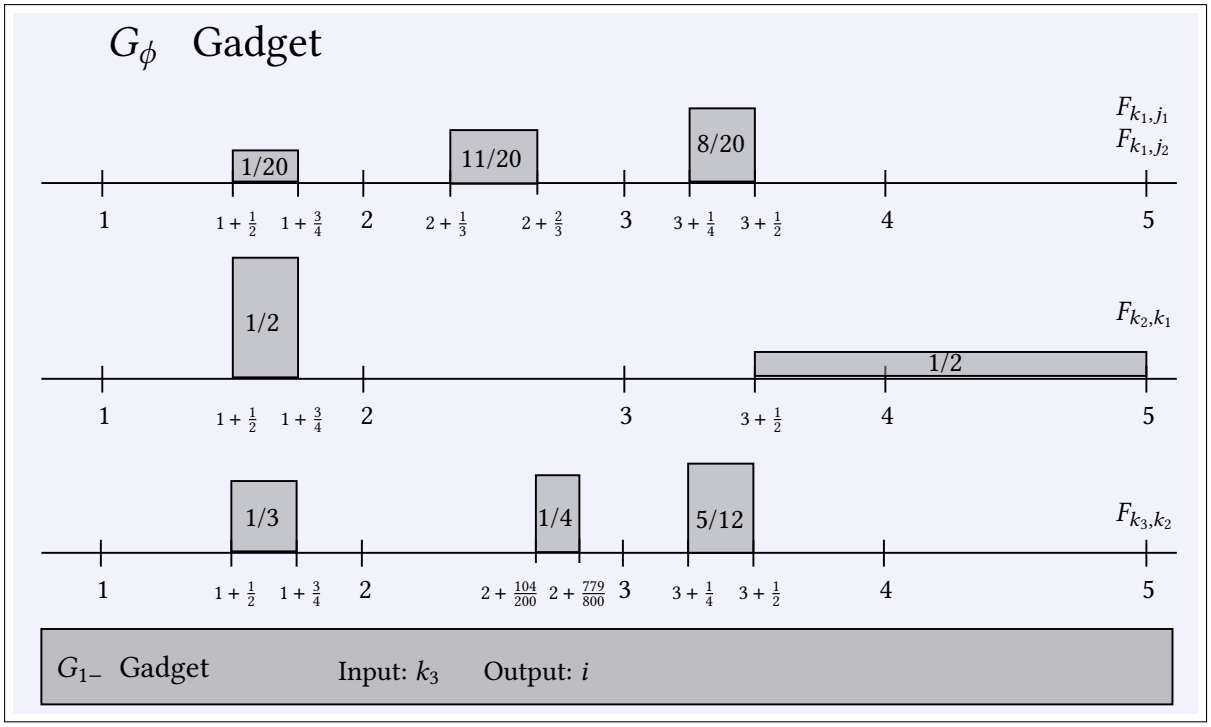

Figure 4: The G_ϕ gadget. The probability density functions of the corresponding subjective priors are shown.

where we used the bounds we have on these probabilities and $\varepsilon \le 1/4$. Finally, we have

$$\begin{aligned}\alpha_{k_1}(2) = \mathrm{T}_{[\alpha_{k_1}(1),\alpha_{k_1}(3)]}\left(3 + \frac{H_{k_1}(2) \pm \varepsilon}{H_{k_1}(3) - H_{k_1}(2)}\right) &= \mathrm{T}_{[3,5]}\left(3 + \frac{1/3 + (p_1+q_1)/6 + p_1q_1/3 \pm \varepsilon}{1 - (1/3 + (p_1+q_1)/6 + p_1q_1/3)}\right) \\ &= 3 + \frac{1 + (p_1+q_1)/2 + p_1q_1}{2 - (p_1+q_1)/2 - p_1q_1} \pm 3\varepsilon \\ &= 3 + \frac{1}{2} + \frac{3}{2}\frac{(p_1+q_1)/2 + p_1q_1}{2 - (p_1+q_1)/2 - p_1q_1} \pm 3\varepsilon\end{aligned}$$

where we used the fact that $\frac{(p_1+q_1)/2+p_1q_1}{2-(p_1+q_1)/2-p_1q_1} \le 1$, since $p_1, q_1 \le 12/20$. As $p_1, q_1 \ge 1/20$ and $\varepsilon \le 1/60$, we also have that $\alpha_{k_1}(2) \ge 3 + 1/2$. In particular, k_1 is almost-valid. Note that since j_1 and j_2 are valid, we have $p_1 = 1/20 + 11\mathbf{v}[j_1]/20$ and $q_1 = 1/20 + 11\mathbf{v}[j_2]/20$.

Next, we consider bidder k_2. Let p'_b denote the probability that bidder k_1 bids b, as perceived by bidder k_2. By the previous paragraph, we have $p'_0 = 0$, $p'_1 = 1/2$, $p'_4 = 0$ and

$$p'_2 = \frac{1}{3}\frac{3}{2}\frac{(p_1+q_1)/2 + p_1q_1}{2 - (p_1+q_1)/2 - p_1q_1} \pm 3\varepsilon = \frac{1}{2}\frac{(p_1+q_1)/2 + p_1q_1}{2 - (p_1+q_1)/2 - p_1q_1} \pm 3\varepsilon$$

where we used the fact that the height of the block of volume of F_{k_2,k_1} in $[3 + 1/2, 5]$ is $1/3$. Since the density function of F_{k_2,k_1} has a block of volume $1/2$ in $[1 + 1/2, 1 + 3/4]$, as before we obtain that $\alpha_{k_2}(0) \le 1 + 1/2$. Using (11) and (12), we also have

$$\alpha_{k_2}(3) \ge 4 + \frac{p'_0 + p'_1 + p'_2 + p'_3/2 - \varepsilon}{p'_3/2 + p'_4/2} \ge 5$$

as well as

$$\alpha_{k_2}(2) \ge \mathrm{T}_{[\alpha_{k_2}(1),\alpha_{k_2}(3)]}\left(3 + \frac{2p'_0 + 2p'_1 + p'_2 \pm 2\varepsilon}{p'_2 + p'_3}\right) \ge 3 + 1/2.$$

Finally, (13) yields

$$\begin{aligned}\alpha_{k_2}(1) &= \mathrm{T}_{[\alpha_{k_1}(0),\alpha_{k_1}(2)]}\left(2+\frac{2p_0'+p_1'\pm 2\varepsilon}{p_1'+p_2'}\right)\\ &= \mathrm{T}_{[\alpha_{k_1}(0),\alpha_{k_1}(2)]}\left(2+\frac{1/2}{1/2+\frac{1}{2}\frac{(p_1+q_1)/2+p_1q_1}{2-(p_1+q_1)/2-p_1q_1}\pm 3\varepsilon}\right)\pm 4\varepsilon\\ &= 2+\frac{2-(p_1+q_1)/2-p_1q_1}{2}\pm 10\varepsilon.\end{aligned}$$

Substituting in $p_1 = 1/20 + 11\mathbf{v}[j_1]/20$ and $q_1 = 1/20 + 11\mathbf{v}[j_2]/20$, we compute

$$\begin{aligned}\alpha_{k_2}(1) &= 2+1+1/8-\frac{1}{2}(11/20+11\mathbf{v}[j_1]/20)(11/20+11\mathbf{v}[j_2]/20)\pm 10\varepsilon\\ &= 2+9/8-\frac{121}{200}\phi(\mathbf{v}[j_1],\mathbf{v}[j_2])\pm 10\varepsilon.\end{aligned}$$

Note that we have $\alpha_{k_2}(1) \in [2+104/200, 2+779/800] \pm 10\varepsilon$. In particular, bidder k_2 is almost-valid.

Since bidder k_3 is the output of a base gadget with input k_2 and parameters $(\gamma_\ell, \gamma_r, \ell, r) = (1/3, 1/3(1+1/4), 104/200, 779/800)$, it follows by Claim 1 that k_3 is valid and

$$\begin{aligned}\mathbf{v}[k_3] &= 3(1-\gamma_\ell-\gamma_r)\frac{\mathrm{T}_{[2+\ell,2+r]}(\alpha_{k_2}(1))-(2+\ell)}{r-\ell}\pm 6\varepsilon\\ &= \frac{200}{121}(\alpha_{k_2}(1)-(2+\ell))\pm 6\varepsilon\\ &= \frac{200}{121}\left(2+9/8-\frac{121}{200}\phi(\mathbf{v}[j_1],\mathbf{v}[j_2])-(2+104/200)\right)\pm 26\varepsilon\\ &= 1-\phi(\mathbf{v}[j_1],\mathbf{v}[j_2])\pm 26\varepsilon.\end{aligned}$$

Finally, it is easy to see that the G_{1-} gadget with input k_3 and output i ensures the desired value for bidder i (Claim 4). □

Finishing the proof. Using the gadgets we have described above we can now enforce the constraints of the GCIRCUIT instance. Indeed, for each gate $g_i = (G, j, k)$ where $G \in \mathcal{G} = \{G_{\times 2}, G_{1-}, G_\phi\}$, it suffices to use the gadget corresponding to the gate-type G, with output bidder i and input bidder j (as well as k, in the case $G = G_\phi$). Since the distributions are subjective, we can re-use a bidder j as an input to multiple different gadgets, without any interference. By Claims 3 to 5 it immediately follows that the gate-bidders $1, 2, \ldots, \mathfrak{m}$ must all be valid, since each of them is the output of some gadget. But this means that for any gate $g_i = (G, j, k)$, the input bidder j (and k, if applicable) will be valid, because she is also a gate-bidder. As a result, again by Claims 3 to 5, it follows that the gadgets will correctly enforce their constraints on all values $\mathbf{v}[i]$.

To obtain a solution, it suffices to set $\mathbf{v}[g_i] := \mathbf{v}[i]$ for all $i \in [\mathfrak{m}]$. For the case $\varepsilon = 0$, note that since every gate-bidder i is valid, we have that $\alpha_i(1) \in [2+1/3, 2+2/3]$ and as a result $\mathbf{v}[i] = \mathrm{T}_{[0,1]}(3(\alpha_i(1)-2-1/3)) = 3(\alpha_i(1)-2-1/3)$, which indeed yields an SL-reduction [Etessami and Yannakakis, 2010]. By scaling back to the original value space $[0,1]$, the proof yields that for all $\varepsilon \in [0, 1/10^5]$, from any ε-BNE of the auction we can extract an 500ε-satisfying assignment for the generalized circuit. As discussed at the beginning of the section, this yields both PPAD- and FIXP-hardness.

6 An Efficient Algorithm for a Constant Number of Bidders and Bids

In this section, we design an algorithm which computes an ε-Bayes-Nash equilibrium of the FPA when (a) the number of bidders n is constant, (b) the size of the bidding space $|B|$ is constant, and (c) the value

distributions $F_{i,j}$ of the bidders are *piecewise polynomial*.

To be more precise, our input comprises of:

- a set of bids[9] $B = \{b_0, b_1, \ldots, b_{|B|-1}\} \subset [0,1]$
- a partition[10] of $[0,1]$ into K intervals $[x_{\ell-1}, x_\ell]$, $\ell = \{1, 2, \ldots, K\}$, with rational endpoints
- for each distribution $F_{i,j}$ and each subinterval $[x_{\ell-1}, x_\ell]$, a vector of rationals $(a_0^{i,j,\ell}, a_1^{i,j,\ell}, \ldots, a_d^{i,j,\ell})$.

Then, (the cumulative distribution function of) $F_{i,j}$ is defined as

$$F_{i,j}(z) = F_{i,j}^{\ell}(z), \qquad \text{for } \; z \in [x_{\ell-1}, x_\ell],$$

where

$$F_{i,j}^{\ell}(z) = \sum_{\kappa=0}^{d} a_\kappa^{i,j,\ell} z^\kappa \tag{14}$$

is the polynomial representation of $F_{i,j}$ in the ℓ-th interval. Of course, the input should respect the conditions

$$F_{i,j}^1(0) \geq 0, \quad F_{i,j}^K(1) = 1, \quad F_{i,j}^{\ell}(x_\ell) = F_{i,j}^{\ell+1}(x_\ell) \;\text{ for }\; \ell = 1, 2, \ldots, K-1,$$

and that each $F_{i,j}^{\ell}$ is nondecreasing on $[x_{\ell-1}, x_\ell]$.

Finally, when we say that n and $|B|$ are fixed, we mean that they are constant functions of the other parameters of the input.

We have the following theorem.

Theorem 6.1. *For a fixed number of bidders, a fixed bidding space, and piecewise polynomial value distributions, an ε-BNE of the first-price auction can be computed in polynomial time, even for subjective priors and even when ε is inversely-exponential in the input size.*

The remainder of the section is devoted to developing the algorithm that will prove Theorem 6.1.

At a high level, the algorithm will perform the following four steps:

1. It "guesses", for each bidder, an assignment of the jump points of her best-response strategy to the K sub-intervals $[x_{\ell-1}, x_\ell]$ above; intervals may be allocated zero or multiple jump points. Since the number of bidders and the size of the bidding space are constant, there is a total constant number of jump points for all bidders. Therefore, this "guessing" step is an enumeration of all such possible assignments; the subsequent steps of the algorithm are run for any such assignment.

2. It "guesses" a set of *effective* jump points and bids. This is a technical corner case, to eliminate degenerate cases in which multiple jump points coincide. Again, this can be done via enumeration given that the number of jump points is constant.

3. It formulates the problem of finding the *exact positions* of the effective jump points (within the intervals corresponding to the guessed allocation above) as a system of polynomial inequalities of polynomially-large degree. A δ-approximate solution to this system can be found using standard methods, in time polynomial in $\log(1/\delta)$ and the input parameters.

4. It "projects" the approximate solution to the "equilibrium space", as defined by the constraints of the aforementioned system, ensuring that the resulting object is indeed an ε-BNE, for some ε that can be made as small as needed, by making δ as small as needed.

Below we describe these steps in more detail.

[9] Recall that here $|B|$ is *fixed*, i.e., not part of the input.

[10] Our assumption here of a *common* interval partition for the piecewise polynomial representation of all subjective priors $F_{i,j}$ is for the sake of simplicity, and it is not critical for the positive results of this section. In particular, it is not difficult to see that our model can handle different partitions $[x_{\ell-1}^{i,j}, x_\ell^{i,j}]$ with just a polynomial blow-up in the size of the representation; essentially one needs to take the interval partition induced by all points $\{x_\ell^{i,j}\}$.

