

First-Price Path Auctions

Nicole Immorlica*
MIT CS&AI Laboratory
Cambridge, MA 02139
nickle@csail.mit.edu

Evdokia Nikolova‡
MIT CS&AI Laboratory
Cambridge, MA 02139
enikolova@csail.mit.edu

David Karger†
MIT CS&AI Laboratory
Cambridge, MA 02139
karger@csail.mit.edu

Rahul Sami§
MIT CS&AI Laboratory
Cambridge, MA 02139
sami@csail.mit.edu

ABSTRACT

We study first-price auction mechanisms for auctioning flow between given nodes in a graph. A *first-price auction* is any auction in which links on winning paths are paid their bid amount; the designer has flexibility in specifying remaining details. We assume edges are independent agents with fixed capacities and costs, and their objective is to maximize their profit. We characterize all *strong ϵ -Nash equilibria* of a first-price auction, and show that the total payment is never significantly more than, and often less than, the well known dominant strategy Vickrey-Clark-Groves mechanism. We then present a randomized version of the first-price auction for which the equilibrium condition can be relaxed to ϵ -Nash equilibrium. We next consider a model in which the amount of demand is uncertain, but its probability distribution is known. For this model, we show that a simple *ex ante* first-price auction may not have any ϵ -Nash equilibria. We then present a modified mechanism with 2-parameter bids which does have an ϵ -Nash equilibrium.

Categories and Subject Descriptors

F.m [Theory of Computation]: Miscellaneous

General Terms

Economics, Theory

Keywords

path auctions, first-price auctions

1. INTRODUCTION

*Supported by NSF doctoral fellowship.

†Supported in part by NSF grants ANI-0225660 and ITR-0219018.

‡Supported in part by NSF grants ANI-0225660 and ITR-0219018.

§Supported NSF grant ITR-0219018.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.

Copyright 2005 ACM 1-59593-049-3/05/0006 ...\$5.00.

In this paper, we study variants of the *path auction* problem. The basic problem can be described as follows: We are given a directed graph G with two distinguished vertices s and t . Each link in the graph is a self-interested agent whom we assume to be risk-neutral. All links have capacity 1, but each link i also has a cost c_i that is known only to the link itself. A customer wants to buy 1 (or more generally, some integer k) paths from s to t . For this, she holds an auction in which each link can bid; the auction should end with the customer announcing a path, as well as the payments to each link. The questions we are chiefly concerned with are: (1) What is the form of bids, and how are the path and payments selected? (2) How much does the customer end up paying, given that the links have an informational advantage (the customer does not know the true link costs)?

Previous work on path auctions has studied the Vickrey-Clarke-Groves (VCG) mechanism [17, 12, 9, 2, 8]. Roughly speaking, the VCG mechanism pays each edge on a winning path an amount equal to the highest bid with which it could still have won, all other bids being unchanged. The VCG mechanism has the attractive property that each link's dominant strategy is to bid exactly its cost. Thus, no bargaining or communication between bidders is required to stabilize on bids. Also the buyer does end up using the path of lowest true cost, which can be seen as optimizing social utility.

On the negative side, the VCG mechanism can lead to the customer paying far more than the true cost of the cheapest path. The tendency to overpay is exaggerated in path auctions (as compared to simple auctions) because a bonus needs to be paid to every agent on the path. Thus, the payment to the lowest-cost path may even greatly exceed the cost of the second-cheapest path. For example, in Figure 1, VCG selects the bottom path and pays 4 to it, even though the alternate path has cost 3. Archer and Tardos showed that a more general class of dominant strategy mechanisms can be forced to make arbitrarily high overpayments [2]. Their result was strengthened to hold for every truthful mechanism by Elkind, Sahai and Steiglitz [8].




Figure 1: Any ϵ -Nash equilibrium selects the lower path and pays $3 - \epsilon$ while VCG pays 4 for it.

In this paper, we are interested in finding techniques to rein in

the cost to the consumer, even when the information is completely asymmetric—the links know the customer’s valuation, but the customer does not know the links’ valuation. If we restrict ourselves to dominant-strategy mechanisms, we cannot hope to do better than the VCG mechanism. In this paper, we instead consider variants on *first-price* auctions and less restrictive solution concepts.¹

First-price auctions open the possibility of paying less than VCG auctions, but they do so by sacrificing valuable properties of the VCG mechanism. In particular, in a first-price auction, a risk-neutral edge may have incentive to lie, bidding a price higher than its cost. Also, in the absence of a dominant strategy, it may be necessary for bidders to communicate and bargain to achieve a stable set of bids.

1.1 Our Results

We begin by exploring the sets of bids that are stable under a first-price auction mechanism. The most natural solution concept is that of a Nash equilibrium. 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 for us. Unfortunately, we will not necessarily have a Nash Equilibrium in pure strategies, as the following simple example shows. Consider a network of two parallel links, one of cost 2 and another of cost 1. Also assume that ties are broken deterministically by assigning the flow to the link with cost 2. In this case, the lower-cost edge would bid less than 2; however, for any bid $2 - \epsilon$, it could always do better by increasing its bid by a further $\epsilon/2$. Hence there is no pure Nash equilibrium in this case.

