First-Price Procurement Auctions

Nicole Immorlica* David R. Karger† Evdokia Nikolova‡ Rahul Sami§

Abstract

We study markets in which an auctioneer wishes to assemble a team of agents to accomplish some task. These agents offer fixed services that incur some privately known cost. The auctioneer must select a team, or *feasible set* of agents, that is capable of performing the task. To this end, he designs a *procurement auction* in which he solicits bids from the agents and then selects some feasible set of agents (the *winners*) to perform the task at hand and pays them according to the rules of the auction. One possible mechanism for procurement auctions is the Vickrey-Clark-Groves (VCG) mechanism. However, a drawback of this mechanism is the potentially large overpayment: the payment of the auctioneer may greatly exceed the true cost of even the second-cheapest feasible set.

We thus propose and analyze variants on *first-price* auctions, or auctions in which the team with the lowest bid is selected and pays their bid. These auctions are not truthful; instead, we motivate analyzing their properties in a *strong ϵ-Nash equilibrium*. We show that in general procurement settings, strong ϵ-Nash exist, and the feasible set of agents selected in any strong ϵ-Nash equilibrium is approximately efficient. For path and flow auctions, we bound the total payment to the winning agents by relating it to the true cost of routing one additional unit of demand (assuming all edges have unit capacity). Finally, we study the setting in which the demand of the auctioneer is not known, but rather the auctioneer and bidders share a *common prior belief* regarding the amount of demand. For this model, we design a first-price mechanism involving *two-parameter* bids and derive a bound on the payments of this mechanism similar to that of the known demand case.

Keywords: First-price auctions, procurement auctions, VCG, path auctions, flow auctions

1 Introduction

In this paper, we study markets in which an auctioneer wishes to assemble a team of agents to accomplish some task. These agents offer fixed services that incur some privately known cost. The auctioneer must select a team, or *feasible set* of agents, the combination of which is capable of performing the task. To this end, he designs a *procurement auction* in which he solicits bids from the agents and then selects some feasible set of agents (the *winners*) to perform the task at hand and pays them according to the rules of the auction.

Path and *flow* auctions are important special cases of procurement auctions. In path auctions, the auctioneer seeks to buy a path of edges of lowest price between a specified

*Northwestern University, Chicago, IL, USA. email: {nickle}@eecs.northwestern.edu

†Massachusetts Institute of Technology, Cambridge, MA, USA. email:{karger}@theory.csail.mit.edu

‡Massachusetts Institute of Technology, Cambridge, MA, USA. email:{nikolova}@mit.edu

§University of Michigan, Ann Arbor, MI, USA. email: {rsami}@umich.edu

source and destination in a network. Sellers (network edges) have a privately known cost for transmitting traffic, and bid to attract traffic. Path auctions arise naturally in network routing—for example, an Internet Service Provider (ISP) might use a procurement auction to select autonomous systems (ASs) to route his demand. Flow auctions are a generalization of path auctions in which the demand of the auctioneer might exceed the capacity of any single source-destination path. In this case, the auctioneer must buy a set of edges capable of routing his demand.

One plausible mechanism for procurement auctions, proposed for use in path auctions by Nisan and Ronen [32], Hershberger and Suri [18], and Feigenbaum *et al.* [14], is the Vickrey-Clark-Groves (VCG) mechanism [7, 17, 35]. Roughly speaking, the VCG mechanism pays each winning agent the highest bid with which it could still have won, all other bids being unchanged. The utility of an agent is quasi-linear, and so the VCG mechanism is truthful (agents bid their true cost) and efficient (a feasible set of minimum total true cost is selected). However, as observed by Archer and Tardos [1, 2], even in the special case of path auctions, the VCG mechanism (and, in fact, all *min function* mechanisms) can lead to the auctioneer paying far more than the true cost of completing the task at hand. In fact, the payment of the auctioneer may even greatly exceed the true cost of the second-cheapest feasible set. Elkind, Sahai, and Steiglitz [12] generalized the result of Archer and Tardos [1] to prove that *all* truthful mechanisms have high overpayments in general.

We are interested in reining in the cost to the auctioneer. There are two general approaches to this problem. One approach tries to characterize procurement settings in which truthful mechanisms have small overpayments. Talwar [34] and Garg et al. [16] consider restricting the setting by imposing a structure on the collection of feasible sets of agents. Mihail, Papadimitriou, and Saberi [30] show that in a *random graph*, the expected payment of a VCG mechanism for a shortest path is small. Karger and Nikolova [25] give a tighter VCG overpayment bound for Erdos-Renyi random graphs and provide empirical results of small overpayment for random graphs with power-law degree distributions. Feigenbaum *et al.* [14] measure the average overpayment of the VCG mechanism for shortest path auctions in the Internet's autonomous systems (ASs) graph and conclude that it is relatively small. Beyond VCG, one could look for truthful path mechanisms that are more *frugal*, namely that have smaller payments than the corresponding VCG payments [2, 12, 38]. Frugal mechanisms have also been proposed and analyzed for different set systems other than paths, *e.g.,* [34, 26, 11, 4, 27, 13]. The truthful concept could be further refined to be false-name-proof, namely when agents own more than one edge, they have no incentive to misrepresent which edges they own. For path auctions, this direction has been pursued by Du *et al.* [10] and Iwasaki *et al.* [20].

A second approach is to consider alternative solution concepts. Garg *et al.* [16] propose an *ascending price auction* format for procurement auctions that can perform well in settings of incomplete information. For the special case of path auctions, Elkind, Sahai, and Steiglitz [12] present and analyze an optimal *Bayesian-Nash* mechanism. Czumaj and Ronen [9] propose a mechanism that combines dominant and non-dominant strategy mechanisms and has small overpayments under certain assumptions. However they show that it has an arbitrary ratio between the payment of different equilibria and say that overall, "finding a natural and tractable measure of [non-dominant strategy] protocols seems challenging and important." Huang *et al.* [19] propose a Nash Implementation mechanism for path auctions, based on the randomized first-price mechanisms in this paper. They show small overpayments under this mechanism similarly to our analysis here. Small overpayments under certain dominant [29] or non-dominant strategy mechanisms are also shown to result in specific settings of wireless routing (*e.g.,* [36, 37]). A different line of research focuses on bounding the cost of the cheapest Nash equilibria in general procurement settings [6, 5].

In this paper, we follow the second of these approaches. We propose and analyze variants on *first-price* auctions, or auctions in which the team with the lowest bid is selected and paid their bid. First-price auctions are a natural class of auctions often implemented in practice. Therefore, it is interesting to ask if first-price auctions or their variants can reduce the payment of the auctioneer. These auctions are not truthful; instead, we motivate analyzing their properties in a *strong ϵ-Nash equilibrium* (see Definition 2). We show that in general procurement settings, strong ϵ-Nash exist, and the feasible set of agents selected in any strong ϵ-Nash equilibrium is approximately efficient. For path and flow auctions, we then bound the total payment to the winning agents by relating it to the true cost of routing one additional unit of demand (assuming all edges have unit capacity). Finally, we study the setting in which the demand of the auctioneer is not known, but rather the auctioneer and bidders share a *common prior belief* regarding the amount of demand. In other words, there is a publicly known *distribution* of possible demands. For this model, we design a first-price mechanism involving *two-parameter* bids and derive a bound on the payments of this mechanism similar to that of the known demand case.

Our work is also related to the literature on strong Nash and strong ϵ-Nash implementation of the core. In particular, the deterministic first-price procurement auction we consider is similar to the game introduced by Young [39] in the context of cost-sharing. For the random demand path auction introduced in section 5.2.3, we use techniques based on Curiel [8] to show the existence of the core. We also note that Kalai *et al.* [24] presented a strong Nash implementation of the core of any cooperative game. We could have used this implementation in place of the 2-parameter auction in Section 5.2.3; however, the method in [24] is more complex and communication-intensive, and in our case it would essentially require each bidder to report an entire flow.

We structure the presentation of our results as follows: In Section 2, we formalize the setting of procurement auctions and define the path and flow settings which we study later. In Section 3, we motivate the selection of strong ϵ-Nash equilibria as a solution concept for first-price auctions. In Section 4, we show that first-price auctions are approximately efficient in the general procurement setting. Finally, in Section 5, we show how to bound the payment of first-price auctions and their variants in the special case of path and flow settings.

2 Setting

Consider a *procurement setting* in which an auctioneer wishes to hire a team of agents to accomplish a particular task. There is a set U of n agents. Each agent is capable of performing a fixed service. In performing this service, an agent incurs a privately known cost $c_i \in \mathbb{R}^+ \cup \{0\}$. Some subsets of services can be combined to accomplish the auctioneer's task. We call a subset $S \subseteq U$ of agents a *feasible set* if their combined services can accomplish the task. The collection of feasible sets is denoted by $\mathcal{S}$. The collection $\mathcal{S}$ could be publicly known to the auctioneer and all agents, or, more generally, they could share a *common prior* (a publicly known probability distribution over the collection of subsets of U) about $\mathcal{S}$.

A special case of the procurement setting is the *path* or *flow setting.* In this setting, there is a graph G. Each edge (u, v) is an agent capable of sending *one* unit of flow from u to v

at a privately known cost $c_i \in \mathbb{R}^+ \cup \{0\}$. The auctioneer wants to route k units of demand between a known source node s and destination node t (in a *path* setting $k = 1$). Hence the collection $\mathcal{S}$ of feasible sets is the collection of all subgraphs in G that contain a k-flow from s to t. We assume that the structure of the graph G is public knowledge. The demand k could be publicly known to the auctioneer and all edges (the *known demand case*), or it might be drawn from a publicly known probability distribution (the *unknown demand case*).

The unknown demand case is modeled as follows: The demand can take any integral value in the range $[1, r]$, where r is a positive integer. Further, there is a known prior distribution on the demand values; say that the demand is k with probability p_k, for $k = 1, 2 \ldots, r$. We assume for simplicity that $p_k > 0$ for all k; our results easily extend to a situation in which $p_k = 0$ for some values of $k \in \{1, \ldots, r\}$.

In a *procurement auction* (similarly a *path auction* or *flow auction*), the auctioneer selects a feasible set by running an auction. He solicits from each agent a *bid* $b_i \in \mathbb{R}^+ \cup \{0\}$ which is supposed to represent the agent's true cost c_i. He then selects some feasible set S of agents and pays each agent $i \in S$ an amount $payment_i \in \mathbb{R}^+ \cup \{0\}$ and all other agents 0. The set S is called the *winning set*. Each agent $i \in S$ is a *winner*, and all other agents are *losers*. An agent's utility for the outcome is $payment_i - c_i x_i$, where $\mathbf{x}$ is the characteristic vector of S (that is, $x_i = 1$ if $i \in S$ and 0 otherwise).

We will focus on *first-price* auctions. In a first-price auction, the payment of every winner equals his bid. The auctioneer is restricted to select a *minimum price* feasible set S, or one which minimizes $\sum_{i \in S} b_i$. His only flexibility is in the definition of a *tie-breaking rule*, or method to select from among the collection of minimum price feasible sets. Thus, in specifying a first-price auction, we only need to specify a tie-breaking rule. We also consider variants of first-price auctions in which the minimum price feasible set is almost always selected and the winners are paid a quantity close to their bid.

To avoid confusion between the true costs and the prices of sets, we will adopt the following terminology: the *cost* of a set S is $\sum_{i \in S} c_i$, sometimes written $c(S)$. Similarly, the *price* of a set S is $\sum_{i \in S} b_i$, sometimes written $b(S)$. Additional notation will be introduced for the path and flow settings in Section 5.