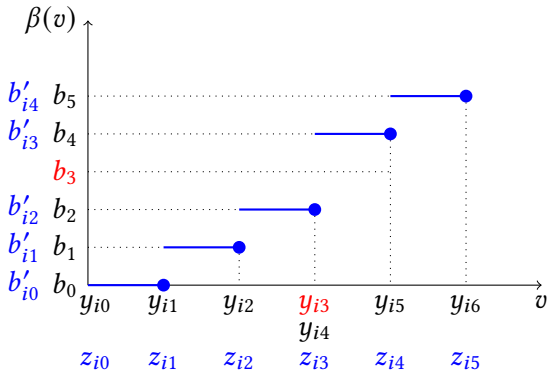

Figure 5: An illustration of the selection of effective jump points (b') and effective bids (z'), for $|B| = 6$. In the figure, jump points 3 and 4 coincide, and therefore among those, only jump point 4 will be in the sequence used in the next step. Also, bid b_3 is never used in the best-response function, as the strategy jumps directly from b_2 to b_4, and therefore b_3 will be excluded from the set of effective bids. In the end, the effective jump points would be 1, 2, 4 and 5 and the effective bids will be b_0, b_1, b_2, b_4 and b_5.

Step 1: Guessing an allocation of jump points to intervals

Recall the definition of the jump points $\alpha_i(b)$ from Section 2, which represent the equilibrium strategy of bidder i. Intuitively, $\alpha_i(b)$ is the largest value for which bidder i would bid b or lower. Since $|B|$ is constant, there is a constant number of such jump points for each bidder, and since n is also constant, there is a constant number of jump points overall. The algorithm enumerates over all the possible ways of assigning the $n \cdot (|B| - 1)$ jump points to the intervals $[x_{\ell-1}, x_\ell]$, for $\ell = 1, \ldots, K$; this can be done in time $O(K^{n|B|})$. Then, for any possible such allocation, it moves to the next step. We introduce variables $y_{i,j}$, $j = 1, 2, \ldots, |B| - 1$ for the positions of the jump points of the strategy of bidder i in $[0, 1]$, and we set $y_{i,0} = 0$, $y_{i,|B|} = 1$.

Step 2: Guessing a set of effective jump points and bids

We "guess" possible "collisions" of sequential jump points, where a collision happens when the positions of two or more jump points coincide. In that case, we would like to only keep a *single* representative from each coinciding jump point; the positions of these representatives are denoted using the variables z_{ij}. We also use the variables $b'_{i,j}$ to denote the corresponding bids, as subscribed by the chosen jump points. We refer to the chosen jump points and bids as *effective* jump points and bids respectively. See Figure 5 for an illustration.

Formally, this corresponds to picking, for each bidder i, an (increasing) subsequence $\mu_i(j) \subseteq \{1, \ldots, |B| - 1\}$, such that

$$z_{i,j} = y_{i,\mu_i(j)} \quad \text{and} \quad \{0 = z_{i,0} < z_{i,1} < \cdots < z_{i,m_i} = 1\} = \{0 = y_{i,0} \le y_{i,1} \le y_{i,2} \le \cdots \le y_{i,|B|} = 1\}.$$

Notice that $m_i \le |B|$. Given the "guessing" in the current step, we let $L_{i,j}, R_{i,j}$ denote the left and right, respectively, endpoints of the sub-interval in which the j-th effective break point of player i lies; i.e., $z_{i,j} \in [L_{i,j}, R_{i,j}]$. For ease of notation, we also use the shortcut $b'_{i,j} = b_{\mu_i(j)}$ for the j-th effective bid of player i.

Again, since $|B|$ is constant, we can enumerate over all possible effective jump point subsequences μ_i in constant time and for each such subsequence, we proceed to the next step.

Step 3: Solving a system of polynomial inequalities

From the previous two steps we have, for each bidder i, an assignment of effective jump points $z_{i,0}, \ldots, z_{i,m_i}$ to intervals $[x_{\ell-1}, x_\ell]$. In particular, $z_{i,j}$ is mapped to $[L_{i,j}, R_{i,j}]$. Below, we express all the

properties that must be satisfied by the effective jump points at an (exact) BNE of the FPA as a system of polynomial inequalities; the system includes inequalities to ensure

- that the positions of the jump points of each bidder i respect the ordering implied by the set of indices, i.e., $z_{i,j-1} < z_{i,j}$ for all $j = 1, \dots, m_i$,
- that the bidding strategies are non-overbidding,
- that the variables $z_{i,j}$ indeed correspond to jump points of best-responses, in terms of the implications to the utility functions.

$$
\begin{aligned}
& z_{i,j-1} < z_{i,j} && \forall i,\ \forall j && (15)\\
& L_{ij} \le z_{i,j} \le R_{ij} && \forall i,\ \forall j && (16)\\
& z_{i,j} \ge b'_{i,j} && \forall i,\ \forall j && (17)\\
& u_i(b'_{i,j}, z_{-i}; z_{i,j}) \ge u_i(b, z_{-i}; z_{i,j}) && \forall i,\ \forall j,\ \forall b < b'_{i,j} && (18)\\
& u_i(b'_{i,j-1}, z_{-i}; z_{i,j}) \ge u_i(b, z_{-i}; z_{i,j}) && \forall i,\ \forall j,\ \forall b > b'_{i,j-1} && (19)
\end{aligned}
$$

Lemma 6.2. *Fix a bidder i and a bid $b \in B$. Then, for every $j = 1, \dots, m_i$, her utility $u_i(b, z_{-i}; z_{i,j})$ can be expressed (in polynomial time) as a polynomial of degree at most dn with respect to the effective jump point variables $\{z_{i',j'}\}_{i' \in N,\ j'=0,\dots,m_i}$.*

Proof. Without loss of generality, similar to what we did in the proof of Lemma 3.2, we will show the lemma from the perspective of bidder n. Fix an index $j = 0, \dots, m_n$ for an effective jump point $z_{n,j} \in [L_{n,j}, R_{n,j}]$, and consider a bid b. Then, importing some notation from our proof of Lemma 3.2, the utility of player n when she has a true value of $v_n = z_{n,j}$ is

$$u_n(b, z_{-n}; z_{n,j}) = H_n(b, z_{-n})(z_{n,j} - b),$$

where $H_n(b, z_{-n})$ is the probability that bidder n wins the item. Due to (5) and (6) (and the fact that n is now constant), it is enough to show that, for any bidder $i \le n-1$, the quantities G_{i,b^-} and $g_{i,b}$, defined in the proof of Lemma 3.2, are polynomials of the jump point variables $z_{i',j'}$. Furthermore, to guarantee a maximum degree of dn, as in the statement of our lemma, it is enough to show that each of these polynomials are of degree at most d: the number of factors in the products appearing as summands in (6) are at most n.

Recall that G_{i,b^-} and $g_{i,b}$ are the probabilities (from the perspective of bidder n) that bidder i bids below b and exactly b, respectively. So, if $b = b'_{i,j'}$ for some index $j' = 0, 1, \dots, m_i - 1$, then $G_{i,b^-} = F_{n,j}(z_{i,j'})$ and $g_{i,b} = F_{n,j}(z_{i,j'+1}) - F_{n,j}(z_{i,j'})$. If, on the other hand, $b'_{i,j'} < b < b'_{i,j'+1}$ for an index j', then $G_{i,b^-} = F_{n,j}(z_{i,j'})$ and $g_{i,b} = 0$. In any case, deploying the representation from (14), quantities G_{i,b^-} and $g_{i,b}$ can indeed be written (in polynomial time with respect to the input of the problem) as polynomials, of degree at most d, of the jump point variables. □

As the following lemma suggests, a solution to System (15)–(19) corresponds to a BNE of the first-price auction. Note that although the existence of a BNE is guaranteed by Theorem 4.1, it might be the case that the equilibrium strategies are *not* consistent with the specific "preliminary" guesses of Steps 1 and 2 that gave rise to the particular instantiation of System (15)–(19) above. However, there has to exist *some guess* for which the system has a solution, and since we are enumerating over all possible choices, we are guaranteed to find it.

Lemma 6.3. *Given the "guessed" allocations of jump points to intervals and the "guessed" effective jump points and bids, a compatible BNE of the FPA exists if and only if System (15)–(19) has a solution.*

Proof. Immediate by the characterization of BNE in Lemma 3.1 (using $\varepsilon = 0$), by setting $\alpha_i(b^-) = z_{i,j}$ in condition (3) and $\alpha_i(b) = z_{i,j}$ in (4). □

Step 4: "Projecting" back to the equilibrium domain

From Step 3 above, we know that by solving System (15)–(19) we can compute an exact BNE of the auction. More precisely, we can compute a δ-approximation to System (15)–(19) in time polynomial in $\log(1/\delta)$, by making use of the following result by Grigor'ev and Vorobjov [1988, Remark, p. 38]:

Theorem 6.4. *For any $\delta \in (0, 1]$, it is possible to find a rational δ-approximation to System (15)–(19) in time polynomial in $\log(1/\delta)$ and the size of the input.*

By δ-approximation here, we mean a point which is geometrically close, with respect to the max norm, to an exact solution of System (15)–(19). This is almost a strong approximation to an exact BNE; if we were to translate this point to a feasible strategy profile, it would yield jump points which are close to the jump points of an exact equilibrium strategy. However, these would only *approximately* satisfy the conditions in System (15)–(19); in particular special care should be taken for the monotonicity and no-overbidding conditions, which we want to be satisfied *exactly*, rather than approximately.

To remedy this, we must first "project" the δ-approximate solution of System (15)–(19) back to the equilibrium domain $\mathcal{D}$ introduced in the proof of Theorem 4.1. Formally, let us denote by z the δ-approximate solution of System (15)–(19), and by z^* the exact solution which it approximates, so that $\|z - z^*\|_\infty \leq \delta$. We compute the projection $\tilde{z}$ from z as

$$\tilde{z}_{i,0} = 0 \quad \text{and} \quad \tilde{z}_{i,j} = \mathrm{T}_{[\max\{b'_{i,j}, \tilde{z}_{i,j-1}\}, 1]}(z_{i,j}).$$

Our next claim is that $\|\tilde{z} - z^*\|_\infty \leq \delta$ as well. This is equivalent to saying that $|\tilde{z}_{i,j} - z^*_{i,j}| \leq \delta$ for every i, j, which can be done by induction on j, the base case $j = 0$ being trivial. For $j > 0$, observe that $\tilde{z}_{i,j}$ must coincide with one of $b'_{i,j}, \tilde{z}_{i,j-1}, 1, z_{i,j}$.

- If $\tilde{z}_{i,j} = z_{i,j}$ then obviously $|\tilde{z}_{i,j} - z^*_{i,j}| \leq \delta$.
- If $\tilde{z}_{i,j} = b'_{i,j}$ then we must have had $z_{i,j} \leq b'_{i,j}$. Since $z^*_{i,j} \geq b'_{i,j}$ and $|z_{i,j} - z^*_{i,j}| \leq \delta$, we must also have $|\tilde{z}_{i,j} - z^*_{i,j}| \leq \delta$.
- Similarly, if $\tilde{z}_{i,j} = 1$ then we must have had $z_{i,j} \geq 1$. Since $z^*_{i,j} \leq 1$ and $|z_{i,j} - z^*_{i,j}| \leq \delta$, we must also have $|\tilde{z}_{i,j} - z^*_{i,j}| \leq \delta$.
- Finally, suppose $\tilde{z}_{i,j} = \tilde{z}_{i,j-1}$. Then we must have had $z_{i,j} \leq \tilde{z}_{i,j-1}$. Using the induction hypothesis, we have that $\tilde{z}_{i,j-1} \leq z^*_{i,j-1} + \delta \leq z^*_{i,j} + \delta$; thus we also have $|\tilde{z}_{i,j} - z^*_{i,j}| \leq \delta$.

Therefore, $\tilde{z}$ constitutes a valid monotone non-decreasing, non-overbidding joint strategy profile, which is within distance δ of the exact BNE z^*. In other words, $\tilde{z}$ is a valid joint strategy profile that is a *strong* δ-approximation to a BNE.

Finally, we need to show that if δ is chosen to be sufficiently small, then any strong δ-approximate BNE is also an ε-BNE of the auction. For this, we use the fact that the family of piecewise polynomial distributions is polynomially continuous (see Appendix A for the formal definition). Indeed, given such a piecewise polynomial distribution, it is easy to see that it must be Lipschitz-continuous, and, crucially, we can in polynomial time compute a corresponding Lipschitz-constant. (Note that any polynomial function $F(z_j) = a_0 + a_1 z_j + \ldots + a_d z_j^d$ is L-Lipschitz-continuous over $[0, 1]$, where $L = |a_1| + 2|a_2| + \ldots d|a_d|$.) With this observation in hand, we can now use Lemma 4.3 to efficiently construct $\delta > 0$ sufficiently small such that for all $i \in N$, $b \in B$ and $v_i \in [0, 1]$

$$\|z - z'\|_\infty \leq \delta \implies |u_i(b, z_{-i}; v_i) - u_i(b, z'_{-i}; v_i)| \leq \varepsilon/2.$$

Since $\tilde{z}$ is a strong δ-approximation, i.e., $\|\tilde{z} - z^*\|_\infty \leq \delta$, it immediately follows that inequalities (18) and (19) of the System are satisfied with additive error at most ε. Using Lemma 3.1, it immediately follows that $\tilde{z}$ is an ε-BNE.

As a result, to summarize, given $\varepsilon > 0$ and the problem instance, we can in polynomial time compute $\delta > 0$ such that running the algorithm described in this section is guaranteed to find an ε-BNE. Since the number of agents and bids is fixed, and the algorithm runs in polynomial time in $\log(1/\delta)$ and the instance size, Theorem 6.1 follows.

7 Conclusion and Future Directions

In this paper, we have classified the complexity of computing a Bayes-Nash equilibrium of the first-price auction with subjective priors, by proving that it is PPAD-complete. As we explained in the introduction, our result contributes fundamentally to our understanding of this celebrated auction format, as well as the literature on total search problems and TFNP. The challenging next step is to move towards the special case of the common priors assumption, where the value distribution of each bidder is common knowledge ($F_{i,j} = F_{i',j}$ for all i, i'). Our PPAD-membership result obviously already extends to this case, as it is a special case of the subjective priors setting. The really intriguing question is to extend our PPAD-hardness result to this case as well. To this end, we state the following open problem, which we consider to be one of the most important problems both in computational game theory and in the literature of total search problems.

Open Problem. *What is the complexity of computing an ε-Bayes-Nash equilibrium of the first-price auction with common priors? Is it* PPAD*-complete? Is it polynomial-time solvable? Or could it be complete for some other (smaller) sub-class of* PPAD*?*

A potential candidate for such a smaller class could be the class PPAD ∩ PLS, which was recently shown by Fearnley et al. [2021] and Babichenko and Rubinstein [2021] to capture the complexity of interesting problems related to optimization via gradient descent, and computing mixed Nash equilibria in congestion games [Rosenthal, 1973] respectively. The class PLS was introduced by Johnson et al. [1988] and captures the computation of local minima of some objective function, and notably characterizes the complexity of finding *pure* Nash equilibria in congestion games [Fabrikant et al., 2004].

A possible "intermediate" step before settling the open problem above for common priors would be to consider priors that are still subjective, but *consistent*, meaning that there exists some common prior distribution P (a "ground truth") over the set of value profiles, such that each bidder's subjective prior distribution given her own value can be directly computed from P. As Myerson [1997, Sec. 2.8] argues, when the subjective priors are consistent, the differences in beliefs can be explained by differences in information, rather than differences in opinion (which are captured even by inconsistent beliefs). In settings like the First-Price Auction, it is meaningful to assume that beliefs are often formed based on observing public signals (e.g., the bidding history of the competitors), possibly with varying degrees of information, and hence subjective priors are quite meaningful.

Another very interesting question is to study the case where both the value distributions and the bidding space are discrete. A special case of this setting was studied by Escamocher et al. [2009], but they only obtained conclusive results for the case of two bidders with bi-valued distributions. We believe that some of our technical contributions (e.g., the computation of the best response functions or the gadgets used in the PPAD-hardness proof) can be adapted to show similar results for that case as well; we leave the details for future work. Finally, it would be very interesting to identify further (in)tractable special cases for our problem; for example, can we obtain a positive result similar to Theorem 6.1 for more general value distributions? Do the hardness results also hold in the setting where the number of bidders is constant, but the bidding space is allowed to be large?

References

Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, July 2001. doi:10.1111/1468-0262.00223.

Yakov Babichenko and Aviad Rubinstein. Settling the complexity of Nash equilibrium in congestion games. In *Proceedings of the 53rd ACM Symposium on Theory of Computing (STOC)*, pages 1426–1437, 2021. doi:10.1145/3406325.3451039.

Pierpaolo Battigalli and Danilo Guaitoli. Conjectural equilibria and rationalizability in a game with incomplete information. In Pierpaolo Battigalli, Aldo Montesano, and Fausto Panunzi, editors, *Decisions, Games and Markets*, pages 97–124. Springer, 1997. doi:10.1007/978-1-4615-6337-2_4.

Pierpaolo Battigalli, Mario Gilli, and M. Cristina Molinari. Learning and convergence to equilibrium in repeated strategic interactions: An introductory survey. *Ricerche Economiche*, 46:335–378, 1992.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017. doi:10.3982/ecta13958.

Kshipra Bhawalkar and Tim 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, January 2011. doi:10.1137/1.9781611973082.55.

Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a Nash equilibrium. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1480–1498, October 2015. doi:10.1109/focs.2015.94.

Gangshu Cai, Peter R. Wurman, and Xiting Gong. A note on discrete bid first-price auction with general value distribution. *International Game Theory Review*, 12(01):75–81, 2010. doi:10.1142/s0219198910002520.

Yang Cai and Christos Papadimitriou. Simultaneous Bayesian auctions and computational complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 895–910, June 2014. doi:10.1145/2600057.2602877.

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory*, 156:343–388, March 2015. doi:10.1016/j.jet.2014.04.010.

Shuchi Chawla and Jason D. Hartline. Auctions with unique equilibria. In *Proceedings of the 14th ACM conference on Electronic Commerce (EC)*, pages 181–196, 2013. doi:10.1145/2492002.2483188.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14:1–14:57, May 2009. doi:10.1145/1516512.1516516.

Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. *Journal of the ACM*, 64(3):1–56, June 2017. doi:10.1145/3064810.

Harrison Cheng. Ranking sealed high-bid and open asymmetric auctions. *Journal of Mathematical Economics*, 42(4-5):471–498, aug 2006. doi:10.1016/j.jmateco.2006.05.008.

Arka Rai Choudhuri, Pavel Hubáček, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. Finding a Nash equilibrium is no easier than breaking Fiat-Shamir. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*, pages 1103–1114, June 2019. doi:10.1145/3313276.3316400.

George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *Journal of the ACM*, 63(2), April 2016. doi:10.1145/2835172.

Michael Suk-Young Chwe. The discrete bid first auction. *Economics Letters*, 31(4):303–306, December 1989. doi:10.1016/0165-1765(89)90019-0.

Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008. doi:10.1016/j.geb.2008.02.015.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009. doi:10.1137/070699652.

Argyrios Deligkas, John Fearnley, Themistoklis Melissourgos, and Paul G. Spirakis. Computing exact solutions of consensus halving and the Borsuk-Ulam theorem. *J. Comput. Syst. Sci.*, 117:75–98, 2021. doi:10.1016/j.jcss.2020.10.006.

Digiday.com. What to know about Google's implementation of first-price ad auctions, 2019. URL https://digiday.com/media/buyers-welcome-auction-standardization-as-google-finally-goes-all-in-on-first-price. Accessed: 2019-09-06.

Guillaume Escamocher, Peter Bro Miltersen, and Rocio Santillan R. Existence and computation of equilibria of first-price auctions with integral valuations and bids. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1227–1228, 2009. URL https://dl.acm.org/doi/10.5555/1558109.1558225.

Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, January 2010. doi:10.1137/080720826.

Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004. doi:10.1145/1007352.1007445.

John Fearnley, Paul W. Goldberg, Alexandros Hollender, and Rahul Savani. The complexity of gradient descent: CLS = PPAD ∩ PLS. In *Proceedings of the 53rd ACM Symposium on Theory of Computing (STOC)*, pages 46–59, 2021. doi:10.1145/3406325.3451052.

Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions without complements are (almost) efficient. *Games and Economic Behavior*, 123:327–341, September 2020. doi:10.1016/j.geb.2015.11.009.

Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. On the complexity of equilibrium computation in first-price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 454–476, 2021. doi:10.1145/3465456.3467627.

Rafael Frongillo and Jens Witkowski. A geometric method to construct minimal peer prediction mechanisms. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 502–508, 2016. doi:10.5555/3015812.3015888.

Drew Fudenberg and David Levine. Limit games and limit equilibria. *Journal of Economic Theory*, 38(2): 261–279, April 1986. doi:10.1016/0022-0531(86)90118-3.

Jugal Garg, Ruta Mehta, and Vijay V. Vazirani. Dichotomies in equilibrium computation and membership of PLC markets in FIXP. *Theory of Computing*, 12(20):1–25, 2016a. doi:10.4086/toc.2016.v012a020.

Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the cryptographic hardness of finding a Nash equilibrium. In *Proceedings of the 36th Annual International Cryptology Conference (CRYPTO)*, pages 579–604, 2016b. doi:10.1007/978-3-662-53008-5_20.

Paul W. Goldberg. A survey of PPAD-completeness for computing Nash equilibria. In Robin Chapman, editor, *Surveys in Combinatorics 2011*, London Mathematical Society Lecture Note Series, pages 51–82. Cambridge University Press, 2011. doi:10.1017/CBO9781139004114.003.

Paul W. Goldberg and Alexandros Hollender. The Hairy Ball problem is PPAD-complete. *Journal of Computer and System Sciences*, 122:34–62, 2021. doi:10.1016/j.jcss.2021.05.004.

Paul W. Goldberg, Alexandros Hollender, Ayumi Igarashi, Pasin Manurangsi, and Warut Suksompong. Consensus halving for sets of items. *Mathematics of Operations Research*, 2022. doi:10.1287/moor.2021.1249.

Georg Gottlob, Gianluigi Greco, and Toni Mancini. Complexity of pure equilibria in bayesian games. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence (IJCAI)*, pages 1294–1299, 2007. doi:10.5555/1625275.1625485.

James H. Griesmer, Richard E. Levitan, and Martin Shubik. Toward a study of bidding processes part IV – games with unknown costs. *Naval Research Logistics*, 14(4):415–433, 1967. doi:10.1002/nav.3800140402.

D. Yu. Grigor'ev and N.N. Vorobjov. Solving systems of polynomial inequalities in subexponential time. *Journal of Symbolic Computation*, 5(1-2):37–64, 1988. doi:10.1016/s0747-7171(88)80005-1.

Frank Hahn. *On the Notion of Equilibrium in Economics: An Inaugural Lecture [By] F.H. Hahn.* Cambridge University Press, 1973.

John C. Harsanyi. Games with incomplete information played by "Bayesian" players, I–III: Part I. The basic model. *Management Science*, 14(3):159–182, 1967. doi:10.1287/mnsc.1040.0270.

Jason D. Hartline. Bayesian mechanism design. *Foundations and Trends in Theoretical Computer Science*, 8(3):143–263, 2012. doi:10.1561/0400000045.

Geoffrey A. Jehle and Philip J. Reny. *Advanced Microeconomic Theory.* Financial Times/Prentice Hall, 2001.

David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988. doi:10.1016/0022-0000(88)90046-3.

Ehud Kalai and Ehud Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993. doi:10.2307/2951492.

Ehud Kalai and Ehud Lehrer. Subjective games and equilibria. *Games and Economic Behavior*, 8(1): 123–163, 1995. doi:10.1016/s0899-8256(05)80019-3.

Vijay Krishna. *Auction Theory.* Academic Press, 2nd edition, 2009.

Bernard Lebrun. Existence of an equilibrium in first price auctions. *Economic Theory*, 7:421–443, 1996. doi:10.1007/BF01213659.

Bernard Lebrun. First price auctions in the asymmetric N bidder case. *International Economic Review*, 40(1):125–142, 1999. doi:10.1111/1468-2354.00008.

Bernard Lebrun. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, 55(1):131–151, April 2006. doi:10.1016/j.geb.2005.01.006.

Alessandro Lizzeri and Nicola Persico. Uniqueness and existence of equilibrium in auctions with a reserve price. *Games and Economic Behavior*, 30(1):83–114, January 2000. doi:10.1006/game.1998.0704.

Brendan Lucier and Allan Borodin. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 537–553, January 2010. doi:10.1137/1.9781611973075.46.

Robert C. Marshall, Michael J. Meurer, Jean-Francois Richard, and Walter Stromquist. Numerical analysis of asymmetric first price auctions. *Games and Economic Behavior*, 7(2):193–220, 1994. doi:10.1006/game.1994.1045.

Eric Maskin and John Riley. Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67 (3):439–454, 2000. doi:10.1111/1467-937X.00138.

Eric Maskin and John Riley. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395–409, 2003. doi:10.1016/S0899-8256(03)00150-7.

Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2):317–324, 1991. doi:10.1016/0304-3975(91)90200-l.

Ruta Mehta. Constant rank two-player games are PPAD-hard. *SIAM Journal on Computing*, 47(5): 1858–1887, 2018. doi:10.1137/15M1032338.

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994. doi:10.2307/2951479.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. doi:10.1287/moor.6.1.58.

Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.

Renato Paes Leme and Éva Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 735–744, October 2010. doi:10.1109/focs.2010.75.

Renato Paes Leme, Balasubramanian Sivan, and Yifeng Teng. Why do competitive markets converge to first-price auctions? In *Proceedings of The World Wide Web Conference (WWW)*, pages 596–605, 2020. doi:10.1145/3366423.3380142.

Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994. doi:10.1016/s0022-0000(05)80063-7.

Michael Plum. Characterization and computation of Nash-equilibria for auctions with incomplete information. *International Journal of Game Theory*, 20(4):393–418, December 1992. doi:10.1007/bf01271133.

Itzhak Rasooly and Carlos Gavidia-Calderon. The importance of being discrete: on the inaccuracy of continuous approximations in auction theory. *arXiv:2006.03016*, 2020. URL http://arxiv.org/abs/2006.03016.

Philip J. Reny and Shmuel Zamir. On the existence of pure strategy monotone equilibria in asymmetric first-price auctions. *Econometrica*, 72(4):1105–1125, July 2004. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2004.00527.x.

John G. Riley and William F. Samuelson. Optimal auctions. *The American Economic Review*, 71(3): 381–392, 1981. URL https://www.jstor.org/stable/1802786.

Alon Rosen, Gil Segev, and Ido Shahaf. Can PPAD hardness be based on standard cryptographic assumptions? *Journal of Cryptology*, 34(1), 2021. doi:10.1007/s00145-020-09369-6.

Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973. doi:10.1007/BF01737559.

Ariel Rubinstein and Asher Wolinsky. Rationalizable conjectural equilibrium: Between Nash and rationalizability. *Games and Economic Behavior*, 6(2):299–311, March 1994. doi:10.1006/game.1994.1016.

Aviad Rubinstein. Inapproximability of Nash equilibrium. *SIAM Journal on Computing*, 47(3):917–959, 2018. doi:10.1137/15m1039274.

Vijay V. Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM*, 58(3):1–25, May 2011. doi:10.1145/1970392.1970394.

William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16 (1):8–37, March 1961. doi:10.1111/j.1540-6261.1961.tb02789.x.

Zihe Wang, Weiran Shen, and Song Zuo. Bayesian Nash equilibrium in first-price auction with discrete value distributions. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1458–1466, 2020. URL https://dl.acm.org/doi/abs/10.5555/3398761.3398929.

Jens Witkowski and David C. Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981, 2012. doi:10.1145/2229012.2229085.

Mihalis Yannakakis. Equilibria, fixed points, and complexity classes. *Computer Science Review*, 3(2): 71–85, May 2009. doi:10.1016/j.cosrev.2009.03.004.