3 Solution Concepts

First-price auctions are clearly not truthful. This raises the question of how we expect agents to bid. We want to retain the property that agents can see each others' bids, so that the bidding could be performed through posted prices. Thus, mixed-strategy equilibria are not very meaningful in our setting. Instead, we look for a pure strategy equilibrium solution concept which *always exists* and is *arguably reasonable* in that agents can be expected to reach that equilibrium. This section motivates the selection of *strong ϵ-Nash equilibria* (see Definition 2) as that solution concept through a series of examples. First we note that not every first-price procurement auction has a Nash equilibrium (Example 3.1), and those that do are impractical (Example 3.2). Both of these examples heavily rely on the continuity of the bid and payment space. In reality, bids and payments are restricted to a discrete space as they should be some multiple of a unit of money, like cents, for example. Thus it is simply not possible for agents to arbitrarily improve their payoffs, and so we suggest studying *ϵ-Nash equilibrium* (see Definition 1) where an agent deviates only if it improves his payoff by at least ϵ. Unfortunately, Example 3.3 shows that the overpayments in such an equilibrium can be quite high. In this example, however, if certain subsets of agents could arrange to jointly reduce their bids, all of them would benefit. This leads us to study *strong ϵ-Nash equilibria* (see Definition 2), or ϵ-Nash equilibria that are robust to such manipulations. As proven in Theorem 3.1, strong ϵ-Nash equilibria exist in all deterministic first-price procurement auctions (but may fail to exist in randomized ones as evidenced by Example 3.4). It remains to be seen if one can devise a bidding protocol that helps agents converge to a strong ϵ-Nash equilibrium.

3.1 Nash Equilibria

The most natural solution concept is that of a Nash equilibrium. Unfortunately, as the following example shows, not every first-price auction has a Nash equilibrium. A similar observation in a more general setting was made by Jackson *et al.* [22].

Example 3.1. *Suppose there are two agents A and B, either of whom forms a feasible set (that is, $\mathcal{S} = \{\{A\}, \{B\}\}$). Consider any auction in which ties are broken by selecting agent B with probability p, $0 < p \leq 1$, independent of the bid values. Now suppose the costs of the agents are $c_A = 1$ and $c_B = 2$, and so in case of a tie the auction selects the higher-cost agent with positive probability.*

Suppose agent A bids x and B bids y. If $x \geq y$, then the expected payment of agent A is at most $(1-p)y$. As B has positive probability of winning, $y \geq c_B = 2$, and so the bid $y - \epsilon$ for $\epsilon < \min(yp, \frac{1}{2})$ is a better bid than x for agent A. If $x < y$, then the payment to A is x and so $x + (y-x)/2$ is a better strategy than x for A.

This example relies on the assumption that the tie-breaking rule is not a function of the bid values (otherwise we would have been unable to assume that the auction selects the higher-cost agent with positive probability). In fact, for a carefully chosen tie-breaking rule which is a function of the bid values, we can design first price auctions with pure strategy Nash equilibria, as the following example shows.

Example 3.2. *For ease of exposition, suppose all subsets 2^U of the set of agents U are feasible and index the subsets so $2^U = \{B_1, \ldots, B_{2^n}\}$. Partition the real numbers into 2^n subsets $S_1, \ldots, S_{2^n}$ such that each subset is dense in the reals. Let p be the price of the minimum price set and suppose $p \in S_k$. If B_k has price p, choose B_k. Otherwise, choose randomly among the collection of minimum price sets.*

We can construct a pure strategy Nash equilibrium for this tie-breaking rule as follows. If the minimum cost set is not unique, then it is a Nash equilibrium for all agents to bid $b_i = c_i$, their true cost. Otherwise, let B_1 be the minimum cost set and B_2 be the next cheapest set (in terms of true costs). Find a $p \in S_1 \cap [c(B_1), c(B_2))$ (thus B_1 wins in the case of a tie at price p). Consider a set of bids $\mathbf{b}$ such that $b(B_1) = b(B_2) = p$, $b_i \geq c_i$ for $i \in B_1$, $b_i \leq c_i$ for $i \in B_2$, and $b_i = c_i$ for $i \notin B_1 \cup B_2$. Then $\mathbf{b}$ is a Nash equilibrium.

However, this auction is arguably impractical as are the deviations discussed in the last example because they both assume that the bids and payments can be any real number. Yet, in many problems, payments are discrete, so it is simply not possible for agents to improve their utilities by arbitrarily small amounts. This motivates us to use the solution concept of *ϵ-Nash equilibrium.*

Remark 3.1. *The results in this paper can be proved using tie-breaking rules such as that in Example 3.2 or using ϵ-equilibria concepts presented below. However, for clarity of presentation, we present our results in terms of ϵ-equilibria.*

3.2 ϵ-Nash Equilibria

In an ϵ-Nash equilibrium, we assume agents are indifferent to deviations that improve their payoff by a small amount.

Definition 1. *An* ϵ-Nash equilibrium *is a set of strategies, one for each agent, such that no agent can unilaterally deviate in a way that improves his payoff by at least* ϵ.

Unfortunately, there is a drawback to the ϵ-Nash solution concept as well. As the following example shows, when the winning set contains many agents, it may have a *price* higher than the *cost* of the best competing set.

Example 3.3. *Consider any first-price auction. Suppose there are four agents, A, B, C, and D with costs* 1, 2, 2, *and* 6 *respectively, and the collection $\mathcal{S}$ of feasible sets is* $\{\{A\},\{B,C\},\{D\}\}$. *Then it is an ϵ-Nash equilibrium for agent A to bid $6-\epsilon$, and the rest to bid 6. In this case, the price to the auctioneer for the winning set $\{A\}$ is $6-\epsilon$ which is higher than the cost, 2, of the best competing set $\{B,C\}$.*

This defeats our goal of reducing customer overpayment. We might argue that this solution would not be sustained in practice, since the agents in the second lowest-cost set are likely to each reduce their price. This leads us to explore the concept of *strong ϵ-Nash equilibria.*

3.3 Strong ϵ-Nash Equilibria

Strong ϵ-Nash equilibria, first introduced by Aumann [3] and used by Young [39], require that there is no *group* of agents who can deviate in a way that improves the payoff of each member by at least ϵ.

Definition 2. *A* strong ϵ-Nash equilibrium *is a set of strategies, one for each agent, such that no group of agents (called a* coalition*) can deviate in a way that improves the payoff of each member by at least* ϵ.

This definition captures the notion that agents might collude to win the auction if it is beneficial for each of them. For example, the bid vector in Example 3.3 is not a strong ϵ-Nash equilibrium as agents B and C could collude and bid $3-\epsilon$, thus improving each of their payoffs by at least ϵ (assuming $\epsilon < \frac{1}{2}$).

Strong ϵ-Nash equilibria have several advantages over Nash and ϵ-Nash equilibria. First, although *randomized* first-price auctions may fail to have strong ϵ-Nash equilibria (see Example 3.4), Theorem 3.1 shows that every *deterministic* first-price auction has a strong ϵ-Nash equilibrium. Second, as demonstrated by Lemma 3.1, in a strong ϵ-Nash equilibrium of a deterministic first-price auction, we can bound the bids of the winning agents by the true costs of the losing agents, furthering our goal of reducing payments and allowing us to prove that the winning set is approximately efficient (see Section 4). The rest of this section contains proofs of Theorem 3.1 and Lemma 3.1 and Example 3.4.

First, we show that any first-price auction with a *deterministic* tie-breaking rule has a strong ϵ-Nash equilibrium. Our proof is constructive. We consider the minimum cost feasible set and fix the bids of all items outside this set to be equal to their true cost. For the items in this set, we adjust their bids so that the price of the set is just less than the cost of the second-lowest cost feasible set.

Theorem 3.1. *Any first price auction with a deterministic tie-breaking rule has a strong ϵ-Nash equilibrium.*

Proof. Our proof is constructive. Let c_i be the cost of agent i, $\mathcal{S}$ be the collection of feasible sets, and S^* be the minimum cost feasible set selected by the auction under bid vector $\mathbf{c}$. Define a variable x_i for each $i \in S^*$ and consider the following linear program (LP for short):

$$\begin{array}{ll} \text{maximize} & \sum_{i \in S^*} x_i \\ \text{subject to} & \forall\, S \in \mathcal{S}: \ \sum_{i \in S^* - S} x_i \le \sum_{i \in S - S^*} c_i \\ & \forall\, i \in S^*: \ x_i \ge c_i \end{array}$$

The strong ϵ-Nash equilibrium that we construct will be a slightly modified optimal solution to this LP. The first constraint guarantees that S^* will be a minimum price set in this equilibrium, and the second that every agent has non-negative utility in this equilibrium. By setting $x_i = c_i$ for all i, we see that the LP is feasible.

Let x_i^* be an optimum solution of the LP, and define bid vector $\mathbf{b}$ where $b_i = \max\{c_i, x_i^* - \epsilon/(2n)\}$ for $i \in S^*$ and $b_i = c_i$ for all other i. Notice that our minimum cost set S^* is also a minimum price set with respect to bids $\mathbf{b}$.

We prove that $\mathbf{b}$ is a strong ϵ-Nash equilibrium. Note that only agents who are guaranteed winners (that is, agents in every minimum price set) are submitting a bid other than their true cost. For agents outside S^*, this is evident from the definition of $\mathbf{b}$. Consider an agent i in S^* that is not in every minimum price set, and let S be a minimum price set that does not contain i. Corresponding to this S is an inequality of type 1. This inequality together with those of type 2 for all $j \in S^* - S$ imply that $x_i = c_i$ and so $b_i = c_i$. Thus the bidders in a successful coalition can only increase their bids.

Let T be a successful coalition and $\mathbf{b}'$ be the bid vector when T deviates (so $b'_i = b_i$ for all $i \notin T$). Recall the notation $b(S) = \sum_{i \in S} b_i$. Then

$$b'(S^*) = b(S^*) + \sum_{i \in T \cap S^*} (b'_i - b_i). \tag{1}$$

In order for each member of the coalition to benefit by at least ϵ, he must increase his bid by at least ϵ, so

$$\forall\, i \in T,\ b'_i - b_i \ge \epsilon, \tag{2}$$

and T must be a subset of the selected minimum price set S'. Therefore,

$$b'(S') = b(S') + \sum_{i \in T} (b'_i - b_i). \tag{3}$$

Furthermore, as S^* is a minimum price set according to $\mathbf{b}$ and S' is a minimum price set according to $\mathbf{b}'$,

$$b(S^*) \le b(S'), \tag{4}$$

and

$$b'(S') \le b'(S^*). \tag{5}$$

Inequalities 1, 3, 4, and 5 imply

$$\sum_{i \in T \cap S^*} (b'_i - b_i) \geq \sum_{i \in T} (b'_i - b_i). \tag{6}$$

Together with inequality 2, inequality 6 implies $T \subseteq S^*$. As $T \subseteq S'$ as well, inequalities 4 and 5 imply

$$b'(S^*) = b'(S'). \tag{7}$$

Now consider the solution to the LP which sets each variable x_i to agent i's bid in bid vector $\mathbf{b}'$. By inequality 2,

$$\begin{aligned}
\sum_{i \in S^*} x_i &= \sum_{i \in S^*} b'_i \\
&\geq \sum_{i \in S^* - T} b_i + \sum_{i \in T \cap S^*} (b_i + \epsilon) \\
&\geq \sum_{i \in S^* - T} (x_i^* - \frac{\epsilon}{2n}) + \sum_{i \in T \cap S^*} (x_i^* - \frac{\epsilon}{2n} + \epsilon) \\
&\geq \sum_{i \in S^*} x_i^* + \frac{\epsilon}{2}.
\end{aligned}$$

By maximality of $\mathbf{x}^*$, this implies that $\mathbf{x}$ is not feasible. Since each $b'_i \geq c_i$ by construction, $\mathbf{x}$ must violate an inequality of type 1. Letting S be the set in the violating constraint, we see $b'(S) < b'(S^*)$ which, by equality 7, implies $b'(S) < b'(S')$, contradicting the optimality of S'. □

Next, we show that in a strong ϵ-Nash equilibrium, the price of the winning set can be bounded by the cost of losing feasible sets. The intuition for this proof is that if the winning agents are bidding significantly more than the losing agents, the losing agents can undercut the bidding agents and win at a profitable price. One powerful consequence of this definition is that, from the point of view of the total price, it lets us assume without loss of generality that items who are not winning in a strong ϵ-Nash equilibrium are bidding within ϵ of their cost. This notion is formalized in the following lemma.

Lemma 3.1. *Fix a strong ϵ-Nash equilibrium $\mathbf{b}$ and let S be the feasible set that wins with bids $\mathbf{b}$. Let T be any set (not necessarily feasible) such that $T \cap S = \emptyset$ and for all $i \in T$, $b_i > c_i + \epsilon$, where c_i is the true cost of item i. Consider the altered bid vector $\mathbf{b}'$ in which*

$$b'_i = \begin{cases} c_i + \epsilon & \text{for } i \in T, \\ b_i & \text{otherwise.} \end{cases}$$

Let S' be a minimum price feasible set with respect to bids $\mathbf{b}'$. Then $b'(S') = b(S)$.

Proof. Since $T \cap S = \emptyset$, $b(S) = b'(S)$, and so $b'(S') \leq b(S)$. Suppose $b'(S') < b(S)$. This means for all minimum price sets with respect to bids $\mathbf{b}'$, there are items in the set T. Let R' be a minimum price set with respect to bids $\mathbf{b}'$ which minimizes $|R' \cap T|$ (by the previous statement, this minimum is at least one). We will show that the agents in $R' \cap T$ form a coalition when the bids are $\mathbf{b}$, contradicting the assumption that $\mathbf{b}$ was a strong ϵ-Nash equilibrium.

Consider the bid vector $\mathbf{b}''$ constructed from $\mathbf{b}$ in which just the agents in $R' \cap T$ lower their bids to $c_i + \epsilon$:

$$b''_i = \begin{cases} c_i + \epsilon & \text{for } i \in R' \cap T, \\ b_i & \text{otherwise.} \end{cases}$$

We will argue that the agents in $R' \cap T$ benefit by at least ϵ in this deviation. As $T \cap S = \emptyset$, all agents in $R' \cap T$ were losing agents with bid vector $\mathbf{b}$ and so their utility with bids $\mathbf{b}$ was zero. We argue that in bid vector $\mathbf{b}''$ the agents in $R' \cap T$ all win the auction and, therefore, as $b''_i = c_i + \epsilon$ for $i \in R' \cap T$, increase their utility by ϵ. As first-price auctions choose a winning set from among the minimum price feasible sets, we must show that the agents in $R' \cap T$ are contained in *any* minimum price feasible set R'' with respect to bids $\mathbf{b}''$. As $b'_i \leq b''_i$ for all agents i,

$$b''(R'') \geq b'(R'') \geq b'(R') = b''(R') \geq b''(R''),$$

and so all statements hold with equality. Since items in $T - R'$ increased in price from $\mathbf{b}'$ to $\mathbf{b}''$, $b''(R'') = b'(R'')$ implies R'' does not contain any element of $T - R'$. Since $b'(R'') = b'(R')$, R'' is also a minimum price set with respect to bids $\mathbf{b}'$. As R' was chosen to minimize the intersection with T among all minimum price sets, this means R'' must contain $R' \cap T$. Therefore, the agents in $R' \cap T$ are winners in bid vector $\mathbf{b}''$ and so increase their utility by at least ϵ. Furthermore, as $b_i = b''_i$ for all agents outside of $R' \cap T$, the agents in $R' \cap T$ can form a successful coalition in $\mathbf{b}$, contradicting the assumption that $\mathbf{b}$ was a strong ϵ-Nash equilibrium. □

Remark 3.2. *In the proof of Theorem 3.1, we used the determinism of the mechanism in assuming that there was a unique winner for every bid vector. As the following example shows, this assumption was necessary. Strong ϵ-Nash equilibria do not necessarily exist for randomized first-price auctions. Randomized tie-breaking rules pose a problem for the solution concept as a minor adjustment in bid value can drastically affect a bidder's expected utility.*

Example 3.4. *Suppose there are two agents, A and B, either of whom forms a feasible set for the auctioneer (that is, $\mathcal{S} = \{\{A\}, \{B\}\}$). In the case of a tie, assume the auctioneer chooses uniformly at random between the two items. Suppose the cost of each agent is 0. Note for any set of bids $\{b_A, b_B\}$, the agents can form a coalition and each bid $2\max(b_A, b_B) + 2\epsilon$. In this way they both profit by at least ϵ in expectation. Therefore no pure strategy bid vector forms a strong ϵ-Nash equilibrium.*

3.4 Bayesian Nash equilibria

In this paper, we consider full information equilibrium concepts: Nash equilibria and strong Nash equilibria. For the procurement settings we study, an alternative model with Bayesian partial information may be appealing. Here, we observe that the well-known result of Myerson and Satterthwaite [31] immediately limit the possible results in the Bayesian setting. The following example demonstrates that, in the Bayesian setting, it is impossible for any individually rational mechanism to implement the efficient outcome while bounding the total payments by the cost of the best alternative team that has no agents in the winning team.

Example 3.5. *Consider four agents $\{A, B, C, D\}$ with three feasible sets $\{A, B\}$, $\{C\}$ and $\{D\}$. Assume that the costs c_A, c_B, c_C, and c_D are drawn independently from the following* *prior distributions:* c_A *and* c_B *drawn uniformly at random from the interval* $[0, 1]$*, while* c_C *has a certain value of* 1*, and* c_D *has a certain value of* $1 + \delta$*, for an arbitrarily small positive* δ*.*

We identify this scenario with the model of bilateral bargaining with a broker [31]: Let A be a buyer with value $v_A = c_A$ *and B be a seller with value* $v_B = 1 - c_B$*. Then, trade is efficient* iff $v_A \leq v_B \Leftrightarrow c_A + c_B \leq 1$*. This is the same condition under which it is efficient to pick the team* $\{A, B\}$*.*

Next, we identify the ex-post surplus of A and B in the procurement setting with the ex-post surpluses x_A, x_B *in the bargaining model; the broker in the bargaining model receives a surplus of* $x_{broker} = v_B - v_A - x_A - x_B = 1 - (c_A + x_A) - (c_B + x_B)$ *when trade takes place, and* $-x_A - x_B$ *when trade does not take place. When trade does not take place, C is the efficient outcome in the procurement problem. In this case, the individual rationality condition and the requirement that total payment is less than the cost of the best alternative imply that* $x_A + x_B \leq \delta$*.*

For any procurement mechanism that always pays less than the best available alternative, the right hand side of the equation above must be always greater than $-\delta$*. This would yield a bargaining mechanism in which the broker had expected profit greater than* $-\delta$*. However, the results of Myerson and Satterthwaite [31] imply that for any efficient and individually rational mechanism, the expected value of* x_{broker} *is bounded above by a strictly negative value (in other words, there is a minimum subsidy required). For a suitably small* δ*, this is a contradiction.*

4 Approximate Efficiency of First-Price Combinatorial Auctions

It is often desirable to design auctions that choose *efficient* allocations. A procurement auction is *efficient* if it always select the minimum cost feasible set. The VCG mechanism guarantees that the set it selects is efficient. The strong ϵ-Nash equilibria of first-price procurement auctions are not necessarily efficient. For example, if the minimum cost and second-minimum cost feasible sets have costs within ϵ of one another, then it is a strong ϵ-Nash for the second-minimum cost set to bid truthfully and for the minimum cost set to overbid by ϵ. In such a scenario, the first price procurement auction will select the second-minimum cost set. Still, the winning set is approximately efficient as its cost is within ϵ of the minimum cost set. In this section, we prove that this holds in general, that is the strong ϵ-Nash equilibria of first-price procurement auctions are approximately efficient.

Theorem 4.1. *Let* $\mathbf{b}$ *be a strong* ϵ*-Nash equilibrium of a deterministic first-price procurement auction. Then the cost* $c(S)$ *of the winning set* S *in a first-price procurement auction is at most the cost* $c(S^*)$ *of the minimum cost feasible set* S^* *plus an additive factor of* ϵn*:*

$$c(S) \leq c(S^*) + \epsilon n.$$

Proof. The proof is by contradiction. Assume the winning set S is not approximately efficient, that is, $c(S) > c(S^*) + \epsilon n$. Define a new bid vector $\mathbf{b}'$ in which the agents who are not winning but are in a minimum cost feasible set lower their bids to just above their cost:

$$b'_i = \begin{cases} \min\{b_i, c_i + \epsilon\} & \text{for } i \in S^* - S, \\ b_i & \text{otherwise.} \end{cases}$$

In this bid vector, S^* is cheaper than S:

$$\begin{aligned} b'(S) - b'(S^*) &= b(S - S^*) - b'(S^* - S) \\ &\geq c(S - S^*) - (c(S^* - S) + \epsilon n) \\ &= c(S) - c(S^*) - \epsilon n \\ &> 0. \end{aligned}$$

This contradicts Lemma 3.1 with $T = \{i \in S^* - S : b_i > c_i + \epsilon\}$. □

5 Payment Bounds for Flow Auctions

In this section, we bound the overpayments of first-price flow auctions. We assume that we have a deterministic tie-breaking rule so that if there is more than one cheapest feasible flow, we take the lexicographically first integral one. We consider two settings. In the *known demand path auction* studied in Section 5.1, the total demand of the auctioneer is known to the auctioneer and all the bidders at the time of the auction. It is easy to imagine that the assumptions of this model might be unrealistic in practice. Can the total demand really be known before it is realized? What if the auctioneer wishes to buy flow in advance? In our second model, the *unknown demand path auction* studied in Section 5.2, the auctioneer and bidders instead know a probability distribution over possible demand values.

Notation For a graph G, let $\mathbf{c}$ be the vector of edge costs, $\mathbf{b}$ be the vector of edge bids, and $F_{\mathbf{w}}(k, G)$ be the set of edges in the winning k-flow[1] in G with respect to edge weights $\mathbf{w}$ (as we only consider deterministic first-price auctions, this is well-defined). We will refer to $F_{\mathbf{c}}(k, G)$ as the minimum cost k-flow and $F_{\mathbf{b}}(k, G)$ as the minimum price k-flow with respect to bid vector $\mathbf{b}$. When the bids, costs, or graph is clear from the context, we will sometimes drop them from the notation. As a shorthand, we sometimes write $c(k)$ for the (cost of the) lowest cost k-flow. Finally, to be consistent with the previous notation, we denote the number of agents, or edges in G, by n.

5.1 Known Demand Path Auction

In the known demand setting, we assume that the auctioneer has a publicly-known demand k. We will show that in such settings, the payments in a strong ϵ-Nash equilibrium of deterministic first-price auctions is bounded. In particular, we show that the overpayment to each unit of flow is (approximately) at most the true marginal cost of sending an additional unit of flow (see Theorem 5.1). Together with the observation that the VCG mechanism pays each edge a bonus at least as large as this marginal (see Theorem 5.2), this shows that the payments in first-price auctions are (approximately) bounded by the payments in the VCG auction. We saw in Section 4 that the winning set in the first-price auction is also (approximately) efficient. These statements regarding the payments and efficiency of first-price auctions suggest that first-price auctions perform better than VCG auctions. However, first-price auctions have a significant drawback; it is not clear how agents might converge to a strong ϵ-Nash equilibrium. We partially address this concern by proposing another auction whose ϵ-Nash equilibria have the same properties as the strong ϵ-Nash equilibria of a first-price auction.

[1]The weight of this flow is equal to the weight of the minimum weight k-flow, that is requiring integrality does not change the value of the optimal solution.