APPENDIX

A The Input Model for the Value Distributions

Let $\mathcal{F}$ be a class of cumulative distribution functions on the interval $[0, 1]$. In other words, for any $F \in \mathcal{F}$ and any $x \in [0, 1]$, $F(x)$ is the probability of the interval $[0, x]$ according to F. For every $F \in \mathcal{F}$, let $\text{size}(F)$ denote the representation size of F, i.e., the number of bits needed to represent F. (Here we implicitly assume that some representation scheme is given in the definition of $\mathcal{F}$.)

For any rational number x, let $\text{size}(x)$ denote the representation size of x, namely the length of the binary representation of the denominator and numerator of x. The definitions in this section are based on the corresponding notions introduced by Etessami and Yannakakis [2010].

Definition 4. A class of cumulative distribution functions $\mathcal{F}$ is *polynomially computable*, if there exists some polynomial p such that for all $F \in \mathcal{F}$ and all rational $x \in [0, 1]$, $F(x)$ can be computed in time $p(\text{size}(F) + \text{size}(x))$.

In order to guarantee the existence of approximate equilibria with polynomial representation size we add an extra requirement on $\mathcal{F}$.

Definition 5. A class of cumulative distribution functions $\mathcal{F}$ is *polynomially continuous*, if there exists some polynomial q such that for all $F \in \mathcal{F}$ and all rational $\varepsilon > 0$, there exists rational $\delta > 0$ with $\text{size}(\delta) \leq q(\text{size}(F) + \text{size}(\varepsilon))$ such that

$$|F(x) - F(y)| \leq \varepsilon$$

for all $x, y \in [0, 1]$ with $|x - y| \leq \delta$.

Note that distribution functions given by piecewise-constant density functions on the interval $[0, 1]$ are an example of such a class of polynomially-computable and polynomially-continuous $\mathcal{F}$. The density functions are represented explicitly, i.e., as a list of "blocks", where for every block we give the sub-interval of $[0, 1]$ that it occupies and the height of the block.

B Impossibilities for Implicit Bidding Spaces

In Section 2, we emphasized that it is necessary for our computational problem to have the bidding space explicitly as part of the input, as otherwise it is hard to even compute the best responses of the auction. We provide more details on this topic in this section.

If the bidding space $B \subseteq [0, 1]$ is discrete but represented in some implicit way, this immediately gives rise to some computational obstacles. When we proved in Section 3 that best-responses could be computed efficiently, our procedure essentially goes over all possible bids, and checks which bid achieves the highest utility. If the bidding space is large (say, exponential in the input size), this approach is no longer efficient. In fact, in this subsection we will prove that, essentially, one cannot hope to find a better approach; in particular, we provide lower bounds from an information-theoretical as well as a computational perspective.

For simplicity, in this subsection we will assume that the bidding space is the set of all rational numbers in $[0, 1]$ that have denominator 2^m,

$$B = \left\{ \frac{p}{2^m} \mid 0 \leq p \leq 2^m \right\},$$

where m is part of the input and given in unary representation. Notice that each bid can then be encoded by a binary string of size m (with the exception of the bid 1, which can be encoded with $m + 1$ bits). We will also assume that there are only two bidders, each having a valuation over the unit interval, $V = [0, 1]$. This is arguably the simplest natural example one could consider.

As we explained in Section 2 we can identify a strategy by its set of jump points

$$\alpha_i(b) = \sup\{v \mid \beta_i(v) \leq b\}.$$

Intuitively, $\alpha_i(b)$ is the largest value for which player i would bid b or lower. At this point we have two options on how to represent the functions α_i:

- **Black-box model**: in the black-box model we have access to an oracle that, given a bid $b \in B$, returns the corresponding jump point $\alpha_i(b)$.
- **White-box model**: in the white-box model we have an algorithm that, given a bid $b \in B$, computes the jump point $\alpha_i(b)$. For example, this could be given by a circuit. Alternatively, we can assume that α_i is a function computable in polynomial time.

In both cases we need to describe how the jump points themselves are represented. For simplicity, we just assume that all jump points are rational quantities (as we are going for a hardness result).

Besides the inverse bidding strategies, we also need to represent the cumulative density functions $F_i : [0,1] \to [0,1]$. Here similar considerations apply, or we can use the notions in Appendix A.

Now, given F_i and α_i, an important quantity of interest is

$$\Pi_i(b) = F_i(\alpha_i(b));$$

since $\alpha_i(b)$ is the largest value for which bidder i will bid b or lower, and $F_i(\alpha_i(b))$ is the probability that bidder i's valuation is at most this value, it turns out that $\Pi_i(b)$ can be very naturally interpreted as the probability that player i bids on or below b. Notice that we can then get the probability that player i bids exactly b as $\Pi_i(b) - \Pi_i(b^-)$, where b^- is the bid immediately below b, for $b > 0$. Regarding the computation of Π_i, it will be either a black-box or white-box computation, depending on whether we have assumed α_i and F_i to be given in a black-box or white-box fashion.

As we already mentioned, in our reduction we will consider only two bidders. We shall fix the second bidder's bidding strategy and cumulative distribution function throughout the reduction, and look at the best-response of bidder 1. For ease of notation, we will drop the subscript 2 and write α, F, Π instead of α_2, F_2, Π_2; there will be no confusion since we will never look at bidder 1's valuation distribution or bidding strategy. Given a bid b, we can express the probability that bidder 1 wins the auction when bidding b, denoted by $H(b)$, via

$$\begin{aligned} H(0) &= \frac{1}{2}\Pi(0); \\ H(b) &= \Pi(b^-) + \frac{1}{2}\left(\Pi(b) - \Pi(b^-)\right) = \frac{1}{2}\left(\Pi(b^-) + \Pi(b)\right), \qquad \text{for } \ b > 0. \end{aligned}$$

Finally, we wish to maximize the utility of bidder 1; when she has a valuation of v and bids b, this is given by $u(v,b) = H(b)(v-b)$.

Now that we have given the preliminaries of our reduction, let us go into the construction. Let us fix some $m \geq 3$ and define a baseline instance. We will want to choose a bidding strategy α and distribution F for bidder 2, so that the resulting function $\Pi(\cdot)$ is given as follows.

$$\begin{aligned} \Pi(0) = \Pi(2^{-m}) &= \Pi(2 \cdot 2^{-m}) = 0; \\ \Pi(b - 2^{-m}) = \Pi(b) &= \Pi(b + 2^{-m}) = \Pi(b + 2 \cdot 2^{-m}) \\ &= \frac{1}{2(1-b)}, \quad \text{for} \quad b = p \cdot 2^{-m},\ p \text{ a multiple of 4, and } b \leq \frac{1}{2}; \\ \Pi(b) &= 1 \quad \text{for} \quad b \geq \frac{1}{2}. \end{aligned}$$

Our function Π essentially corresponds to a discrete probability distribution on the bids with the following properties. First, it only has mass at points of the form $(4k-1) \cdot 2^{-m}$, for positive integer

k, where $4k - 1 < 2^{m-1}$. Second, the mass at $3 \cdot 2^{-m}$ equals $\frac{1}{2(1-4\cdot 2^{-m})}$, whereas for $k \geq 2$ the mass at $(4k-1)\cdot 2^{-m}$ equals $\frac{1}{2(1-4k\cdot 2^{-m})} - \frac{1}{2(1-(4k-1)\cdot 2^{-m})}$. To yield the desired Π, we can for example take $F(x) = x$, corresponding to the uniform distribution on $[0,1]$, and $\alpha(b) = \Pi(b)$ defined as above.

Given the probability distribution Π on the bids of player 2, we are interested in computing the best-response strategy for player 1. In fact, we will do so for the case that player 1's valuation equals 1. If we can show it is hard to compute the best-response for this value, then it follows that it is hard to compute the best-response strategy function in general. Using the definition of $H(b)$, we can write

$$H(0) = 0; \quad H(2^{-m}) = 0; \quad H(2 \cdot 2^{-m}) = 0; \quad H(3 \cdot 2^{-m}) \frac{1}{4(1 - 4 \cdot 2^{-m})};$$

for $b = p \cdot 2^{-m}$, p a multiple of 4, and $b \leq \frac{1}{2} - 4 \cdot 2^{-m}$,

$$H(b) = \frac{1}{2(1-b)}; \quad H(b + 2^{-m}) = \frac{1}{2(1-b)}; \quad H(b + 2 \cdot 2^{-m}) = \frac{1}{2(1-b)};$$

$$H(b + 3 \cdot 2^{-m}) = \frac{1}{4(1-b)} + \frac{1}{4(1-b-4\cdot 2^{-m})};$$

finally, for $b \geq 1/2$,

$$H(b) = 1.$$

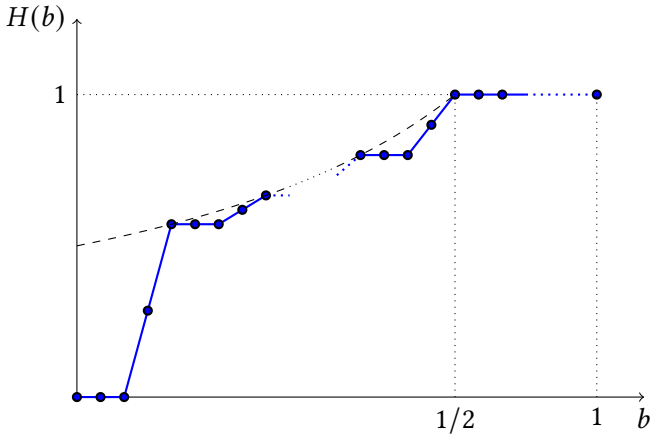

Figure 6: Depiction of the baseline construction. $H(b)$ denotes the probability of player 1 winning the auction when bidding b, and is represented by the blue circles. We also plot in dashed line the auxiliary function $x \mapsto \frac{1}{2(1-x)}$.

A graphical depiction of $H(b)$ can be found in Figure 6. It is not hard to check that, for every bid b, we have that

$$H(b) \leq \frac{1}{2(1-b)};$$

moreover, this is achieved with equality for every bid of the form $b = p \cdot 2^{-m}$, for p a multiple of 4, as long as $b \leq 1/2$. Therefore, the maximum utility that player 1 can achieve is $1/2$, and all such multiple-of-four bids are equally best-responses.

Now that we understand the baseline instance, we can construct a family of "perturbed" instances that will be used in our reduction. For a fixed subset $S \subseteq \{0,1\}^{m-3}$ of binary strings of size $m-3$, we will define a corresponding Π_S, H_S as follows. Π_S and H_S coincide with Π and H on every bid $b \geq 1/2$. For the bids smaller than $1/2$, we can write their binary expansion as a sequence of m bits, the first of which is 0. For example, if $m = 4$, then the bid $3/2^4$ can be written as 0011. For every $x \in \{0,1\}^{m-3}$, if $x \notin S$, then Π_S and H_S coincide with Π and H for bids of the form $0xb_1b_2$; in particular, if $b = x \cdot 2^{m-2}$,

- if $x = 0 \cdots 0$, then we have $\Pi_S(0) = 0, \Pi_S(2^{-m}) = 0, \Pi_S(2 \cdot 2^{-m}) = 0, \Pi_S(3 \cdot 2^{-m}) = \frac{1}{2(1-4\cdot 2^{-m})}$;
- otherwise, we have $\Pi_S(b) = \frac{1}{2(1-b)}, \Pi_S(b+2^{-m}) = \frac{1}{2(1-b)}, \Pi_S(b+2\cdot 2^{-m}) = \frac{1}{2(1-b)}, \Pi_S(b+3\cdot 2^{-m}) = \frac{1}{2(1-b-4\cdot 2^{-m})}$.

On the other hand, for $x \in S$ and $b = x \cdot 2^{-m+2}$, Π_S is obtained from Π by shifting the mass at $b + 3 \cdot 2^{-m}$ to $b + 2^{-m}$; in other words,

- if $x = 0 \cdots 0$, then we have $\Pi_S(0) = 0, \Pi_S(2^{-m}) = \frac{1}{2(1-4\cdot 2^{-m})}, \Pi_S(2\cdot 2^{-m}) = \frac{1}{2(1-4\cdot 2^{-m})}, \Pi_S(3\cdot 2^{-m}) = \frac{1}{2(1-4\cdot 2^{-m})}$;
- otherwise, we have $\Pi_S(b) = \frac{1}{2(1-b)}$, $\Pi_S(b + 2^{-m}) = \frac{1}{2(1-b-4\cdot 2^{-m})}$, $\Pi_S(b + 2 \cdot 2^{-m}) = \frac{1}{2(1-b-4\cdot 2^{-m})}$, $\Pi_S(b + 3 \cdot 2^{-m}) = \frac{1}{2(1-b-4\cdot 2^{-m})}$.

This gives rise to a change in H_S as well:

- if $x = 0 \cdots 0$, then we have $H_S(0) = 0, H_S(2^{-m}) = \frac{1}{4(1-4\cdot 2^{-m})}, H_S(2\cdot 2^{-m}) = \frac{1}{2(1-4\cdot 2^{-m})}, H_S(3\cdot 2^{-m}) = \frac{1}{2(1-4\cdot 2^{-m})}$;
- otherwise, we have $H_S(b) = \frac{1}{2(1-b)}$, $H_S(b + 2^{-m}) = \frac{1}{4(1-b)} + \frac{1}{4(1-b-4\cdot 2^{-m})}$, $H_S(b + 2 \cdot 2^{-m}) = \frac{1}{2(1-b-4\cdot 2^{-m})}$, $H_S(b + 3 \cdot 2^{-m}) = \frac{1}{2(1-b-4\cdot 2^{-m})}$.