5.1.1 Payment Bound

We first bound the payments in a strong ϵ-Nash equilibrium (see Definition 2) of a deterministic first-price auction. The edges announce bids and the auctioneer runs a first-price auction to select a cheapest k-flow according to the bid vector, paying each edge on the flow an amount equal to his bid. By Theorem 3.1, strong ϵ-Nash equilibria exist for such auctions. Given the existence of strong ϵ-Nash equilibria, we can bound the payments in any such equilibrium. In order to develop some intuition for the proof, it is useful to first consider sending 1 unit of flow in a graph consisting of just two parallel edges from the source s to the sink t of costs a and b, $a > b + \epsilon$. The lower-true-cost edge must be allocated the flow in equilibrium since he can bid just under the true cost of the higher cost edge and be guaranteed a profit of at least ϵ. Therefore, by the conditions of a strong ϵ-Nash equilibrium, we can assume that the bid of the higher cost edge is at most ϵ more than his true cost, and so the overpayment of any equilibrium will be at most $a + \epsilon - b$. The crux of this argument was to bound the bid of the winning path by the bid of an augmenting path. Since the augmenting path does not receive flow, Lemma 3.1 permitted us to assume, for the purposes of bounding the price, that the bid of this path was close to its true cost. This proof idea easily extends to auctions for k-flows in general graphs as can be seen below.

Theorem 5.1. *The total payment of the deterministic first price k-flow auction in a strong ϵ-Nash equilibrium is at most*

$$k\left[c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k))\right] + kn\epsilon,$$

where $\mathbf{c}$ *is the vector of true edge costs.*

Proof. Fix a strong ϵ-Nash equilibrium vector of bids $\mathbf{b}$ and define bid vector $\mathbf{b}'$ such that

$$b'_i = \begin{cases} b_i & \text{for } i \in F_{\mathbf{b}}(k), \\ \min\{b_i, c_i + \epsilon\} & \text{otherwise.} \end{cases}$$

By Lemma 3.1, $F_{\mathbf{b}}(k)$ is a minimum price k-flow with respect to $\mathbf{b}'$. Consider the (non-integral) flow $(k/(k+1))F_{\mathbf{c}}(k+1)$, that is the flow which sends $k/(k+1)$ units of flow along the flow paths determined by $F_{\mathbf{c}}(k+1)$. Since $F_{\mathbf{b}}(k)$ is a lowest-price k-flow with respect to $\mathbf{b}'$ and using the integrality of optimal network flows [33], we have

$$\left(\frac{k}{k+1}\right) b'(F_{\mathbf{c}}(k+1)) - b'(F_{\mathbf{b}}(k)) \geq 0. \tag{8}$$

Define edge sets

$$\begin{aligned} E_+ &= \{e \in F_{\mathbf{c}}(k+1) - F_{\mathbf{b}}(k)\} \\ E_o &= \{e \in F_{\mathbf{c}}(k+1) \cap F_{\mathbf{b}}(k)\} \\ E_- &= \{e \in F_{\mathbf{b}}(k) - F_{\mathbf{c}}(k+1)\} \end{aligned}$$

Then equation 8 reduces to

$$\left(\frac{k}{k+1}\right) b'(E_+) - \left(\frac{1}{k+1}\right) b'(E_o) - b'(E_-) \geq 0$$

which, solving for $b'(E_o) + b'(E_-)$, gives

$$\begin{aligned} b(F_{\mathbf{b}}(k)) &= b'(E_o) + b'(E_-) \\ &\leq k(b'(E_+) - b'(E_-)) \\ &\leq k(c(E_+) + n\epsilon - c(E_-)) \quad (9) \\ &\leq k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{b}}(k)) + n\epsilon) \\ &\leq k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k)) + n\epsilon) \quad (10) \end{aligned}$$

where 9 follows from the fact that for any edge $b'_i \geq c_i$ and for all $i \in E_+$, $b'_i \leq c_i + \epsilon$; and 10 follows from the optimality of $F_{\mathbf{c}}(k)$ with respect to $\mathbf{c}$. □

In addition, it is easy to see that this bound is tight. Consider a graph with $(k+1)$ parallel edges where the cost of the bottom k edges is c and the cost of the remaining top edge is $c' > c$. Let all k lower cost edges bid $c' - \epsilon$ for a small $\epsilon > 0$, so their bid is less than the bid of the remaining higher cost edge (whose bid is at least c'). The minimum price k-flow with respect to this bid vector will use the bottom k edges for a total price of $k(c' - \epsilon)$ which approaches $k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k)))$.

Finally, we emphasize that the total payment of our first price mechanism in a strong ϵ-Nash equilibrium is at most $kn\epsilon$ more than the VCG payment for the same graph in a Nash equilibrium.

Theorem 5.2. *Given a graph G with source s and sink t, the VCG payment for k units of flow from s to t is at least $k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k)))$.*

Proof. Let $P_1, \ldots, P_k$ be the k disjoint paths in the selected minimum cost k flow. Fix one path P_i with, say, l edges. We will prove that the sum of payments to edges on this path is at least $c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k))$. Recall that the VCG payment for an edge e on a minimum cost k-flow is

$$c_e + c(F_{\mathbf{c}}(k, G - \{e\})) - c(F_{\mathbf{c}}(k, G)). \quad (11)$$

We construct a new directed multi-graph on the same vertex set as G as follows. We use the term *forward* to mean an edge directed from s to t along the flow path and *backward* to mean an edge directed from t to s. For each edge e on path P_i, add a backward copy of each edge in $F_{\mathbf{c}}(k, G)$ and a forward copy of each edge in $F_{\mathbf{c}}(k, G - \{e\})$, retaining multiplicities. Now add a forward copy of the path P_i to the graph. Label each forward edge e with the cost c_e of the corresponding edge in G and each backward edge with the cost $-c_e$. Then, by equation 11, the sum of edge weights in this graph equals the VCG sum of payments to edges on path P_i. Note that this graph is a union of l $s-t$ flows, l $t-s$ flows, and one s-t path. Thus, the in-degree of every vertex except s and t is equal to its out-degree, and for s (t), the out-degree is one more (less) than its in-degree. For every pair of vertices, *cancel* the 2-edge cycles connecting them. That is, if the vertices are connected by k_1 forward edges and k_2 backward edges, replace the edges by $k_1 - k_2$ forward edges if $k_1 > k_2$, $k_2 - k_1$ backward edges if $k_1 < k_2$, or simply remove the edges if $k_1 = k_2$ (this does not change the degree or edge weight properties of the graph discussed above). Call the resulting graph G'. As the sum of edge weights in G' equals the sum of VCG payments to edges on path P_i, we can bound the sum of VCG payments to edges on path P_i by bounding the sum of edge weights in G'.

First note every edge of $F_{\mathbf{c}}(k, G)$ is either non-existent or directed backward in G': for edges $e \in F_{\mathbf{c}}(k, G) - P_i$, e is added exactly once in the backward direction and at most once in the forward direction by each of the l edges in P_i; for edges $e \in P_i$, e is added exactly once in the backward direction and at most once in the forward direction by each of the $l-1$ edges e' in P_i, $e' \neq e$. Furthermore, edge e is added once in the backward direction by itself and once in the forward direction in the last step of the construction of G'.

Select a path P from s to t in G' (such a path exists by the degree properties discussed above). As edges of $F_{\mathbf{c}}(k,G)$ exist only in the backward direction, our path P is a valid augmenting path in the original graph G, and so its weight is at least $c(F_{\mathbf{c}}(k{+}1,G){-}c(F_{\mathbf{c}}(k,G))$ by minimality of $F_{\mathbf{c}}(k+1,G)$. We claim the weight of P is at most the sum of edge weights in G' (which equals the VCG payment), proving the result. This follows from the fact that, due to its degree properties, G' can be written as a union of P and a set of disjoint cycles, and, since $F_{\mathbf{c}}(k,G)$ is a minimum cost k-flow in G, the sum of edge weights on any cycle must be non-negative. Otherwise we could construct a cheaper k-flow in G by replacing the backward edges of a negative cycle with the forward edges in $F_{\mathbf{c}}(k,G)$: specifically, if C is a negative cycle in G' with backward edges A and forward edges B, then $F_{\mathbf{c}}(k,G) - A + B$ is a cheaper k-flow in G. □

5.1.2 Implementation in ϵ-Nash

The simple first-price auction may have costly ϵ-Nash equilibria, as shown in Example 3.3. In Section 5.1 we used the strong ϵ-Nash solution concept to get around this problem. However, assuming that the bidders will reach an strong ϵ-Nash equilibrium is perhaps too strong an assumption: it requires extensive coordination between agents. In this section, we present a variant of a first-price auction in which every ϵ-Nash equilibrium results in a low price.

One idea to achieve this is to pay edges a bonus that increases as their bid decreases. This encourages edges to submit low bids. However, this has the side-effect of giving edges incentives to bid even below their true cost, as long as they remain off the winning flow. This would make the bargaining problem that edges must solve much more complex, as it would include bargains between winning and losing edges. Alternatively, we could instead send flow on each edge with some probability that increases as the bid decreases. Thus an edge will not bid below its true cost, but it might have an incentive to bid very high. Using a combination of these two ideas, we can construct a payoff function such that an edge will bid close to its true cost if it is not on the lowest true cost flow. This is known as *virtual implementation* in the economics literature (see, for example, Jackson [21]). If the bonuses and probabilities are small enough, then the extra payment will not be very large in expectation, and we can prove a bound on the total payment of the mechanism similar to that in Theorem 5.1.

We describe the techniques in this section in the setting of path auctions, although they extend to more general settings as noted. Assume that there is a value B such that no edge bids more than B. (Alternatively, B can be the maximum amount that the buyer is willing to pay.) Further, we assume that the edges are risk-neutral. The mechanism starts by computing a collection of (not necessarily simple) paths $\{P_e\}$. The mechanism then solicits a bid b_e from each edge e. The lowest-price path is almost always picked; however, with a small probability, one of the paths from the collection is picked instead. In addition, each edge is paid a small bonus that depends on the bids. The selection probability and bonus are chosen to ensure that it is optimal for every edge that is *not* on the lowest-price path to bid its true cost. For simplicity, we present the mechanism and analysis for a single unit of flow; the results can easily be extended to any constant $k > 1$ units of flow.

Mechanism RandomPath: The parameters α and τ are selected to be small positive constants such that $\alpha < \min\{n^{-2}B^{-1}, \frac{2}{1+2n}\}$ and $\tau < \alpha n^{-1}B^{-1}$.

1. For each edge e, find P_e, a (not necessarily simple) path from s to t through e. Let $\mathcal{P} = \{P_e\}_{e\in G}$. Note that an edge e may appear in multiple paths in $\mathcal{P}$.
2. Solicit bids $\mathbf{b} = (b_1, \ldots, b_e, \ldots, b_n)$ from the edges.
3. For each path $P \in \mathcal{P}$, compute
$$\sigma_P = \alpha - \tau \sum_{e\in P} b_e$$
4. Select each path $P \in \mathcal{P}$ with probability σ_P; with probability $(1 - \sum_{P\in\mathcal{P}} \sigma_P)$, select the lexicographically first lowest price path. Call the selected path P^*. Pay each edge $e \in P^*$ its bid b_e.
5. In addition to any payment edge e may get in step 4, pay each edge $e \in G$ the sum $f_e(\mathbf{b}) = \sum_{P\in\mathcal{P}, P\ni e} f_e^P(\mathbf{b})$, where
$$f_e^P(\mathbf{b}) = \alpha(B - b_e) + \tau b_e \sum_{j\in P} b_j - \tau \frac{b_e^2}{2}$$

Our payment rule is constructed in a way that encourages bidders not receiving flow to bid their true cost. Note that the bonus increases as the bid decreases, but the expected selection payment decreases as the bid decreases. Intuitively, we design the bonus and selection probabilities so that the total payoff function is maximized when $b_i = c_i$. Note that if an edge is selected, it incurs its true cost. In this way, the true cost automatically enters his expected payoff function—the mechanism does not need to know the cost in order to achieve the maximum at $b_i = c_i$.