Similarly as above, we can define an α_S for player 2 that give rise to this choice of Π_S and H_S. The net effect of our construction is that, for $x \notin S$, the bids of the form $0x00$, $0x01$, $0x10$ and $0x11$ achieve the same utility in both the baseline and the perturbed instances (and thus, at most $1/2$); but if $x \in S$, the bids of the form $0x01$, $0x10$, $0x11$ now achieve higher utility; in fact, if $x \in S$, then bidding $0x10$ achieves a utility strictly higher than $1/2$. Writing $b = x \cdot 2^{-m+2}$, we can see that

$$u(1, b + 2 \cdot 2^{-m}) = \frac{1 - b - 2 \cdot 2^{-m}}{2(1 - b - 4 \cdot 2^{-m})} > \frac{1}{2}.$$

We want to find an ε that bounds the utility gap, in order to show that computing ε-best-responses is hard. Using the trivial bound that $1 - b - 4 \cdot 2^{-m} < 1$, it turns out that $\varepsilon \leq 2^{-m}$ is small enough:

$$u(1, b + 2 \cdot 2^{-m}) - \frac{1}{2} = \frac{1 - b - 2 \cdot 2^{-m}}{2(1 - b - 4 \cdot 2^{-m})} - \frac{1 - b - 4 \cdot 2^{-m}}{2(1 - b - 4 \cdot 2^{-m})} = \frac{2 \cdot 2^{-m}}{2(1 - b - 4 \cdot 2^{-m})} > 2^{-m}.$$

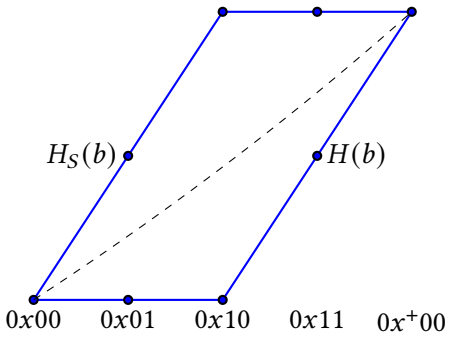

Figure 7: Depiction of the baseline construction. $H_S(b)$ denotes the probability of player 1 winning the auction when bidding b, and is represented by the upper blue circles. These are higher than the probabilities in $H(b)$ (lower blue circles), and go above the function $x \mapsto \frac{1}{2(1-x)}$ (dashed line). Here $0x^+00$ represents the binary string immediately after $0x11$.

We can depict the change from function H to function H_S as in Figure 7. To conclude this section, we just need to prove that one cannot distinguish between H and H_S unless we explicitly compute utilities for a large number possible bids.

Theorem B.1. *Consider a FPA where the bidding space corresponds to all dyadic rationals of order m, and bidding strategies are represented implicitly according to the black-box model. Then, any algorithm that computes ε-best-responses, for $\varepsilon \leq 2^{-m}$, makes an exponential number of queries in the worst-case.*

Proof. Let A be an algorithm that computes exact best-responses. Fix an integer $m \geq 3$, the number of players to be 2, and run the algorithm A for the baseline instance, where player 2 bids according to function H, and player 1's value is fixed to 1. Suppose that A makes less than $2^{m-3} - 1$ queries; let Q be the set of queries made by A, and b be the bid returned by A. Next, notice that there are 2^{m-3} disjoint sets of bids of the form $\{0x01, 0x10, 0x11\}$, one for each $x \in \{0,1\}^{m-3}$. Since A makes less than $2^{m-3} - 1$ queries, it follows that there must exist some x for which none of $0x01$, $0x10$, $0x11$ belongs to $Q \cup \{b\}$. Now consider the perturbed instance H_x, that is, we take $S = \{x\}$. Notice that H_S and Π_S coincide with H and Π everywhere except at $\{0x01, 0x10, 0x11\}$; therefore, running A on the instance H_x would produce the same answers on all queries, and so would produce the same best-response bid of b. However, by our construction we know that bidding b gives an utility of at most $1/2$, whereas bidding according to the string $0x1$ gives an utility strictly higher than $1/2 + \varepsilon$. Hence, the algorithm would not give a correct answer. We conclude that any algorithm for computing best-responses would have to make at least $2^{m-3} - 1$ queries. □

Theorem B.2. *Consider a FPA where the bidding space corresponds to all dyadic rationals of order m, and bidding strategies are represented implicitly according to the white-box model. Then, computing ε-best-responses, for exponentially small ε, is an NP-hard optimization problem.*

Proof. Let $\mathcal{P}$ be any problem in NP. Without loss of generality assume that certificates for instances of size n must all have size $p(n)$, for some polynomial p. Given an input y for $\mathcal{P}$, let $S(y) = \{x \in \{0,1\}^{p(n)} \ : \ x \text{ is a valid certificate for } y\} \subseteq \{0,1\}^{p(n)}$ be the set of valid certificates for y. In other words, y is a yes-instance if and only if $S(y) \neq \emptyset$; and there is a polynomial-time algorithm that, given x, y, decides whether $x \in S(y)$.

We can define our reduction, from $\mathcal{P}$ to the problem of computing best-responses, as follows. Given an input y of size n, consider a first-price auction where:

- $m = p(n) + 3$ and $\varepsilon = 2^{-m}$;
- the bidding space corresponds to all dyadic rationals of order m;
- there are two players; the second player has a bidding distribution according to the perturbed instance $H_{S(y)}$;
- the first player has a valuation of 1.

Notice that we can indeed construct this auction in polynomial time. In particular, there is an algorithm that computes $\Pi_{S(y)}(b)$ as follows. If $b \geq 1/2$, then $\Pi_{S(y)}(b) = 1$. Otherwise, write b in the form $0xb_1b_2$; decide whether $x \in S(y)$ (in polynomial time); depending on the answer, compute $\Pi_{S(y)}(b)$ according to the formulas above.

To complete the proof, suppose y is a no-instance. Then $\Pi_{S(y)} = \Pi_\emptyset = \Pi$, and the best-response for player 1 in this auction must achieve utility of exactly $1/2$, so that any ε-best-response achieves utility at most $1/2$. On the other hand, suppose y is a yes-instance. Then $S(y) \neq \emptyset$, and the best-response for player 1 in this auction must achieve utility strictly higher than $1/2 + 2^{-m}$, so that any ε-best-response achieves utility strictly higher than $1/2$. □

C Exact Equilibria Can Be Irrational

In this section we provide the technical details on Example 1, which shows that a FPA can have *only* irrational equilibria. Recall that in Section 2 we imposed two standard assumptions in the literature, namely that equilibrium strategies are *monotone non-decreasing* and exhibit *no overbidding*. Here, since

we would like to argue that *all equilibria* are irrational, to make our statement even stronger, we will show that in the example that we construct, there is a unique equilibrium which is not only irrational, but also in non-decreasing strategies and in which the bidders are not overbidding.

To this end, we start with the following proposition that states that essentially, violations of overbidding and monotonicity only occur in trivial corner cases. In our subsequent construction, such cases will not occur.

Proposition C.1. *Let $\boldsymbol{\beta}$ be an exact equilibrium of a FPA. For a bid b_i by player i, let $H_i(b_i, \boldsymbol{\beta}_{-i})$ denote the (perceived) probability that player i gets the item, given this bid and the bidding strategies by the other players. Then, strategies will always be no over-bidding and monotone non-decreasing except only possibly when the probability of winning is zero. Formally,*

1. *let v_i be a valuation by player i and $b_i = \beta_i(v_i)$. If $b_i > v_i$, then $H_i(b_i, \boldsymbol{\beta}_{-i}) = 0$;*
2. *let v_i, v_i' be valuations by player i and $b_i = \beta_i(v_i)$, $b_i' = \beta_i(v_i')$. If $v_i < v_i'$ and $b_i > b_i'$, then $H_i(b_i, \boldsymbol{\beta}_{-i}) = H_i(b_i', \boldsymbol{\beta}_{-i}) = 0$.*

Proof.

1. If $b_i > v_i$ and $H_i(b_i, \boldsymbol{\beta}_{-i}) > 0$, then player i achieves a strictly negative utility by bidding b_i when her valuation is v_i. However, player i could achieve non-negative utility by bidding below v_i (e.g. by bidding 0). Hence $\boldsymbol{\beta}$ would not be an equilibrium.

2. Suppose that $v_i < v_i'$ and $b_i > b_i'$. As $\boldsymbol{\beta}$ is an exact equilibrium, we know that b_i, b_i' are the best bidding responses by player i. In other words, $u_i(b_i, \boldsymbol{\beta}_{-i}; v_i) \geq u_i(b_i', \boldsymbol{\beta}_{-i}; v_i)$ and $u_i(b_i', \boldsymbol{\beta}_{-i}; v_i') \geq u_i(b_i, \boldsymbol{\beta}_{-i}; v_i')$. Moreover, as $b_i > b_i'$ we also know that $H_i(b_i, \boldsymbol{\beta}_{-i}) \geq H_i(b_i', \boldsymbol{\beta}_{-i})$. Putting these together, we find that

$$\begin{aligned}
u_i(b_i, \boldsymbol{\beta}_{-i}; v_i') + u_i(b_i', \boldsymbol{\beta}_{-i}; v_i) &= (v_i' - b_i)H_i(b_i, \boldsymbol{\beta}_{-i}) + (v_i - b_i')H_i(b_i', \boldsymbol{\beta}_{-i}) \\
&= (v_i' - v_i)H_i(b_i, \boldsymbol{\beta}_{-i}) + (v_i - b_i)H_i(b_i, \boldsymbol{\beta}_{-i}) + (v_i - b_i')H_i(b_i', \boldsymbol{\beta}_{-i}) \\
&\geq (v_i' - v_i)H_i(b_i', \boldsymbol{\beta}_{-i}) + (v_i - b_i)H_i(b_i, \boldsymbol{\beta}_{-i}) + (v_i - b_i')H_i(b_i', \boldsymbol{\beta}_{-i}) \\
&= (v_i' - b_i')H_i(b_i', \boldsymbol{\beta}_{-i}) + (v_i - b_i)H_i(b_i, \boldsymbol{\beta}_{-i}) \\
&= u_i(b_i', \boldsymbol{\beta}_{-i}; v_i') + u_i(b_i, \boldsymbol{\beta}_{-i}; v_i).
\end{aligned}$$

   From this, we conclude that all steps in the above derivation must hold with equality, implying that $u_i(b_i, \boldsymbol{\beta}_{-i}; v_i) = u_i(b_i', \boldsymbol{\beta}_{-i}; v_i)$, $u_i(b_i', \boldsymbol{\beta}_{-i}; v_i') = u_i(b_i, \boldsymbol{\beta}_{-i}; v_i')$ and $H_i(b_i, \boldsymbol{\beta}_{-i}) = H_i(b_i', \boldsymbol{\beta}_{-i})$. But then $0 = u_i(b_i', \boldsymbol{\beta}_{-i}; v_i) - u_i(b_i, \boldsymbol{\beta}_{-i}; v_i) = (b_i - b_i')H_i(b_i, \boldsymbol{\beta}_{-i})$. As $b_i > b_i'$ we conclude that $H_i(b_i, \boldsymbol{\beta}_{-i}) = H_i(b_i', \boldsymbol{\beta}_{-i}) = 0$.

□

We are now ready to proceed with the example showing that all equilibria of the FPA can be irrational. Consider a first-price auction with $n = 3$ bidders and common priors, whose valuations are independently and identically distributed according to the uniform distribution on $[0, 1]$; that is, $F_i(x) = x$ for $i = 1, 2, 3$. Let the bidding space be $B = \{0, 1/2\}$. Clearly, this auction can be represented with piecewise-constant density functions (with a single piece) and with a finite number of rational quantities. We shall show that the auction has a unique equilibrium, and that this equilibrium is described by an irrational jump point.

First observe that, at an exact equilibrium, the probability of a player winning when bidding 0 is positive. Otherwise, one of the other players would be bidding $1/2$ with probability 1, and would achieve expected negative utility when having a valuation in $[0, 1/2)$, which contradicts the best-response conditions. Since the probability of winning is never zero, Proposition C.1 implies that any equilibrium

must consist of non-overbidding, monotone non-decreasing strategies. In particular, the best response strategy of a player i can be described by a single jump point a_i, that is,

$$\beta_i(x) = \begin{cases} 0 & \text{if } 0 \le x \le a_i; \\ 1/2 & \text{if } a_i < x \le 1. \end{cases}$$

Since strategies are non-overbidding, we must have that $1/2 \le a_i \le 1$. Moreover, a joint strategy profile can be described by the jump points of each player, which form a triple (a_1, a_2, a_3).

Next we show that, at an exact equilibrium, each of the a_i must be strictly less than 1. Suppose that bidder 1 has a valuation of v_1 and that bidders 2 and 3 have played according to (a_2, a_3). This means that bidder 2 bids 0 with probability a_2 and bids $1/2$ with probability $(1 - a_2)$, and similarly for bidder 3. Thus, the probability of player 1 winning when bidding 0 or when bidding $1/2$ is, respectively,

$$\begin{aligned} H(0; a_2, a_3) &= \frac{1}{3} a_2 a_3, \\ H(1/2; a_2, a_3) &= a_2 a_3 + \frac{1}{2} a_2 (1 - a_3) + \frac{1}{2}(1 - a_2) a_3 + \frac{1}{3}(1 - a_2)(1 - a_3) \\ &= \frac{1}{3} + \frac{1}{3} a_2 a_3 + \frac{1}{6} a_2 + \frac{1}{6} a_3. \end{aligned}$$

From this we can compute the utility of player 1 when bidding 0 or when bidding $1/2$,

$$\begin{aligned} u_1(v_1, 0; a_2, a_3) &= \frac{1}{3} a_2 a_3 v_1 \\ u_1(v_1, 1/2; a_2, a_3) &= \left(\frac{1}{3} + \frac{1}{3} a_2 a_3 + \frac{1}{6} a_2 + \frac{1}{6} a_3 \right) \left(v_1 - \frac{1}{2} \right). \end{aligned}$$

We can compute the jump point v_1 for which player 1 is indifferent between bidding 0 or bidding $1/2$, by solving the equation

$$\begin{aligned} & u_1(v_1, 0; a_2, a_3) = u_1(v_1, 1/2; a_2, a_3) \\ \Rightarrow \quad & \frac{1}{3} a_2 a_3 v_1 = \left(\frac{1}{3} + \frac{1}{3} a_2 a_3 + \frac{1}{6} a_2 + \frac{1}{6} a_3 \right) \left(v_1 - \frac{1}{2} \right) \\ \Rightarrow \quad & 2 v_1 + v_1 a_2 + v_1 a_3 = 1 + a_2 a_3 + \frac{1}{2} a_2 + \frac{1}{2} a_3 \qquad (20) \\ \Rightarrow \quad & v_1 = \frac{1}{2} + \frac{a_2 a_3}{2 + a_2 + a_3}. \qquad (21) \end{aligned}$$

Next observe that the expression $\frac{a_2 a_3}{2 + a_2 + a_3}$ is increasing in both a_2 and a_3; hence, by setting $a_2 = 1$, $a_3 = 1$ we get that the right hand side of (21) is at most $\frac{1}{2} + \frac{1 \times 1}{2+1+1} = \frac{3}{4}$. In other words, the break-even point must occur in the interval $[1/2, 3/4]$, and thus in particular setting $v_1 = a_1$ must give a solution to (20).

Repeating this argument for players 2 and 3 we obtain similarly that a_2 and a_3 must lie in $[1/2, 3/4]$, and that these jump points must be the solutions of equations similar to (20). In order for (a_1, a_2, a_3) to define an equilibrium, each player's jump point must be optimal in response to the other players' strategies. Thus, (a_1, a_2, a_3) must be a solution of the system of equations

$$\begin{aligned} 2 a_1 + a_1 a_2 + a_1 a_3 &= 1 + a_2 a_3 + \frac{1}{2} a_2 + \frac{1}{2} a_3 \qquad (22) \\ 2 a_2 + a_1 a_2 + a_2 a_3 &= 1 + a_1 a_3 + \frac{1}{2} a_1 + \frac{1}{2} a_3 \qquad (23) \\ 2 a_3 + a_2 a_3 + a_1 a_3 &= 1 + a_1 a_2 + \frac{1}{2} a_1 + \frac{1}{2} a_2 \end{aligned}$$

Finally, we show that the above system has a unique solution. By subtracting (23) from (22), we get

$$\begin{aligned} & 2(a_1 - a_2) + a_3(a_1 - a_2) = (a_2 - a_1)a_3 + \frac{1}{2}(a_2 - a_1) \\ \Rightarrow \quad & \left(\frac{5}{2} + 2a_3\right)(a_1 - a_2) = 0 \\ \Rightarrow \quad & a_3 = -\frac{5}{4} \quad \text{or} \quad a_1 = a_2. \end{aligned}$$

Since we know that $a_3 \in [1/2, 3/4]$, we conclude that $a_1 = a_2$. By the same argument, we must have $a_2 = a_3$ and $a_1 = a_3$, that is, any equilibrium must be symmetric. Now letting $a := a_1 = a_2 = a_3$, we get that a must be a solution to the equation

$$\begin{aligned} & 2a + a^2 + a^2 = 1 + a^2 + \frac{1}{2}a + \frac{1}{2}a \\ \Rightarrow \quad & a^2 + a - 1 = 0 \\ \Rightarrow \quad & a = \frac{-1 \pm \sqrt{5}}{2}. \end{aligned}$$

Since a must be positive, we conclude that the unique equilibrium of this auction is given by the jump point $a = \frac{-1+\sqrt{5}}{2} \approx 0.618$ (the inverse of the golden ratio), which is irrational.

D Proof of Lemma 4.3

Since the distributions are polynomially-continuous, it follows that given any $\varepsilon > 0$, we can compute $\delta > 0$ in polynomial time such that $|F_{i,j}(x) - F_{i,j}(y)| \le \varepsilon/2^{n+1}$ for all x, y with $|x - y| \le \delta$ and all $i, j \in N$ $(i \ne j)$.

Consider any $\boldsymbol{\alpha}, \boldsymbol{\alpha}' \in \mathcal{D}$ (see the proof of Theorem 4.1) with $\|\boldsymbol{\alpha} - \boldsymbol{\alpha}'\|_\infty \le \delta$. Then, we have

$$\begin{aligned} \left| \Pr_{v_i \sim F_{j,i}} [\beta_i(v_i) \le b] - \Pr_{v_i \sim F_{j,i}} \left[\beta_i'(v_i) \le b\right] \right| &\le \left| \Pr_{v_i \sim F_{j,i}} [v_i \le \alpha_i(b)] - \Pr_{v_i \sim F_{j,i}} \left[v_i \le \alpha_i'(b)\right] \right| \\ &\le \left| F_{j,i}(\alpha_i(b)) - F_{j,i}(\alpha_i'(b)) \right| \\ &\le \varepsilon/2^{n+1} \end{aligned}$$

for all $i, j \in N$ $(i \ne j)$ and $b \in B$. It follows that $\Pr_{v_i \sim F_{j,i}} [\beta_i(v_i) < b]$ differs from $\Pr_{v_i \sim F_{j,i}} \left[\beta_i'(v_i) < b\right]$ by at most $\varepsilon/2^{n+1}$. Similarly, $\Pr_{v_i \sim F_{j,i}} [\beta_i(v_i) = b]$ differs from $\Pr_{v_i \sim F_{j,i}} \left[\beta_i'(v_i) = b\right]$ by at most $\varepsilon/2^n$.

Let $T_i(b, \ell; \boldsymbol{\alpha}_{-i})$ denote the probability that, from the perspective of bidder i, exactly ℓ out of the bidders $N \setminus \{i\}$ bid exactly b, and the remaining $n - 1 - \ell$ bidders bid below b. We can write

$$T_i(b, \ell; \boldsymbol{\alpha}_{-i}) = \sum_{\substack{S \subseteq N \setminus \{i\} \\ |S| = \ell}} \prod_{k \in S} \Pr_{v_k \sim F_{i,k}} [\beta_k(v_k) = b] \prod_{k \in N \setminus (\{i\} \cup S)} \Pr_{v_k \sim F_{i,k}} [\beta_k(v_k) < b] .$$

From this it follows that $T_i(b, \ell; \boldsymbol{\alpha}_{-i})$ and $T_i(b, \ell; \boldsymbol{\alpha}'_{-i})$ differ by at most $\binom{n-1}{\ell} n\varepsilon/2^n$, for all $i \in N$, $b \in B$ and $\ell \in \{0, 1, \ldots, n-1\}$. As defined in Section 3, recall that $H_i(b, \boldsymbol{\alpha}_{-i})$ denotes the probability that bidder i wins if she bids b and the other bidders bid according to $\boldsymbol{\alpha}_{-i}$. Then, we can write

$$H_i(b, \boldsymbol{\alpha}_{-i}) = \sum_{\ell=0}^{n-1} \frac{1}{\ell + 1} T_i(b, \ell; \boldsymbol{\alpha}_{-i}).$$

It follows that $H_i(b, \boldsymbol{\alpha}_{-i})$ differs from $H_i(b, \boldsymbol{\alpha}'_{-i})$ by at most

$$\sum_{\ell=0}^{n-1} \frac{1}{\ell + 1} \binom{n-1}{\ell} n\varepsilon/2^n = \sum_{\ell=0}^{n-1} \binom{n}{\ell + 1} \varepsilon/2^n \le \varepsilon$$

for all $i \in N$ and $b \in B$. Finally, note that $u_i(b, \boldsymbol{\alpha}_{-i}; v_i) = H_i(b, \boldsymbol{\alpha}_{-i}) \cdot (v_i - b)$. Thus, we obtain

$$\left|u_i(b, \boldsymbol{\alpha}_{-i}; v_i) - u_i(b, \boldsymbol{\alpha}'_{-i}; v_i)\right| \leq \left|H_i(b, \boldsymbol{\alpha}_{-i}) - H_i(b, \boldsymbol{\alpha}'_{-i})\right| |v_i - b| \leq \varepsilon$$

for all $i \in N$, $b \in B$ and $v_i \in [0, 1]$, since $|v_i - b| \leq 1$.

E PPAD and FIXP-completeness of Generalized Circuit Variants

E.1 PPAD-completeness (Proof of Proposition 5.3)

Membership in PPAD follows from the fact that a generalized circuit with gates $g_1, \ldots, g_{\mathfrak{m}}$ can be interpreted as defining an algebraic circuit $F : [0, 1]^{\mathfrak{m}} \to [0, 1]^{\mathfrak{m}}$, where for $x \in [0, 1]^{\mathfrak{m}}$ and $i \in [\mathfrak{m}]$ we let $F_i(x) = G(x_j, x_k)$, where $g_i = (G, j, k)$. Then, it is known that the problem of computing an ε-approximate fixed point of such a function F lies in PPAD [Etessami and Yannakakis, 2010] (and in fact, even when ε is provided in the input in binary representation). Finally, note that an ε-approximate fixed point of F exactly corresponds to an ε-satisfying assignment for the generalized circuit.

In order to prove PPAD-hardness, consider the ε-Gcircuit problem with gate-types $\mathcal{G} = \{G_{1-}, G_+\}$, for some sufficiently small constant $\varepsilon > 0$ (which will be set later). We begin by showing that additional gate-types can be simulated if we allow a larger (but still constant) error.

$G_=$: Copy. The goal of such a gate is to copy the value of some gate g_1. For this, we use the fact that $1 - (1 - x) = x$. Thus, we introduce a gate g_2 of type G_{1-} with input g_1 and a gate g_3 of type G_{1-} with input g_2. It holds that $\mathbf{v}[g_3] = 1 - \mathbf{v}[g_2] \pm \varepsilon = \mathbf{v}[g_1] \pm 2\varepsilon$. In other words, we can simulate a copy gate with error at most 2ε.

G_1: Constant 1. In order to obtain a gate that has value 1, we use the fact that $x + (1 - x) = 1$. First, we introduce an arbitrary gate g_1. Then, we introduce a gate g_2 of type G_{1-} with input g_1, and a gate g_3 of type G_+ with inputs g_1 and g_2. It holds that $\mathbf{v}[g_3] = \mathrm{T}(\mathbf{v}[g_1] + \mathbf{v}[g_2]) \pm \varepsilon = 1 \pm 2\varepsilon$. Thus, we can simulate a constant 1 with error at most 2ε.

G_-: Subtraction. The goal of this gate is to compute $\mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2])$. For this, we use the identity

$$\mathrm{T}(x - y) = 1 - \mathrm{T}\left((1 - x) + y\right)$$

which allows us to express subtraction using only addition and the complement operation. With this in hand, we can implement subtraction as follows. We introduce a gate g_3 of type G_{1-} with input g_1, a gate g_4 of type G_+ with inputs g_3 and g_2, and finally a gate g_5 of type G_{1-} with input g_4. Then, it holds that $\mathbf{v}[g_5] = 1 - \mathbf{v}[g_4] \pm \varepsilon = 1 - \mathrm{T}(\mathbf{v}[g_3] + \mathbf{v}[g_2]) \pm 2\varepsilon = 1 - \mathrm{T}(1 - \mathbf{v}[g_1] + \mathbf{v}[g_2]) \pm 3\varepsilon = \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2]) \pm 3\varepsilon$. Thus, we can simulate a subtraction gate with error at most 3ε.

$G_{/2}$: Division by 2. The goal of this gate is to compute $\mathbf{v}[g_1]/2$. This is achieved by constructing a cycle. Namely, we introduce two gates g_2 and g_3. The gate g_2 is of type G_- with inputs g_1 and g_3, and the gate g_3 is of type $G_=$ with input g_2. As a result, it holds that

$$\mathbf{v}[g_3] = \mathbf{v}[g_2] \pm 2\varepsilon = \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_3]) \pm 5\varepsilon.$$

From this, it follows that $\mathbf{v}[g_3] = \mathbf{v}[g_1]/2 \pm 5\varepsilon$. To see this, note that if $\mathbf{v}[g_3] \geq \mathbf{v}[g_1]$, then $\mathbf{v}[g_3] = 0 \pm 5\varepsilon = \mathbf{v}[g_1]/2 \pm 5\varepsilon$, since $[0, 5\varepsilon] \subseteq [\mathbf{v}[g_1]/2 - 5\varepsilon, \mathbf{v}[g_1]/2 + 5\varepsilon]$ (because $\mathbf{v}[g_1]/2 \leq \mathbf{v}[g_3]/2 \leq 5\varepsilon$). On the other hand, if $\mathbf{v}[g_3] < \mathbf{v}[g_1]$, then we obtain that $2\mathbf{v}[g_3] = \mathbf{v}[g_1] \pm 5\varepsilon$, which again yields the same conclusion, namely $\mathbf{v}[g_3] = \mathbf{v}[g_1]/2 \pm 5\varepsilon$. Thus, we can simulate division by 2 with error at most 5ε.

$G_{\times\zeta}$: Multiplication by $\zeta \in [0, 1]$. If $\zeta = 0$, then we can simply output $G_{1-}(G_1) = 0 \pm 3\varepsilon$. If $\zeta = 1$, we can simply use a $G_=$ gate that has error at most 2ε. Consider now the case where $\zeta \in (0, 1)$. Let $k = \lceil \log_2(1/\varepsilon) \rceil$. Recall that ε will be a fixed constant, so k will also be a fixed constant. It is easy to see that in polynomial time (in the representation size of ζ) we can find $a \in \{1, 2, \ldots, 2^k - 1\}$ such that $|\zeta - a/2^k| \leq \varepsilon$.