Lemma 5.1. *For any edge e not on the lowest-price path with bids $\mathbf{b}$, if $b_e \notin [c_e - \sqrt{2\epsilon/\tau}, c_e + \sqrt{2\epsilon/\tau}]$, then $b_e = c_e$ will increase the expected payoff to e by at least ϵ.*

Proof. With the bid vector b, e's expected payoff is

$$\begin{aligned} f_e(\mathbf{b}) \quad + \quad \sum_{P\ni e} \sigma_P(b_e - c_e) \quad &= \sum_{P\ni e} [f_e^P(\mathbf{b}) + \sigma_P(b_e - c_e)] \\ &= \sum_{P\ni e} [\alpha B - \tau \frac{b_e^2}{2} + \tau c_e \sum_{j\in P} b_j - \alpha c_e] \end{aligned}$$

Let $g(b_e) = [\alpha(B - c_e) - \tau\frac{b_e^2}{2} + \tau c_e \sum_{j\in P} b_j]$. Then, $g(b_e)$ is a quadratic function of b_e. Observe that $\frac{\partial g(b_e)}{\partial b_e} = -\tau b_e + \tau c_e = 0$ when $b_e = c_e$; at this point, $\frac{\partial^2 g(b_e)}{\partial^2 b_e} = -\tau < 0$. This is true for all paths P containing e. Further, for $\Delta > 0$,

$$g(c_e) - g(c_e + \Delta) = \tau c_e \Delta + \tau\Delta^2/2 - \tau c_e \Delta = \tau\Delta^2/2$$

Similarly, $g(c_e) - g(c_e - \Delta) = \tau\Delta^2/2$. Thus, if $b_e < c_e - \sqrt{2\epsilon/\tau}$, then edge e has incentive to raise his bid to $b_e = c_e$. Similarly, if $b_e > c_e + \sqrt{2\epsilon/\tau}$, then edge e has incentive to decrease his bid to $b_e = c_e$ (even if this puts him on the lowest-price path, then his payoff is still $g(c_e)$ per path so the above calculation still holds). □

Lemma 5.1 implies that if ϵ-Nash equilibria exist in mechanism RandomPath, then any edge not on the lowest-price flow must bid close to its true cost. This will help us bound the total expected payment in an ϵ-Nash, but first we must prove that ϵ-Nash equilibria exist in this mechanism. Indeed the same construction as in Theorem 3.1 yields an ϵ-Nash equilibrium.[2]

Lemma 5.2. *For any cost vector* $\mathbf{c}$ *and any* $\epsilon > 0$*, an* ϵ*-Nash equilibrium always exists in the mechanism RandomPath.*

Proof. Construct a bid vector $\mathbf{b}$ as in Theorem 3.1. By this construction, the lowest-cost path equals the lowest-price path. We have $b_e = c_e$ for any edge e that is *not* on the lowest-price path. Edges on the lowest-price path bid close to the maximum they can while still remaining on the lowest-price path (see the proof of Theorem 3.1 for the precise construction).

Following the analysis of $g(b_e)$, the expected payoff in Lemma 5.1, b_e maximizes e's payoff. (Note that e can only get onto the lowest-price path by bidding below its cost, which would result in a loss.) It remains to show that every edge i on the lowest-price path would not significantly benefit by changing it's bid. Note that, by construction of the bid vector, if i increased its bid by more than $\epsilon/2$, it would no longer be on the lowest-price path. Further, because of the shape of the bonus payoff function, i's expected gain $g(b_e)$ from the bonus and probability of off-path selection would also drop. Thus, i cannot possibly gain more than ϵ by raising its bid. Consider the possibility that i lowers its bid by x. Then, i would still be on the lowest-price path. It would lose at least $(1 - n\alpha)x$ in profit from being on the lowest-price path, and gain at most $(g_e(b_e) - g_e(b_e - x)) = \tau(\frac{1}{2}x^2 + (c_e - b_e)x)$ in $g_e(b_e)$ per path. As $b_e \geq c_e$ in $\mathbf{b}$, its total gain is at most $\frac{n\tau}{2}x^2$. As $x \leq B$, the loss is more than the gain for any choice of τ less than $2(1 - n\alpha)/(nB)$ or, rewriting in terms of α, $\alpha < \frac{2}{1+2n}$. These conditions can be guaranteed by the choice of α and τ. □

Now, we observe that the values α and τ can be chosen small enough to make the probabilities $\{\sigma_P\}$ and bonuses $f_e^P(\mathbf{b})$ arbitrarily small. Thus, the total payment to edges not on the shortest path is very small. The bound on the payment of the mechanism RandomPath is more sensitive to the value of ϵ because edges not on the lowest-price path get very small payments in expectation. However, we can show that, in the limit as $\epsilon \to 0$, the maximum expected payment in any ϵ-Nash equilibrium is bounded. The following proof can be generalized to the flow setting to derive a bound similar to that in Theorem 5.1.

Theorem 5.3. *Choose any* $\alpha < n^{-2}B^{-1}, \tau < \alpha n^{-1}B^{-1}$*. For these values of* α *and* τ*,*

$$\lim_{\epsilon \to 0} \max_{\epsilon\text{-}NE\ \mathbf{b}} \{\textit{Total payments with bids } \mathbf{b}\} \to c(2) - c(1) + 3\alpha n^2 B.$$

Proof. Let $\mathbf{b}$ be an ϵ-Nash equilibrium bid vector, for sufficiently small ϵ. The total probability that the mechanism picks a path other than the lowest-price path is bounded by $n\alpha$. Any such path can have at most n edges on it, each with price at most B. Thus, the expected payment for using one of these paths is at most $\alpha n^2 B$. Similarly, we can bound the bonus $f_e(\mathbf{b})$ paid to any edge e: $f_e(\mathbf{b}) \leq n[\alpha B + \tau n B^2]$. This is always less than $2\alpha n B$.

[2]Mechanism RandomPath can be extended to the general procurement setting. The proof of the following theorem can be generalized to prove existence of ϵ-Nash equilibria in this setting as well.

Finally, using Lemma 5.1, we know that any edge not on the lowest-price path bids at most $c_e + \sqrt{2\epsilon/\tau}$. Combining this with a similar argument to Theorem 5.1, we can bound the total payment to edges on the lowest-price path by

$$b(F(1)) \leq c(2) - c(1) + n\sqrt{2\epsilon/\tau}$$

In the limit as $\epsilon \to 0$, the last term is negligible. Adding up all three sources of payment, we get the required result. □

Recall that mechanism RandomPath needs to compute a set of paths $\{P_e\}$, where P_e is a path from s to t that uses e. If e is to be relevant to the path auction, such a path must exist, however, it is not always straightforward to compute. In particular, if the network is a general *directed* graph, it is NP-hard to compute such a path, since it reduces to the two disjoint paths problem, which is NP-complete [15].

However, there are many interesting classes of graphs for which it is possible to compute such a path P_e in polynomial time, including undirected graphs and directed acyclic or planar graphs [15]. We can also modify the mechanism to ask each bidder to exhibit such a path, thus transferring the computational burden on to the bidders. Also, these paths may be precomputed and used in many executions of the mechanism—they do not depend on the costs or bids.

Another possibility is to use a set of covering paths that do not all terminate at t—this can be easily computed, even for general directed graphs. Then, if the path is picked, some arbitrary traffic is sent along this path. After this "audit" traffic has been delivered, the lowest-price path is used for the intended traffic from s to t. As long as the mechanism can verify that the traffic is correctly delivered, the edges would still have an incentive to bid as specified. Similarly, if we could verify the exact path that the traffic used, we could use non-simple paths to cover the edges; again, a set of non-simple covering paths can easily be found.

5.2 Unknown Demand Path Auction

In the previous sections, we studied first-price auctions to meet a known demand, argued that they had stable Nash equilibria, and showed how to adjust this auction so that the equilibria chosen by the auctioneer had relatively small overpayments. In practice, however, it may not be possible to defer the setting of prices until the demand is known. In this section, we examine the problem of achieving stable prices without advance knowledge of the demand. Instead, the bidders and auctioneer share knowledge of a common prior or *probability distribution* over the possible demands.

Ideally, we would like our results for first-price auctions with known demand to carry over. For example, we proved in Section 5.1 that a first price auction for k units of demand led to a payment of $P_k = k[c(F_c(k+1)) - c(F_c(k))]$ in any strong ϵ-Nash equilibrium. It is thus natural to hope that the same auction operating over *random* k also has strong ϵ-Nash equilibria with expected payment $E_k[P_k]$. This turns out to be false—in fact, as we will show, a first-price auction might not even have ϵ-Nash equilibria (recall that strong ϵ-Nash equilibria are a subset of ϵ-Nash equilibria). As ϵ-Nash equilibria do not exist in first-price auctions, we turn to more complex auctions. We will exhibit an auction involving *two parameter bids* that, unlike the single-parameter first-price auction, *does* have ϵ-Nash equilibria. Furthermore, using an indifference-breaking technique similar to that of the mechanism RandomPath, we can restrict the set of equilibria in a variant of this auction to ones with bounded payments. The bound is not quite the $E_k[P_k]$ we hoped to achieve, but does bear a clear resemblance to it. Unfortunately, we are unable to prove that this auction is implementable in polynomial time as it involves solving an integer program. It remains to be seen if further modifications of this auction can result in a polynomial-time auction with bounded payments.

5.2.1 Definitions and Notation

The unknown demand case is modeled as follows: The demand can take any integral value in the range $[1, r]$, where r is a positive integer. Further, there is a known prior distribution on the demand values; say that the demand is k with probability p_k, for $k = 1, 2 \ldots, r$. We assume for simplicity that $p_k > 0$ for all k; our results easily extend to a situation in which $p_k = 0$ for some values of $k \in \{1, \ldots, r\}$.

An auction for the unknown demand case receives bids, and announces flows $F_1, F_2, \ldots, F_r$ for each possible demand value. For a first-price auction in this setting, each $F_k \in \mathcal{F}$ must be a minimum price k-flow. We call the collection $\mathcal{F} = \{F_1, F_2 \ldots, F_r\}$ a candidate solution. We also identify a solution $\mathcal{F}$ with the set of edges in the union $F_1 \cup F_2 \cup \cdots \cup F_r$, and say that $i \in \mathcal{F}$ if $i \in F_k$ for some k.

As before, we use $c(\mathcal{F})$ to denote the total expected cost of a solution $\mathcal{F} = (F_1, \ldots, F_r)$ when the individual edge costs are c, and $\tilde{a}(\mathcal{F})$ to denote the price of the flow $\mathcal{F}$ when the bids are $\tilde{a}$. When the auction is clear from the context, we will denote the auction output by $\hat{\mathcal{F}}(\tilde{a})$.

5.2.2 Impossibility of ϵ-Nash Equilibria in First-Price Auctions

In this section, we show that a first-price auction may have *no* ϵ-Nash in the unknown demand case. Intuitively, this is because edges must tradeoff the *probability* of receiving flow with the *profit* of receiving flow. With a high bid, the profit is large, but the probability of winning the auction is low. If the other bids are also high, an edge will prefer to lower its bid to win with a higher probability. This will lead other edges to lower their bids so as to restore their high winning probability. Now, however, the first edge will increase its bid so as to increase its profit at the expense of its winning probability, and so a cycle emerges in the bidding strategies, as the following example shows.

Consider a graph with four parallel edges W, X, Y, and Z between the source and the sink, with true costs w, x, y, and z respectively. The demand is either 1, 2 or 3; for simplicity, let the probability of each demand value be $\frac{1}{3}$. Assign the costs such that $w + 50\epsilon < x + 42\epsilon = y + 12\epsilon = z$. Suppose there W, X, Y, Z bid a, b, c, d respectively. The proof repeatedly uses the ϵ-Nash conditions to show that one of the following must hold: (1) There is an agent who would gain by raising its bid, or, (2) There is an agent who would gain by undercutting another agent to win with a higher probability.

Theorem 5.4. *There is no pure-strategy ϵ-Nash equilibrium in the unknown demand first-price auction.*

Proof. First we prove a series of inequalities that the bids must satisfy in an ϵ-Nash equilibrium:

Claim 1: $a, b, c \leq d$.
Proof: First, suppose $d > y + 3\epsilon$, and $c > d$. Then, by changing its bid to $d - \delta$, for small

enough δ, Y would be selected with probability 1/3 and so get utility greater than ϵ; thus, any solution in which Y had 0 expected payoff would not be an ϵ-Nash equilibrium. As the same is true for w and x, we must have $a, b, c \leq d$. Now, suppose $d \leq y + 3\epsilon$. Then, $d < z - 3\epsilon$, and as Z is selected with probability 1/3, its payoff is less than $-\epsilon$, which cannot be true in the equilibrium. Thus, in this case too, we have $d \geq a, b, c$.

Claim 2: $d > y - 3\epsilon$.
Proof: If $d \leq y - 3\epsilon$, then Y could not underbid Z without having expected utility less than $-\epsilon$. Hence, Z would be chosen with probability at least $\frac{1}{3}$ (if the demand was 3). But $d < z - 3\epsilon$, and hence Z's expected utility would be less than $-\epsilon$, and hence this cannot be true in an ϵ-Nash equilibrium.

Claim 3: $a, b, c > x + 21\epsilon$.
Proof: Suppose the order of the bids is $a < b < c$. Then, by Claim 1, W wins with probability 1, X with probability 2/3, and Y with probability 1/3. Thus, we must have $(b - a) \leq \epsilon$, $(c - b) \leq 2\epsilon$, and $(d - c) \leq 3\epsilon$, or else one of W, X, Y could increase her profit by ϵ. A similar argument holds if the bids are in a different order. Thus $a, b, c > d - 6\epsilon$. By Claim 2, this implies $a, b, c > y - 9\epsilon$ which equals $x + 21\epsilon$.

Claim 4: $b < c$.
Proof: By Claim 3, $c > x + 21\epsilon$. If we had $b \geq c$, then X could deviate by bidding $c - \delta$. This would involve a bid reduction of at most 6ϵ, but would enable X to win with a $\frac{1}{3}$ additional probability, leading to a net gain of at least ϵ.

Claim 5: $a < b$.
Proof: If $a \geq b$, W could deviate to $b - \epsilon$, resulting in a gain of at least ϵ, as above.

These claims imply that (a, b, c, d) is not an ϵ-Nash equilibrium: We have shown that $a < b < c \leq d$. It also must be true that $(c - b) < 2\epsilon$, and $c > y - 3\epsilon$. Thus, $b > x + 25\epsilon$. Further, $(b - a) < \epsilon$. Hence, X could deviate to $a - \delta$, resulting in a net gain of greater than ϵ. □

5.2.3 Implementation in ϵ-Nash Using a 2-Parameter Bidding Scheme

In this section, we show that by allowing 2-parameter bids, we can define an auction with ϵ-Nash equilibria. Intuitively, a two-parameter auction gets around the problem of a single-parameter auction by letting the edges express their preferences over the entire price-probability space. It allows to an edge to bid a "price" such that the expected payment of any edge with a non-zero probability of winning is equal to its price. In particular, we will allow edges to report their cost along with a *demanded profit* and then guarantee that the *expected* payment of a winning edge is exactly its reported cost plus its demanded profit.

Auction 2-Parameter:

In the following auction, each edge i submits a pair $\tilde{a}_i = (\tilde{c}_i, \tilde{u}_i)$ as its bid, where $\tilde{c}_i$ is interpreted as the reported cost of edge i, and $\tilde{u}_i$ is interpreted as the profit that edge i demands.

1. Define an indicator variable x_{ik} for the event that edge i is on the selected flow F_k, and y_i for the event that edge i is selected to be on some flow. Also, for any node α in the network, let $\text{In}(\alpha)$ denote the set of incoming edges, and $\text{Out}(\alpha)$ denote the set

of outgoing edges. Find an optimal solution to the following integer program (IP for short).

$$\text{minimize} \quad \sum_{k=1}^{r}\left[p_k \sum_{i\in E}\tilde{c}_i x_{ik}\right] + \sum_{i\in E} y_i \tilde{u}_i \tag{12}$$

$$\text{subject to} \quad \forall\, \alpha \neq s,t,\ \forall\, 1 \leq k \leq r:\ \sum_{i\in \mathrm{Out}(\alpha)} x_{ik} - \sum_{i\in \mathrm{In}(\alpha)} x_{ik} = 0 \tag{13}$$

$$\forall\, 1 \leq k \leq r:\ \sum_{i\in \mathrm{Out}(s)} x_{ik} - \sum_{i\in \mathrm{In}(s)} x_{ik} = k \tag{14}$$

$$\forall\, 1 \leq i \leq n,\ \forall\, 1 \leq k \leq r:\ y_i - x_{ik} \geq 0 \tag{15}$$

$$\forall\, 1 \leq i \leq n,\ \forall\, 1 \leq k \leq r:\ x_{ik} \in \{0,1\}$$

$$\forall\, 1 \leq i \leq n,\ \forall\, 1 \leq k \leq r:\ y_i \in \{0,1\}$$

2. Set $F_k = \{i : x_{ik} = 1\}$ and $\mathcal{F} = \{F_1, \ldots, F_r\}$. For each $i \in \mathcal{F}$, calculate the probability $\rho_i = \sum_{\{k|i\in F_k\}} p_k$ that i wins. If the actual demand turns out to be r, use the edges in F_k to route the flow, and pay each edge $i \in F_k$ a sum $\tilde{c}_i + \frac{\tilde{u}_i}{\rho_i}$.

Remark 5.1. *Notice with these payments, IP 12 chooses a flow solution which minimizes the total expected payment for a fixed bid vector: constraints 13 and 14 guarantee that the set $F_k = \{i : x_{ik} = 1\}$ form a feasible k-flow and constraint 15 guarantees that edges selected to be on a flow are paid their reported cost.*

We now prove that this auction has ϵ-Nash equilibria. To develop some intuition for the proof, recall that in the known demand case, only bidders on the cheapest flow had the flexibility to submit a bid significantly more than their cost and still win the auction. A similar statement holds here when the first parameter of all bids are restricted to be equal to the cost. In particular, the following bid vector should intuitively be an ϵ-Nash equilibrium: for edges $i \notin \hat{\mathcal{F}}(\tilde{a})$, set $\tilde{a}_i = (c_i, 0)$; for edges $i \in \hat{\mathcal{F}}(\tilde{a})$, set $\tilde{a}_i = (c_i, \tilde{u}_i)$ where the $\tilde{u}_i$ divide up the available profit (the difference between the price of the cheapest and second cheapest flow). Edges $i \notin \hat{\mathcal{F}}(\tilde{a})$ can not afford to decrease their bids and have no chance of winning by increasing their bids, so they have no profitable deviation. As the expected payment of any edge $i \in \hat{\mathcal{F}}(\tilde{a})$ is the same regardless of their winning probability, these edges also have no incentive to decrease their bid. By an appropriate choice of $\{\tilde{u}_i\}$, we can arrange that if they increase their bid then they will drop out of the solution.

We formalize this argument by using a linear-programming technique similar to the proof of Theorem 3.1. The variables of the linear program (LP) are the profits demanded by the bidders (that is, the second parameter of the bid). The LP constrains the total profit demanded by a set of bidders to be at most the cost-savings induced by this set. Let $\mathcal{F}^*$ be the minimum cost solution and u_i be a variable corresponding to the profit demanded by bidder i. Consider the following linear program.

$$\begin{aligned}
\text{maximize} \quad & \sum_{i=1}^{n} u_i && (16)\\
\text{subject to} \quad & \forall \text{ feasible solutions } \mathcal{F}: \sum_{i \notin \mathcal{F}} u_i \leq c(\mathcal{F}) - c(\mathcal{F}^*) && (17)\\
& \forall\, 1 \leq i \leq n: \; u_i \geq 0
\end{aligned}$$

This LP is clearly feasible as $u_i = 0$ for all i satisfies all constraints. We will show that for an optimal solution $\{u_i^*\}$, the set of bids $\{(c_i, \max\{0, u_i^* - \epsilon/(2n)\})\}$ form an ϵ-Nash equilibrium.

Theorem 5.5. *For any $\epsilon > 0$, let $u_i = \max\{0, u_i^* - \frac{\epsilon}{2n}\}$ and consider the bid profile defined by $a_i^- = (c_i, u_i)$ for each edge i. Then a^- is an ϵ-Nash equilibrium.*

The proof uses three lemmas regarding the bids of the minimum price solution. The first lemma shows that edges $i \notin \mathcal{F}^*$ not in the minimum cost solution have zero demanded profit (that is, $\tilde{u}_i = 0$). This confirms the intuition that, as in the known-demand case, only edges in the minimum-cost solution can demand a payment significantly more than their cost.

Lemma 5.3. *The minimum cost solution includes all i with $u_i^* > 0$.*

Proof. Consider the inequality in LP 16 corresponding to solution $\mathcal{F}^*$. This inequality states that $\sum_{i \notin \mathcal{F}^*} u_i^* \leq 0$. Together with the non-negativity constraints, this implies that $u_i^* = 0$ for all edges i not in the minimum cost solution. Thus the minimum-cost solution includes all edges i with $u_i^* > 0$. □

The second lemma supports the intuition that the minimum cost solution $\mathcal{F}^*$ has minimum price.

Lemma 5.4. *The minimum cost solution is a minimum expected price solution with respect to bids $\tilde{a}_i = (c_i, \tilde{u}_i^*)$.*

Proof. As the first parameter of any bid $\tilde{a}_i$ is c_i, the expected price of any solution $\mathcal{F}$ is equal to its expected cost plus the sum of demanded profits of its edges. Since $u_i^* = 0$ for $i \notin \mathcal{F}^*$, we have

$$\tilde{a}(\mathcal{F}^*) = c(\mathcal{F}^*) + \sum_{i=1}^{n} u_i^*. \tag{18}$$

For any flow $\mathcal{F}$, the inequality 17 of LP 16 corresponding to $\mathcal{F}$ states that $c(\mathcal{F}^*) \leq c(\mathcal{F}) - \sum_{i \notin \mathcal{F}} u_i^*$. Adding $\sum_{i=1}^{n} u_i^*$ to both sides and using equation 18 gives $\tilde{a}(\mathcal{F}^*) \leq \tilde{a}(\mathcal{F})$. □

The third lemma argues that no single edge is essential to the minimum price solution. In other words, for each edge there is a minimum price solution that avoids that edge. Intuitively, if this were not the case, then the edge ought to be able to demand extra profit.

Lemma 5.5. *With bids $\tilde{a} = (c_i, \tilde{u}_i^*)$, for any edge i there is a minimum price solution $\mathcal{F}^{(i)}$ that does not contain i.*

Proof. Let $\mathcal{F}$ be a solution not containing i and suppose every minimum price solution contains i. Then, by Lemma 5.4, the inequality corresponding to $\mathcal{F}$ must be strict. As this holds for *any* solution $\mathcal{F}$ not containing i, *every* inequality containing u_i^* is strict. Therefore $u_i^* + \delta$ is a feasible solution for some $\delta > 0$, contradicting the optimality of solution u_i^*. □

Proof of Theorem 5.5. Suppose a^- is not an ϵ-Nash equilibrium. Then, there is some i which can change its bid to increase its payoff by ϵ. Let (c_i', u_i') be i's successful strategy, and let a' denote the bid profile given by $a_i' = (c_i', u_i')$ and $a_j' = a_j^-$ for all $j \neq i$. Let $\mathcal{F}$ be the solution output by the mechanism with bids a^- and $\mathcal{F}'$ be the solution output by the mechanism with bids a'. Note it must be the case that $i \in \mathcal{F}'$.