Let g_1 denote the input. Our goal now is to compute $(a/2^k) \cdot \mathbf{v}[g_1]$, since this will be ε-close to $\zeta \cdot \mathbf{v}[g_1]$. We compute $(a/2^k) \cdot \mathbf{v}[g_1]$ in a careful manner to ensure that the error remains small. This is achieved as follows. Using the binary representation of $a = \sum_{i=0}^{k-1} a_i 2^i$, $a_i \in \{0,1\}$, we can express the product $(a/2^k) \cdot x$ as

$$\frac{\frac{0+a_0\frac{x}{2}}{2}+a_1\frac{x}{2}}{2} + a_2\frac{x}{2}$$
$$\ddots$$

We implement this as follows. First, introduce g_2 such that $\mathbf{v}[g_2] = \mathbf{v}[g_1]/2 \pm 5\varepsilon$. Next, introduce g_3 such that (i) if $a_0 = 0$, then $\mathbf{v}[g_3] = 0 \pm 3\varepsilon$, (ii) if $a_0 = 1$, then $g_3 = g_2$. In both cases we have

$$\mathbf{v}[g_3] = a_0\mathbf{v}[g_2] \pm 3\varepsilon.$$

Next, introduce g_4 such that (i) if $a_1 = 0$, then $\mathbf{v}[g_4] = \mathbf{v}[g_3]/2 \pm 5\varepsilon = a_0\mathbf{v}[g_2]/2 \pm 5(1+1/2)\varepsilon$, (ii) if $a_1 = 1$, then $\mathbf{v}[g_4] = \mathbf{v}[g_3]/2 + \mathbf{v}[g_2] \pm 6\varepsilon = a_0\mathbf{v}[g_2]/2 + \mathbf{v}[g_2] \pm 6(1+1/2)\varepsilon$. In both cases we have

$$\mathbf{v}[g_4] = a_0\mathbf{v}[g_2]/2 + a_1\mathbf{v}[g_2] \pm 6(1+1/2)\varepsilon = (a_0+2a_1)\mathbf{v}[g_2]/2 \pm 6(1+1/2)\varepsilon.$$

Next, introduce g_5 such that (i) if $a_2 = 0$, then $\mathbf{v}[g_5] = \mathbf{v}[g_4]/2 \pm 5\varepsilon = (a_0+2a_1)\mathbf{v}[g_2]/4 \pm 6(1+1/2+1/4)\varepsilon$, (ii) if $a_2 = 1$, then $\mathbf{v}[g_5] = \mathbf{v}[g_4]/2 + \mathbf{v}[g_2] \pm 6\varepsilon = (a_0+2a_1)\mathbf{v}[g_2]/4 + \mathbf{v}[g_2] \pm 6(1+1/2+1/4)\varepsilon$. In both cases we have

$$\begin{aligned}\mathbf{v}[g_5] &= (a_0+2a_1)\mathbf{v}[g_2]/4 + a_2\mathbf{v}[g_2] \pm 6(1+1/2+1/4)\varepsilon \\ &= (a_0+2a_1+4a_2)\mathbf{v}[g_2]/4 \pm 6(1+1/2+1/4)\varepsilon.\end{aligned}$$

Continuing in the same manner, it follows by induction that after $k-1$ such steps we obtain

$$\mathbf{v}[g_{k+2}] = \left(\sum_{i=0}^{k-1} a_i 2^i\right)\mathbf{v}[g_2]/2^{k-1} \pm 12\varepsilon = \frac{a}{2^k}\left(2\mathbf{v}[g_2]\right) \pm 12\varepsilon = \frac{a}{2^k}\mathbf{v}[g_1] \pm 22\varepsilon = \zeta \cdot \mathbf{v}[g_1] \pm 23\varepsilon.$$

Thus, we can compute multiplication by $\zeta \in [0,1]$ with error at most 23ε. Note that this gadget can be constructed in polynomial time in the representation size of ζ. Furthermore, the number of gates needed to construct the gadget is $O(k)$, which is constant, since $k = \lceil \log_2(1/\varepsilon) \rceil$ and ε will be a fixed constant.

We are now ready to show PPAD-hardness. To do this, we reduce from a slightly modified version of Gcircuit studied by Goldberg et al. [2022], that we call Gcircuit$^{[-1,1]}$. This modified version operates on $[-1,1]$ instead of $[0,1]$, and it uses the gates $G_+^{[-1,1]}$, $G_1^{[-1,1]}$ and $G_{\times-\zeta}^{[-1,1]}$ (where the gates truncate to $[-1,1]$, and $\zeta \in [0,1]$). Goldberg et al. [2022] proved that ε'-Gcircuit$^{[-1,1]}$ is PPAD-hard for some sufficiently small constant $\varepsilon' > 0$. We now set $\varepsilon := \varepsilon'/50$. Below, we show that ε'-Gcircuit$^{[-1,1]}$ reduces to ε-Gcircuit (with gate-types $\mathcal{G} = \{G_{1-}, G_+\}$).

Given a generalized circuit with gates $G_+^{[-1,1]}$, $G_1^{[-1,1]}$ and $G_{\times-\zeta}^{[-1,1]}$, we construct a corresponding circuit with gates G_{1-} and G_+ as follows. Every gate g of the original circuit is replaced by two gates g^+ and g^-. The idea is that the value of g, which lies in $[-1,1]$, will be encoded by the values of g^+ and g^-, which lie in $[0,1]$. Formally, we interpret $\mathbf{v}[g] := \mathbf{v}[g^+] - \mathbf{v}[g^-]$. Next, we show that the constraints of the original circuit can be enforced by corresponding constraints on the new circuit.

Simulating $G_1^{[-1,1]}$. In order to enforce that $\mathbf{v}[g] = 1 \pm \varepsilon'$, we proceed as follows. We simply let $\mathbf{v}[g^+] = 1 \pm 2\varepsilon$ and $\mathbf{v}[g^-] = 0 \pm 3\varepsilon$ (using the constructions described above). Thus, it holds that $\mathbf{v}[g] = \mathbf{v}[g^+] - \mathbf{v}[g^-] = 1 \pm 5\varepsilon = 1 \pm \varepsilon'$.

Simulating $G_{\times-\zeta}^{[-1,1]}$. In order to enforce that $\mathbf{v}[g_2] = -\zeta \cdot \mathbf{v}[g_1] \pm \varepsilon'$, for some $\zeta \in [0,1]$, we proceed as follows. Using the constructions described above, we can enforce that $\mathbf{v}[g_2^+] = \zeta \cdot \mathbf{v}[g_1^-] \pm 23\varepsilon$ and $\mathbf{v}[g_2^-] = \zeta \cdot \mathbf{v}[g_1^+] \pm 23\varepsilon$. Thus, $\mathbf{v}[g_2] = -\zeta \cdot \mathbf{v}[g_1] \pm 46\varepsilon = -\zeta \cdot \mathbf{v}[g_1] \pm \varepsilon'$.

Simulating $G_+^{[-1,1]}$. In order to enforce that $\mathbf{v}[g_3] = \mathrm{T}_{[-1,1]}(\mathbf{v}[g_1] + \mathbf{v}[g_2]) \pm \varepsilon'$, we proceed in two steps. First, using our construction for performing subtraction, we "normalize" the gates by letting $\mathbf{v}[h_1^+] = \mathrm{T}(\mathbf{v}[g_1^+] - \mathbf{v}[g_1^-]) \pm 3\varepsilon$ and $\mathbf{v}[h_1^-] = \mathrm{T}(\mathbf{v}[g_1^-] - \mathbf{v}[g_1^+]) \pm 3\varepsilon$, which yields $\mathbf{v}[h_1] = \mathbf{v}[g_1] \pm 6\varepsilon$. We similarly obtain h_2 from g_2. This "normalization" will ensure that addition is then performed correctly.

In the second step, using the addition gate G_+, we let $\mathbf{v}[g_3^+] = \mathrm{T}(\mathbf{v}[h_1^+] + \mathbf{v}[h_2^+]) \pm \varepsilon$ and $\mathbf{v}[g_3^-] = \mathrm{T}(\mathbf{v}[h_1^-] + \mathbf{v}[h_2^-]) \pm \varepsilon$. Thus, it holds that

$$\begin{aligned}\mathbf{v}[g_3] = \mathbf{v}[g_3^+] - \mathbf{v}[g_3^-] &= \mathrm{T}(\mathbf{v}[h_1^+] + \mathbf{v}[h_2^+]) - \mathrm{T}(\mathbf{v}[h_1^-] + \mathbf{v}[h_2^-]) \pm 2\varepsilon \\ &= \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1^+] + \mathbf{v}[h_2^+]) - \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1^-] + \mathbf{v}[h_2^-]) \pm 2\varepsilon.\end{aligned}$$

Because of the "normalization" step, we know that

$$\min\{\mathbf{v}[h_1^+], \mathbf{v}[h_1^-]\} \le 3\varepsilon \quad \text{and} \quad \min\{\mathbf{v}[h_2^+], \mathbf{v}[h_2^-]\} \le 3\varepsilon.$$

In the case where $\mathbf{v}[h_1^-] \le 3\varepsilon$ and $\mathbf{v}[h_2^-] \le 3\varepsilon$, it holds that $\mathbf{v}[h_1] = \mathbf{v}[h_1^+] \pm 3\varepsilon$ and $\mathbf{v}[h_2] = \mathbf{v}[h_2^+] \pm 3\varepsilon$, which implies that

$$\mathbf{v}[g_3] = \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1] + \mathbf{v}[h_2]) - \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1^-] + \mathbf{v}[h_2^-]) \pm 8\varepsilon = \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1] + \mathbf{v}[h_2]) \pm 14\varepsilon.$$

In the case where $\mathbf{v}[h_1^+] \le 3\varepsilon$ and $\mathbf{v}[h_2^-] \le 3\varepsilon$, it holds that $\mathbf{v}[h_1] = -\mathbf{v}[h_1^-] \pm 3\varepsilon$ and $\mathbf{v}[h_2] = \mathbf{v}[h_2^+] \pm 3\varepsilon$, which implies that

$$\begin{aligned}\mathbf{v}[g_3] = \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1^+] + \mathbf{v}[h_2]) - \mathrm{T}_{[-1,1]}(-\mathbf{v}[h_1] + \mathbf{v}[h_2^-]) \pm 8\varepsilon &= \mathbf{v}[h_1] + \mathbf{v}[h_2] \pm 14\varepsilon \\ &= \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1] + \mathbf{v}[h_2]) \pm 14\varepsilon.\end{aligned}$$

The remaining two cases are handled in the same way, and thus we always obtain that

$$\mathbf{v}[g_3] = \mathrm{T}_{[-1,1]}(\mathbf{v}[h_1] + \mathbf{v}[h_2]) \pm 14\varepsilon = \mathrm{T}_{[-1,1]}(\mathbf{v}[g_1] + \mathbf{v}[g_2]) \pm 26\varepsilon = \mathrm{T}_{[-1,1]}(\mathbf{v}[g_1] + \mathbf{v}[g_2]) \pm \varepsilon'.$$

Clearly, this construction can be performed in polynomial time in the size of the original generalized circuit. Furthermore, given any ε-satisfying assignment of the new generalized circuit, we can easily obtain an ε'-satisfying assignment of the original generalized circuit by setting $\mathbf{v}[g] := \mathbf{v}[g^+] - \mathbf{v}[g^-] \in [-1, 1]$ for all gates g. It follows that the ε-Gcircuit problem with gate-types $\mathcal{G} = \{G_{1-}, G_+\}$ is PPAD-hard.

Finally, note that if we let $\mathcal{G} = \{G_1, G_-\}$ instead, we again obtain the same result, because G_{1-} and G_+ can easily be simulated. Indeed, it is clear that G_{1-} can immediately be simulated. Furthermore, G_+ can be simulated by using the equation $\mathrm{T}(x + y) = 1 - \mathrm{T}((1 - x) - y)$.

E.2 FIXP-completeness (Proof of Proposition 5.4)

Membership in FIXP follows immediately by noting that a generalized circuit with gates $g_1, \dots, g_\mathfrak{m}$ defines an algebraic circuit $F : [0,1]^\mathfrak{m} \to [0,1]^\mathfrak{m}$, where for $x \in [0,1]^\mathfrak{m}$ and $i \in [\mathfrak{m}]$ we let $F_i(x) = G(x_j, x_k)$, where $g_i = (G, j, k)$. Indeed, any fixed point of F corresponds to an assignment that exactly satisfies the gate constraints. In particular, note that all the gate-types we consider can be exactly computed using the usual operations allowed in FIXP, namely $+, \times$, max and rational constants. Furthermore, it is easy to see that this trivially yields an SL-reduction [Etessami and Yannakakis, 2010].

In order to prove FIXP-hardness we will show that our very restricted set of gates is actually enough to simulate various more complex gates. Deligkas et al. [2021, Section 7.2], using a special Brouwer function for the FIXP-complete problem 3-Nash given by Etessami and Yannakakis [2010], proved that computing fixed points of very restricted algebraic circuits is already FIXP-hard. In more detail, they consider functions $F : [0,1]^n \to [0,1]^n$ computed by circuits with a restricted set of gates and such that every gate always has value in $[0,1]$, for any input $x \in [0,1]^n$ to the circuit. Because of this property we can use our gates that truncate to $[0,1]$ without changing any of the computations.

In more detail, they allow the following gates: G_ζ (constant $\zeta \in \mathbb{Q} \cap [0,1]$), G_+, G_- (subtraction truncated to $[0,1]$), $G_\times$, $G_{\times 2}^{[0,1]}$, $G_{\max}$ and $G_{\min}$. We show below that we can simulate all of these gates, using only the gates G_{1-}, $G_{\times 2}$ and $G_\times$ (or alternatively, G_{1-}, G_+ and $G_{(\cdot)^2}$). In particular, $G_{\times 2}^{[0,1]}$ is a restricted gate $G_{\times 2}$ that only works on inputs in $[0,1/2]$. Since our $G_{\times 2}$ gate has the same behavior as that gate for such inputs, it is correctly simulated.

Finally, we simply use copy gates $G_=$ to enforce the fixed point constraint, namely that the ith input to F be equal to its ith output. It is easy to see that this construction yields a polynomial-time reduction, and that it is in fact an SL-reduction [Etessami and Yannakakis, 2010], since we only need to extract the values assigned to the input gates in order to obtain a fixed point of F. In the remainder of this proof, we show how all the required gates can be simulated using our restricted set of gates G_{1-}, $G_{\times 2}$ and $G_\times$.

$G_=$: Copy. In order to copy the value of some gate g_1, we use the complement gate G_{1-} twice. Namely, we first introduce a gate g_2 of type G_{1-} with input g_1, and then another gate g_3 of type G_{1-} with input g_2. Clearly it holds that $\mathbf{v}[g_3] = 1 - \mathbf{v}[g_2] = 1 - (1 - \mathbf{v}[g_1]) = \mathbf{v}[g_1]$.

$G_{1/2}$: Constant 1/2. In order to obtain a gate that has value $1/2$, we create a small cycle. We introduce two gates g_1 and g_2. The gate g_1 is of type $G_=$ with input g_2, and the gate g_2 is of type G_{1-} with input g_1. It follows that $\mathbf{v}[g_1]$ satisfies $\mathbf{v}[g_1] = 1 - \mathbf{v}[g_1]$, which implies $\mathbf{v}[g_1] = 1/2$. Note that together with the $G_\times$ gate we can now also perform multiplication by $1/2$, denoted by $G_{\times 1/2}$.

G_-: Subtraction. In the proof of Lemma 5.5, we show how to construct a subtraction gate given access only to G_{1-}, $G_{\times 2}$ and a special gate G_ϕ, where $\phi : [0,1]^2 \to [0,1]$, $(x,y) \mapsto (x+1)(y+1)/4$. Thus, to obtain the subtraction gate, it is enough for us here to construct a gate G_ϕ. Since we have access to $G_\times$, it suffices to construct a gate that implements the function $x \mapsto (x+1)/2$. Let g_1 be the input gate. We introduce a gate g_2 of type G_{1-} with input g_1, a gate g_3 of type $G_{\times 1/2}$ with input g_2, and finally a gate g_4 of type G_{1-} with input g_3. It follows that $\mathbf{v}[g_4] = 1 - \mathbf{v}[g_3] = 1 - \mathbf{v}[g_2]/2 = 1 - (1 - \mathbf{v}[g_1])/2 = (\mathbf{v}[g_1] + 1)/2$, as desired.