We observe that the change in expected price of $\mathcal{F}'$ from bid vector a' to a^- is at least ϵ. Let ρ_i be the probability (over the demand distribution) that i is in solution $\mathcal{F}'$. Then i's utility increases from u_i to $u_i' + (c_i' - c_i)\rho_i$, and so by assumption $u_i' + (c_i' - c_i)\rho_i - u_i \geq \epsilon$. Therefore, as only i's bid changes and $i \in \mathcal{F}'$,

$$a'(\mathcal{F}') - a^-(\mathcal{F}') = (u_i' + \rho_i c_i') - (u_i + \rho_i c_i) \geq \epsilon. \tag{19}$$

Now, by Lemma 5.5, there is a solution $\mathcal{F}^{(i)}$ not containing i which has minimum price with respect to bids $\tilde{a} = (c_i, \tilde{u}_i^*)$. Let $\mathcal{F}^{(i)}$ be that solution. Then $\tilde{a}(\mathcal{F}^{(i)}) \leq \tilde{a}(\mathcal{F}')$. Note that for any solution $\mathcal{F}$, the price with respect to bids a^- is within $\epsilon/2$ of the price with respect to bids $\tilde{a}$: $a^-(\mathcal{F}) \leq \tilde{a}(\mathcal{F}) \leq a^-(\mathcal{F}) + \epsilon/2$. Therefore

$$\begin{aligned}
a'(\mathcal{F}^{(i)}) &= a^-(\mathcal{F}^{(i)}) \\
&\leq \tilde{a}(\mathcal{F}^{(i)}) \\
&\leq \tilde{a}(\mathcal{F}') \\
&\leq a^-(\mathcal{F}') + \epsilon/2 \\
&< a'(\mathcal{F}'),
\end{aligned}$$

where the last inequality follows from inequality 19. This contradicts the optimality of $\mathcal{F}'$. □

5.2.4 Randomized 2-parameter Auction

The mechanism presented above has an ϵ-Nash equilibrium corresponding to every optimal solution to LP 16, but we cannot guarantee that there are no other ϵ-Nash equilibria. As a result, it was not possible to bound the total payoff to the edges. In this section, we consider a slightly modified mechanism in which we add a small random payment, as in the mechanism RandomPath. We prove that, with this modification, it is possible to bound the total payment. Our mechanism uses Auction 2-Parameter as a subroutine and therefore is not implementable in polynomial-time.

Randomized 2-parameter Auction: As before, each edge i bids a pair $\tilde{a}_i = (\tilde{c}_i, \tilde{u}_i)$ where $\tilde{c}_i$ is interpreted as i's reported cost, and $\tilde{u}_i$ is interpreted as i's demanded profit.

1. **The 2-parameter auction.** This step is conducted exactly as in Auction 2-Parameter by solving IP 12 to select the minimum price solution.

2. **Rejection.** If for any edge not in the selected solution $\tilde{u}_i \neq 0$, reject the bid profile. No edge is selected and no flow is sent.[3]

3. **The randomized audit.** For edges on a random source-destination path, the payoff is based entirely on the $\tilde{c}_i$ component of the bid, and is constructed as in the mechanism RandomPath. The parameters α, τ, and B are as defined in the mechanism RandomPath. If an edge has true cost c_i and bids $(\tilde{c}_i, \tilde{u}_i)$, its expected payoff from this component is $g(\tilde{c}_i) = \tau[c_i\tilde{c}_i - \frac{\tilde{c}_i^2}{2}]$. The exact form of the payoff was derived in the proof of Lemma 5.1.

The audit component of the auction encourages edges to submit bid vectors in which their costs are nearly truthful. The first two steps of the auction help guarantee that the demand profits form a nearly feasible solution to LP 16. These facts allow us to derive bounds on the expected payment as stated in the following theorem.

Theorem 5.6. *The total price paid by the auctioneer in the randomized 2-parameter auction is at most*

$$\left[\sum_{j=1}^{r} jp_j c(F_{r+1})\right] - rc(\mathcal{F}) + nr\sqrt{2\epsilon/\tau} + 3\alpha n^2 B.$$

The result of Theorem 5.6 stands in an interesting relation to that of Theorem 5.1. We do not achieve the intuitively appealing bound of the expectation of the bounds on the known demand auction in Section 5.1, *i.e.*, $E_j[P_j] = \sum_{j=1}^{r} jp_j(c(F_{j+1}) - c(F_j))$ but instead we achieve $\sum_{j=1}^{r} rp_j(c(F_{r+1})(j/r) - c(F_j))$. In other words, the external multiplier j is replaced by r (a larger quantity), while in the first term the quantity $c(F_{j+1})$ is replaced by $c(F_{r+1})(j/r)$, which can also be larger because the cost of j units of flow is a convex function of j. Our Theorem 5.1 is therefore weaker in two important respects than Theorem 2, but it does have a similar overall structure.

To prove Theorem 5.6, we first show that *all* edges are nearly truthful about their costs in equilibrium:

Lemma 5.6. *Let $\tilde{a} = (\tilde{c}, \tilde{u})$ be an ϵ-Nash equilibrium of the randomized 2-parameter auction. Then, for all i,*

$$c_i - \sqrt{2\epsilon/\tau} \leq \tilde{c}_i \leq c_i + \sqrt{2\epsilon/\tau}$$

Proof. We argue that player i can always do better by bidding his true cost; the bounds follow from the ϵ-Nash equilibrium condition and the expected-payoff graph of the randomized path audit. Let ρ_i be the probability of i being included in the lowest price solution in the ϵ-Nash equilibrium $\tilde{a}$. If $\rho_i = 0$, then i's entire expected payoff is due to her expectation of winning in the randomized path audit, and the bounds on $\tilde{c}_i$ follow directly. The same argument holds if $\rho_i > 0$ but i receives a negative expected payoff from the 2-parameter auction (because her bid $\tilde{c}_i$ was too low).

Now, suppose $\rho_i > 0$, and, further, i receives a positive payoff from the 2-parameter auction in the ϵ-Nash equilibrium. Consider the strategy $a'_i = (c_i, u'_i)$ with $u'_i = \tilde{u}_i + \rho_i[\tilde{c}_i - c_i]$. ($i$

[3] This step ensures that, for all edges i not in the winning solution, $\tilde{u}_i$ is 0. Alternatively, we could ensure that these $\tilde{u}_i$ are close to zero (which is enough for our purposes) by charging a small tax to all bidders who submit a positive $\tilde{u}_i$ component of the bid.

received a non-negative profit under $\tilde{a}$, so it follows that u'_i is non-negative.) Let $\mathcal{F}$ be the solution chosen in the 2-parameter part of the mechanism when the bids are $\tilde{a}$. Note that if i were to deviate from $\tilde{a}_i$ to a'_i, the price of $\mathcal{F}$ would not change: the change in the utility component would exactly cancel the change in the cost component. Also, for any other flow $\mathcal{F}'$ that did not use i, the price of $\mathcal{F}'$ would not change with i's deviation; thus, using the consistency of the tie-breaking rule, $\mathcal{F}'$ would not be chosen above $\mathcal{F}$. Thus, we conclude that i remains in the winning solution (which need not be $\mathcal{F}$) under the bids a'_i.

Next, observe that i's expected payoff from the 2-parameter auction (with bid a'_i) is u'_i, because i bids her cost truthfully and is in the winning solution. This is exactly the same as i's payoff $\rho_i[\tilde{c}_i - c_i] + \tilde{u}_i$ from the 2-parameter auction in the ϵ-Nash equilibrium $\tilde{a}$.

To prove the bounds on $\tilde{c}_i$, we compare i's payoff from the randomized part of the mechanism with bids $\tilde{a}_i$ and a'_i. The bounds follow directly from the form of the randomized audit payoffs. □

Using the fact that the costs are nearly truthful, we can show that the utility values are an (almost) feasible solution to LP 16, and hence, derive the following bound on the total payment. We use the linear programming formulation given in LP 16, only this time we define the LP with respect to the *reported* costs rather than the true costs. Rewriting, we get

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{n} u_i \qquad (20) \\ \text{subject to} \quad & \forall \text{ feasible solutions } \mathcal{F}: \sum_{i \notin \mathcal{F}} u_i \leq \tilde{c}(\mathcal{F}) - \tilde{c}(\mathcal{F}^*) \\ & \forall\, 1 \leq i \leq n: \; u_i \geq 0 \end{aligned}$$

where $\mathcal{F}^*$ is now the minimum cost solution with respect to costs $\tilde{c}$.

Let $\tilde{a} = (\tilde{c}, \tilde{u})$ be any ϵ-Nash equilibrium of the Randomized 2-Parameter Auction. Let $\mathcal{F}^*$ be a minimum cost solution with respect to costs $\tilde{c}$, and let F_{r+1} be a minimum cost $(r+1)$-flow with respect to costs $\tilde{c}$.

Lemma 5.7. *Let u be any feasible solution to LP 20. Then for bids $\tilde{a} = \{(\tilde{c}_i, u^*_i)\}$, the minimum price solution $\mathcal{F}$ satisfies*

$$\tilde{a}(\mathcal{F}) \leq \tilde{c}(F_{r+1}) \sum_{j=1}^{r} j p_j \quad - \quad r\tilde{c}(\mathcal{F}^*).$$

Proof. Throughout this proof, minimum cost refers to minimum cost with respect to cost vector $\tilde{c}$. Consider an integral $(r+1)$-flow F_{r+1} minimizing $\tilde{c}(F_{r+1})$. Then F_{r+1} is a minimum cost $(r+1)$-flow and consists of $(r+1)$ disjoint paths $\{P_1, \cdots, P_{r+1}\}$ from s to t. For each $k \in \{1, 2, \cdots, r, r+1\}$, define $F_r^{-k} = F_{r+1} \backslash P_k$, that is, the r-flow obtained by dropping the k'th path. Extend F_r^{-k} to a collection of flows $\mathcal{F}^{-k} = \{F_1^{-k}, F_2^{-k}, \cdots, F_r^{-k}\}$, where F_j^{-k} consists of the j lowest-priced paths in F_r^{-k}. Noting that F_j^{-k} has cost at most $\frac{j}{r}$ that of F_r^{-k},

$$\tilde{c}(\mathcal{F}^{-k}) \leq \tilde{c}(F_r^{-k}) \sum_{j=1}^{r} p_j \frac{j}{r}.$$

Now, summing the inequality corresponding to $\mathcal{F}^{-k}$ over all k, we get:

$$\sum_{k=1}^{r+1} \sum_{i \notin F_r^{-k}} u_i \ \ \leq \ \ \sum_{k=1}^{r+1} \left(\tilde{c}(F_r^{-k}) \sum_{j=1}^{r} p_j \frac{j}{r} - \tilde{c}(\mathcal{F}^*) \right).$$

Note that the left hand side includes each element of F_{r+1} exactly r times. Similarly, the flows F_r^{-k} in the right hand side cover F_{r+1} exactly r times. Thus,

$$(r+1)\sum_{i=1}^{n} u_i - \sum_{i \notin F_r^{-k}} u_i \leq r\tilde{c}(F_{r+1}) \sum_{j=1}^{r} p_j \frac{j}{r} - (r+1)\tilde{c}(\mathcal{F}^*),$$

and so,

$$\begin{aligned} \tilde{a}(\mathcal{F}) &\leq \tilde{a}(\mathcal{F}^*) \\ &\leq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i \\ &\leq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i + r\sum_{i=1}^{n} u_i - \sum_{i \notin F_r^{-k}} u_i \\ &\leq \tilde{c}(F_{r+1}) \sum_{j=1}^{r} j p_j - r\tilde{c}(\mathcal{F}^*). \end{aligned}$$

□

Now, to prove our main theorem, we simply need to prove that the bid profile is a feasible solution of the linear program.

Proof of Theorem 5.6. Similar to Theorem 5.3, the total probability that the mechanism picks a path in the randomized audit is bounded by $n\alpha$. Any such path can have at most n edges on it, each with price at most B. Thus, the expected payment for using one of these paths is at most $\alpha n^2 B$. Similarly, we can bound the bonus $f_e(\mathbf{b})$ paid to any edge e: $f_e(\mathbf{b}) \leq n[\alpha B + \tau n B^2]$. This is always less than $2\alpha n B$.

Now we show that vector $\tilde{u}$ of demanded profits in bid profile $\tilde{a}$ is a feasible solution to LP 20. By assumption, for all losers, the demanded profit is zero. Therefore,

$$\tilde{a}(\mathcal{F}) = \tilde{c}(\mathcal{F}) + \sum_{i=1}^{n} u_i \geq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i.$$

Consider any solution $\mathcal{F}'$ and note that

$$\tilde{c}(\mathcal{F}') + \sum_{i \in \mathcal{F}'} u_i = \tilde{a}(\mathcal{F}') \geq \tilde{a}(\mathcal{F}) \geq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i,$$

and so the constraint corresponding to $\mathcal{F}'$ is satisfied. Therefore $\tilde{u}$ is a feasible solution. Since the $\tilde{u}$ satisfy the conditions of Lemma 5.7, noting that for any set of edges F, $c(F) - n\sqrt{2\epsilon/\tau} \leq \tilde{c}(F) \leq c(F) + n\sqrt{2\epsilon/\tau}$, we can apply Lemma 5.7 to get the result. □

6 Conclusion

In this paper, we showed that first-price auctions entail potentially lower payments than VCG mechanisms. In particular, the results in Section 5 show that for a fixed k-unit path auction, the upper bound on total payments in strong ϵ-equilibria is almost the same as the lower bound on the VCG mechanism payments; further, the bounds are the same in the limit as ϵ tends to 0. It is apparent from the simple examples in Section 3 and results in Archer and Tardos [2] and Elkind *et al.* [12] that the VCG mechanism will often require payments considerably higher than this lower bound (and hence, considerably higher than the strong ϵ-equilibria of the first-price auction).

In Section 5.2.2 and 5.2.3 we considered a model in which the demand is a variable with a known distribution, and we need to select paths *ex ante*. We showed that a simple first-price auction may not even have an ϵ-Nash equilibrium. However, we proved that a variant of the auction with 2-parameter bids induces a surplus-sharing game with a nonempty core, and that every core element can be perturbed slightly to get an ϵ-Nash equilibrium. We also proved a bound on the total payment to links in a core allocation, which suggests that in this domain too it may be possible to prove that the VCG mechanism has higher expected payments.

This leads us to a comparison between first-price and VCG path auctions similar to the comparison between the cost-sharing mechanisms considered by Young [39]. First-price auctions entail potentially lower payments, and have greater collusion resistance than VCG mechanisms. However, they suffer from one major drawback, in that the solution concept (strong ϵ-Nash equilibrium) requires agents to know all costs, and coordinate on the choice of equilibrium. This is much more demanding than the dominant-strategy mechanisms and can lead to inefficiency in practice. Thus, the auction models analyzed here are not completely satisfying, as there is no mechanism prescribed for the agents bids' to reach equilibrium. This is true even for the weaker concept of ϵ-Nash equilibrium.

However, the results in this paper shed new light on the *functions* of overpayment in VCG mechanisms. We can identify three distinct functions of overpayment:

1. Cheaper paths have a competitive advantage and can thus command a surplus.
2. The surplus paid to links eliminates the need for negotiation between links, leading to a simple mechanism without delays or expensive reasoning.
3. The surplus eliminates the externalities of one agent's strategy on other agents, leading to a mechanism that is fair in the sense that uninformed agents can do as well as informed agents.

The first source of overpayment is common to the first-price auction and the VCG mechanism. However, our results show that for path auctions, the VCG mechanism often winds up paying a premium for functions 2 and 3. (In contrast, for single-item auctions, the first-price auction always pays as much in the worst case as the VCG mechanism.)

This premium can be viewed as the "cost of implementation" of the dominant-strategy mechanism, particularly in situations in which this form of fairness is not compelling. We believe that a promising direction for future research is to find bargaining mechanisms to enable the bidders to converge to an equilibrium. When the edges all know each others' costs, an n-party bargaining protocol, such as the one in the Krishna and Serrano [28], could be used. When the edge costs are unknown initially but become revealed eventually, the approach of

Kalai and Kalai [23] for implementing a cooperative solution concept by a noncooperative game may be used. When there is uncertainty, the situation is more complex. Such a mechanism may be subsidized; for example, the links may be given an additional payment that decays with time, to incentivize them to quickly reach an agreement. As long as the subsidy is smaller than the VCG premium, it may be a better alternative.

Acknowledgments. We would like to thank Eva Tárdos for suggesting the bound $c(2) - c(1)$ on the payment in a 1-unit auction, which motivated us to pursue this work. We also thank Joan Feigenbaum for helpful comments on an earlier version and Sergei Izmalkov and Mohammad Mahdian for many fruitful discussions.

References

[1] A. Archer and E. Tardos. Frugal path mechanisms. In *Proceedings of 13th ACM Symposium on Discrete Algorithms*, pages 991–999, 2002.

[2] A. Archer and E. Tardos. Frugal path mechanisms. *ACM Trans. Algorithms*, 3(1):1–22, 2007.

[3] R. Aumann. Acceptable points in general cooperative n-person games. In A. Tucker and R. Luce, editors, *Contribution to the Theory of Games: Annals of Mathematical Studies*, volume 40. Princeton University Press, 1959.

[4] N. Chen, E. Elkind, N. Gravin, and F. Petrov. Frugal mechanism design via spectral techniques. In *FOCS '10: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 472–482, Washington, DC, USA, 2010. IEEE Computer Society.

[5] N. Chen, A. Ghosh, and N. Lambert. Social lending. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 335–344, New York, NY, USA, 2009. ACM.

[6] N. Chen and A. R. Karlin. Cheap labor can be expensive. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 707–715, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.

[7] E. Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

[8] I. Curiel. *Cooperative Game Theory and Applications.* Kluwer Academic Publishers, 1997.

[9] A. Czumaj and A. Ronen. On the expected payment of mechanisms for task allocation. In *Proceedings of 23rd ACM Symposium on Principles of Distributed Computing*, pages 98–106, 2004.

[10] Y. Du, R. Sami, and Y. Shi. Path auctions with multiple edge ownership. *Theoretical Computer Science*, 411(1):293–300, 2010.

[11] E. Elkind, L. A. Goldberg, and P. W. Goldberg. Frugality ratios and improved truthful mechanisms for vertex cover. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 336–345, New York, NY, USA, 2007. ACM.

[12] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *Proceedings of 15th ACM Symposium on Discrete Algorithms*, pages 701–709, 2004.

[13] G. Ercal, R. Izhak-Ratzin, R. Majumdar, and A. Meyerson. Frugal routing on wireless ad-hoc networks. In B. Monien and U.-P. Schroeder, editors, *Algorithmic Game Theory*, volume 4997 of *Lecture Notes in Computer Science*, pages 133–144. Springer Berlin-Heidelberg, 2008.

[14] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A bgp-based mechanism for lowest-cost routing. In *Proceedings of 21st Symposium on Principles of Distributed Computing*, pages 173–182, 2002.

[15] M. Garey and D. Johnson. *Computers and Intractability.* Freeman Press, 1979.

[16] R. Garg, V. Kumar, A. Rudra, and A. Verma. Coalition games on graphs: core structure, substitutes and frugality. In *Proceedings of 4th ACM Conference on Electronic Commerce*, pages 248–249, 2003.

[17] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[18] J. Hershberger and S. Suri. Vickrey prices and shortest paths: what is an edge worth? In *Proceedings of 42nd IEEE Symposium on the Foundations of Computer Science*, pages 129–140, 2001.

[19] C.-C. Huang, M.-Y. Kao, X.-Y. Li, and W. Wang. Using nash implementation to achieve better frugality ratios. In *ISAAC'07: Proceedings of the 18th international conference on Algorithms and computation*, pages 377–389, Berlin, Heidelberg, 2007. Springer-Verlag.

[20] A. Iwasaki, D. Kempe, Y. Saito, M. Salek, and M. Yokoo. False-name-proof mechanisms for hiring a team. In *WINE'07: Proceedings of the 3rd international conference on Internet and network economics*, pages 245–256, Berlin, Heidelberg, 2007. Springer-Verlag.

[21] M. Jackson. A Crash Course in Implementation Theory. *Social Choice and Welfare*, 18:655–708, 2001.

[22] M. O. Jackson, L. K. Simon, J. M. Swinkels, and W. R. Zame. Communication and equilibrium in discontinuous games of incomplete information. *Econometrica*, 70(5):1711–1740, September 2002.

[23] A. T. Kalai and E. Kalai. A cooperative value for bayesian games. *unpublished manuscript.*

[24] E. Kalai, A. Postlewaite, and J. Roberts. A group incentive compatible mechanism yielding core allocations. *Journal of Economic Theory*, 20:13–22, 1979.

[25] D. R. Karger and E. Nikolova. On the expected VCG overpayment in large networks. In *Decision and Control, 2006 45th IEEE Conference on*, pages 2831–2836, Dec. 2006.

[26] A. R. Karlin, D. Kempe, and T. Tamir. Beyond vcg: Frugality of truthful mechanisms. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 615–626, Washington, DC, USA, 2005. IEEE Computer Society.

[27] D. Kempe, M. Salek, and C. Moore. Frugal and truthful auctions for vertex covers, flows, and cuts. In *FOCS '10: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 472–482, Washington, DC, USA, 2010. IEEE Computer Society.

[28] V. Krishna and R. Serrano. Multilateral bargaining. *Review of Economic Studies*, 63:61–80, 1998.

[29] X.-Y. Li, Y. Wu, P. Xu, G. Chen, and M. Li. Hidden information and actions in multi-hop wireless ad hoc networks. In *MobiHoc '08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pages 283–292, New York, NY, USA, 2008. ACM.

[30] M. Mihail, C. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. In *Proceedings of the 44rd IEEE Symposium on Foundations of Computer Science*, 2003.

[31] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[32] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[33] A. Schrijver. *Combinatorial Optimization*. Springer, 2003.

[34] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *20th Annual Symposium on Theoretical Aspects of Computer Science*, pages 608–619, 2003.

[35] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[36] W. Wang, S. Eidenbenz, Y. Wang, and X.-Y. Li. Ours: optimal unicast routing systems in non-cooperative wireless networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 402–413, New York, NY, USA, 2006. ACM.

[37] Y. Wang, W. Wang, and T. Dahlberg. Truthful routing for wireless hybrid networks. In *Global Telecommunications Conference, 2005. GLOBECOM '05. IEEE*, volume 6, pages 5 pp.–3465, Dec. 2005.

[38] Q. Yan. On the price of truthfulness in path auctions. In *WINE'07: Proceedings of the 3rd international conference on Internet and network economics*, pages 584–589, Berlin, Heidelberg, 2007. Springer-Verlag.

[39] H. Young. Cost allocation, demand revelation, and core implementation. *Mathematical Social Sciences*, 36:213–228, 1998.