G_+: Addition. Addition can easily be obtained from subtraction by using the following equality for all $x, y \in [0,1]$

$$\mathrm{T}(x+y) = 1 - \mathrm{T}((1-x) - y) = G_{1-}(G_-(G_{1-}(x), y)).$$

$G_{\max}$, $G_{\min}$: Maximum and Minimum. The function $(x,y) \mapsto \max\{x,y\}$ can easily be simulated with existing gates by noting that

$$\max\{x,y\} = \mathrm{T}(x + \mathrm{T}(y - x)) = G_+(x, G_-(y,x)).$$

Then, $(x,y) \mapsto \min\{x,y\}$ can simply be obtained by $\min\{x,y\} = 1 - \max\{1-x, 1-y\}$.

$G_{\times k}$: Multiplication by integer k. Let k be an integer that is given in binary representation, i.e., $k = \sum_{i=0}^{\ell} a_i 2^i$, where $a_i \in \{0,1\}$. Our goal is to construct a gate that computes $x \mapsto \mathrm{T}(k \cdot x)$. Using the $G_{\times 2}$ gate we can compute $\mathrm{T}(2^i \cdot x)$ for $i = 0, 1, \ldots, \ell$. This requires ℓ separate $G_{\times 2}$ gates. Then, we use the addition gate to compute

$$\mathrm{T}\left(\sum_{i:\, a_i = 1} \mathrm{T}(2^i \cdot x)\right) = \mathrm{T}\left(\sum_{i=0}^{\ell} a_i 2^i x\right) = \mathrm{T}(k \cdot x).$$

This uses at most ℓ separate G_+ gates. Thus, overall we use a number of gates that is polynomial in the representation length of k.

G_ζ: Constant $\zeta \in [0,1] \cap \mathbb{Q}$. If $\zeta = 1$, we can simply do $G_{\times 2}(G_{1/2}) = 1$. If $\zeta = 0$, we can do $G_{1-}(G_{\times 2}(G_{1/2})) = 0$. Now assume that $\zeta \in (0,1)$. Write $\zeta = c/d$ where c and d are positive integers, $c \geq 1$, $c < d$, $d \geq 2$. Clearly, if we can construct the constant $1/d$, then we can use a $G_{\times k}$ gate with $k = c$ to obtain ζ. In order to construct $1/d$, we use a small cycle. We introduce two gates g_1 and g_2. The gate g_1 is of type $G_{\times k}$ with $k = d - 1$ and with input g_2. The gate g_2 is of type G_{1-} with input g_1. Thus, it

holds that $\mathbf{v}[g_2] = 1 - \mathbf{v}[g_1] = 1 - \mathrm{T}((d-1)\cdot\mathbf{v}[g_2])$. It is easy to check that the only solution of this equation is $\mathbf{v}[g_2] = 1/d$.

Finally, let us show that the set of gate-types G_{1-}, G_+ and $G_{(\cdot)^2}$ also suffices to simulate all the gates above, by showing that they can simulate G_{1-}, $G_{\times 2}$ and $G_\times$. As before, G_{1-} can be used to create $G_=$. Then, G_+ and $G_=$ can be used to obtain $G_{\times 2}$. Thus, it remains to simulate $G_\times$.

Note that G_- can be obtained by $\mathrm{T}(x-y) = 1 - (\mathrm{T}((1-x)+y))$. Furthermore, we can construct $G_{\times 1/2}$ on input gate g_1 as follows. We introduce two gates g_2 and g_3. The gate g_2 is of type G_- and has inputs g_1 and g_3. The gate g_3 is of type $G_=$ with input g_2. It follows that $\mathbf{v}[g_3] = \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_3])$, which has the only solution $\mathbf{v}[g_3] = \mathbf{v}[g_1]/2$.

In order to simulate $G_\times$, note that

$$\left(\frac{x}{2} + \frac{y}{2}\right)^2 = (x/2)^2 + (y/2)^2 + xy/2.$$

We can easily compute $x/2 + y/2$ and then square using $G_{(\cdot)^2}$. Similarly, we can also compute $(x/2)^2 + (y/2)^2$. By using G_-, we then obtain $xy/2$, and thus xy after using a $G_{\times 2}$ gate.

E.3 Proof of Lemma 5.5

In order to prove that the problem remains hard with $\mathcal{G} = \{G_{\times 2}, G_{1-}, G_\phi\}$, we will show that other gate-types can be simulated using only these three gate-types. Let $\varepsilon \in [0, 1/14]$ and assume that we have access to gates of type $G_{\times 2}$, G_{1-} and G_ϕ.

G_1: Constant 1. In order to create a constant 1 we use the fact that for any $x, y \in [0,1]$

$$\mathrm{T}\left(2^3 \cdot \phi(x,y)\right) = \mathrm{T}\left(2^3(x+1)(y+1)/4\right) \geq \mathrm{T}(2) = 1.$$

In more detail, we use a gate g_1 of type G_ϕ (with arbitrary inputs), then a gate g_2 of type $G_{\times 2}$ with input g_1, another gate g_3 of type $G_{\times 2}$ with input g_2, and finally another gate g_4 of type $G_{\times 2}$ with input g_3. We have that $\mathbf{v}[g_1] \geq 1/4 - \varepsilon$, $\mathbf{v}[g_2] \geq \mathrm{T}(2\cdot\mathbf{v}[g_1]) - \varepsilon \geq 1/2 - 3\varepsilon$, $\mathbf{v}[g_3] \geq \mathrm{T}(2\cdot\mathbf{v}[g_2]) - \varepsilon \geq 1 - 7\varepsilon$, and $\mathbf{v}[g_4] \geq \mathrm{T}(2\cdot\mathbf{v}[g_3]) - \varepsilon \geq 1 - \varepsilon$, since $\varepsilon \leq 1/14$. Thus, we can construct a gate that has the value $1 \pm \varepsilon$.

$G_{/2}$: Division by 2. In order to divide the value of some gate g_1 by 2, we use the fact that

$$1 - \phi(1 - \mathbf{v}[g_1], 1) = 1 - (2 - \mathbf{v}[g_1])(1+1)/4 = \mathbf{v}[g_1]/2.$$

In more detail, we use a gate g_2 of type G_{1-} with input g_1, then we use a gate g_3 of type G_ϕ with inputs g_2 and a constant $1 \pm \varepsilon$, and finally we use a gate g_4 of type G_{1-} with input g_3. It holds that $\mathbf{v}[g_2] = 1 - \mathbf{v}[g_1] \pm \varepsilon$, $\mathbf{v}[g_3] = \phi(\mathbf{v}[g_2], 1 \pm \varepsilon) \pm \varepsilon = 1 - \mathbf{v}[g_1]/2 \pm 2\varepsilon$, and $\mathbf{v}[g_4] = 1 - \mathbf{v}[g_3] \pm \varepsilon = \mathbf{v}[g_1]/2 \pm 3\varepsilon$. Thus, we can construct a gate that performs division by 2 with error at most 3ε.

$G_=$: Copy. It is easy to see that using two gates of type G_{1-}, one after the other, copies the original value with error at most 2ε.

G_{inv}: Inverse. We now show how to construct the gate G_{inv}, which computes the function $x \mapsto -1 + 4/(2+x)$, and will be very useful to construct the subtraction gate below. The construction of G_{inv} uses a cycle. Let g_1 be the input gate. We first use a gate g_2 of type G_{1-} with input g_1, then we use a gate g_3 of type G_ϕ with input g_2 and g_4, and finally we let gate g_4 be of type $G_=$ with input g_3. We have that $\mathbf{v}[g_2] = 1 - \mathbf{v}[g_1] \pm \varepsilon$, $\mathbf{v}[g_3] = \phi(\mathbf{v}[g_2], \mathbf{v}[g_4]) \pm \varepsilon$, and $\mathbf{v}[g_4] = \mathbf{v}[g_3] \pm 2\varepsilon$. It follows that $\mathbf{v}[g_4]$ must satisfy the equation

$$\mathbf{v}[g_4] = \phi(\mathbf{v}[g_2], \mathbf{v}[g_4]) \pm 3\varepsilon = (\mathbf{v}[g_2]+1)(\mathbf{v}[g_4]+1)/4 \pm 3\varepsilon$$

which implies that

$$\mathbf{v}[g_4] = \frac{1 + \mathbf{v}[g_2]}{3 - \mathbf{v}[g_2]} \pm 6\varepsilon.$$

As a result, we obtain that

$$\mathbf{v}[g_4] = \frac{2 - \mathbf{v}[g_1]}{2 + \mathbf{v}[g_1]} \pm 8\varepsilon = -1 + \frac{4}{2 + \mathbf{v}[g_1]} \pm 8\varepsilon$$

i.e., we can compute the function with error at most 8ε.

G_-: Subtraction. Given gates g_1 and g_2, we want to obtain $\mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2])$. To achieve this, we first use the fact that

$$\phi\left(\phi\left(-1 + \frac{4}{2+y}, 1 - x\right), \frac{y}{2}\right) = \phi\left(\frac{2-x}{2+y}, \frac{y}{2}\right) = \frac{1}{2} + \frac{1}{8}(y - x).$$

In more detail, we first use a gate g_3 of type G_{inv} with input g_2, then a gate g_4 of type G_{1-} with input g_1, then a gate g_5 of type G_ϕ with inputs g_3 and g_4, then a gate g_6 of type $G_{/2}$ with input g_2, and finally a gate g_7 of type G_ϕ with inputs g_5 and g_6. We thus obtain that $\mathbf{v}[g_3] = -1+4/(2+\mathbf{v}[g_2]) \pm 8\varepsilon$, $\mathbf{v}[g_4] = 1-\mathbf{v}[g_1] \pm \varepsilon$, and $\mathbf{v}[g_5] = (2 - \mathbf{v}[g_1])(2 + \mathbf{v}[g_2]) \pm 7\varepsilon$. Furthermore, it holds that $\mathbf{v}[g_6] = \mathbf{v}[g_2]/2 \pm 3\varepsilon$, and thus $\mathbf{v}[g_7] = 1/2 + (\mathbf{v}[g_2] - \mathbf{v}[g_1])/8 \pm 11\varepsilon$.

Next, we can obtain the subtraction operation from this by noting that

$$4\left(1 - \mathrm{T}\left(2\left(\frac{1}{2} + \frac{1}{8}(y - x)\right)\right)\right) = 4\left(1 - \left(1 - \frac{1}{4}\mathrm{T}(x - y)\right)\right) = 4\frac{\mathrm{T}(x-y)}{4} = \mathrm{T}(x - y).$$

This is implemented by using a gate g_8 of type $G_{\times 2}$ with input g_7, then a gate g_9 of type G_{1-} with input g_8, then a gate g_{10} of type $G_{\times 2}$ with input g_9, and finally another gate g_{11} of type $G_{\times 2}$ with input g_{10}. It holds that

$$\mathbf{v}[g_8] = \mathrm{T}(2 \cdot \mathbf{v}[g_7]) \pm \varepsilon = 1 - \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2])/4 \pm 23\varepsilon.$$

As a result, it then holds that $\mathbf{v}[g_9] = \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2])/4 \pm 24\varepsilon$, $\mathbf{v}[g10] = \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2])/2 \pm 49\varepsilon$, and finally $\mathbf{v}[g_{11}] = \mathrm{T}(\mathbf{v}[g_1] - \mathbf{v}[g_2]) \pm 99\varepsilon$. Thus, we can compute subtraction with error at most 99ε.

$G_\times$: Multiplication. Given gates g_1 and g_2, we want to obtain $\mathbf{v}[g_1] \cdot \mathbf{v}[g_2]$. We only perform the construction for the case $\varepsilon = 0$, since we only need this gate for the FIXP-hardness. Note that we can multiply by 4 using two consecutive $G_{\times 2}$ gates. Similarly, we can divide by 4 using two consecutive $G_{/2}$ gadgets. To perform multiplication, we use the fact that

$$\phi(x, y) - \frac{1}{4} - \frac{x}{4} - \frac{y}{4} = \frac{xy}{4}.$$

In more detail, we first use a gate g_3 of type G_ϕ with input g_1 and g_2, then a gate g_4 of type $G_{/4}$ with input the constant 1, then a gate g_5 of type G_- with inputs g_3 and g_4, then a gate g_6 of type $G_{/4}$ with input g_1, then a gate g_7 of type G_- with inputs g_5 and g_6, then a gate g_8 of type $G_{/4}$ with input g_2, then a gate g_9 of type G_- with inputs g_7 and g_8, and finally a gate g_{10} of type $G_{\times 4}$ with input g_9. We have that

$$\mathbf{v}[g_3] = \phi(\mathbf{v}[g_1], \mathbf{v}[g_2]) = (\mathbf{v}[g_1] + \mathbf{v}[g_2] + \mathbf{v}[g_1] \cdot \mathbf{v}[g_2] + 1)/4.$$

Then we obtain that $\mathbf{v}[g_5] = (\mathbf{v}[g_1] + \mathbf{v}[g_2] + \mathbf{v}[g_1] \cdot \mathbf{v}[g_2])/4$, $\mathbf{v}[g_7] = (\mathbf{v}[g_2] + \mathbf{v}[g_1] \cdot \mathbf{v}[g_2])/4$, $\mathbf{v}[g_9] = \mathbf{v}[g_1] \cdot \mathbf{v}[g_2]/4$, and finally $\mathbf{v}[g_{10}] = \mathbf{v}[g_1] \cdot \mathbf{v}[g_2]$. Thus, we can perform exact multiplication when $\varepsilon = 0$.

Hardness. We have shown that we can simulate gates G_1 and G_- with error at most 99ε. Thus, by Proposition 5.3, the PPAD-hardness of our restricted version follows. For the case $\varepsilon = 0$, we have shown that we can exactly simulate gates $G_{\times 2}$, G_{1-} and $G_\times$. As a result, by Proposition 5.4, the exact version of our restricted version is FIXP-hard.