This motivates us to use the solution concept of ϵ -Nash equilibrium, in which no player can deviate in a way that improves his payoff by at least ϵ . Unfortunately, there is a drawback to this solution concept as well. In Figure 2, we see that the winning path may have a price higher than the cost of the best competitor. This defeats our goal of reducing customer overpayment. We might argue that this solution would not be sustained in practice, since the edges on the second lowest-cost path are likely to each reduce their price. This leads us to explore, in Section 3, the concept of a *strong* ϵ -Nash equilibrium, in which there is no *group* of agents who can deviate in a way that improves the payoff of each member by at least ϵ . We prove that a strong ϵ -Nash equilibrium always exists for any $\epsilon > 0$. We then prove an upper bound on the payment of any such equilibrium and show that the payment is essentially not more than that of the corresponding VCG payment, and often it is much less as shown by Figure 1.

Although strong ϵ -Nash equilibria may solve some of the overpayment problem, we cannot guarantee that bidders will reach one. In particular, in the absence of knowledge about other bidders costs, neither losing bidder in the example of Figure 2 may be willing to “blink first” and lower the price. Thus, in Section 4, we present a modified, randomized, first-price auction that explicitly drives the first-price auction towards a strong ϵ -Nash equilibrium.

Another drawback of first-price auctions is that, unlike the VCG mechanism, an edge’s preferred bid may depend on the demand (e.g., if demand is high, an edge can bid higher and still hope to be needed). It is unreasonable to expect edges to delay setting prices until demands are made clear. Thus, in Section 5 we consider a model in which bidders set prices according to a *distribution* of possible demands. We show that, in this model, a simple first-price auction may not have an ϵ -Nash equilibrium. However, we design a

¹By “first-price auction” we refer to any auction in which the links on the winning path (or paths) are paid their bid amount. The designer still has considerable flexibility in designing the details of the auction mechanism.

first-price mechanism involving *two-parameter* bids that *does* have an ϵ -Nash equilibrium. We then sketch a mechanism that combines this two-parameter mechanism with the randomized mechanism of Section 4. For this combined mechanism, we can characterize the set of all ϵ -Nash equilibria, and thereby prove a bound on the total payment in any ϵ -Nash equilibrium.

In order to maintain continuity, most proofs have been deferred to the Appendix.




Figure 2: Costs (left) and Prices (right) in an ϵ -Nash equilibrium. The bottom edge wins and the price is higher than the cost of the second best path.

1.2 Related Work

Path auctions are an instance of the more general class of combinatorial auctions, in which buyers bid for different collections of goods. In path auctions, sellers (in our case, graph edges) bid to attract consumer flow and consumers seek to buy a path of edges of lowest price between a specified source and destination. Finding the winners in general combinatorial auctions is NP-hard [17, 1], for this reason researchers often add restricting assumptions such as symmetric bidders, etc. Path auctions provide one such simple-structured form of combinatorial auctions, which arises naturally in network routing and more generally in any problems with an underlying network structure, such as task allocation to teams of agents.

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 path auction we consider is similar to the game introduced by Young [19] in the context of cost-sharing. For the random demand path auction introduced in section 5.2, we use techniques based on Curiel [5] to show the existence of the core. We also note that Kalai *et al.* [14] 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; however, the method in [14] is more complex and communication-intensive, and in our case it would essentially require each bidder to report an entire flow.

There has also been some previous work on non-dominant strategy mechanisms for path auctions. Elkind *et al.* [8] present and analyze an optimal Bayes-Nash mechanism. Garg *et al.* [7] study use the core concept from cooperative game theory to bound the payments of VCG mechanisms for a large class of problems that includes path auctions. Czumaj and Ronen [6] propose a mechanism which combines dominant and non-dominant strategy mechanisms, 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.”

2. PROBLEM STATEMENT

In the *path auction game*, there is a network G of strategic links, each with a privately-known true cost. All links have unit capacity. A customer wants to buy routes from a source s to a sink t in the

network to guarantee that her integral amount of demand k can be routed. In order to do this, she defines a set of rules, or *mechanism*, that elicits bids from each agent and then allocates flow to each agent in a way that satisfies some natural incentive properties.

One plausible mechanism for this problem is the Vickrey-Clark-Groves (VCG) mechanism [18, 4, 11]. This mechanism is *truthful dominant strategy* or *strategyproof*, i.e. the strategically best bid for an agent is his true cost, independent of others' bids. Thus the bids solicited by the mechanism are in fact the true cost of the agents in an equilibrium. This enables the mechanism to allocate flow to the lowest *true cost* k -flow, a socially desirable goal in many settings. However, in order to guarantee that this allocation rule is truthful dominant strategy, the mechanism must pay a (possibly large) premium to all edges on the selected k -flow. One side effect of dominant strategies is that all bargaining between the strategic agents (links, in our case) is eliminated, and the overpayment to edges on the selected k -flow in the VCG mechanism can be thought of as a side-effect of this fact.

We analyze approaches to reducing the total payment by using a weaker solution concept of a pure strategy equilibrium, the *strong ϵ -Nash equilibrium* first introduced by Aumann [3] and used by Young [19].

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

A strong ϵ -Nash equilibrium for a game is a set of strategies, one for each player, such that no group of players can deviate in a way that improves the payoff of each member by at least ϵ .

In particular, we show that in our models, for any strong ϵ -Nash equilibrium set of bids, there is another strong ϵ -Nash equilibrium set of bids with the same allocation and payment scheme in which each agent bids within ϵ of his true cost unless he is allocated flow (in expectation), and he never bids below his true cost.

Our mechanism is a simple first-price auction. It elicits bids from each agent, computes the cheapest k -flow according to the bids, and then allocates the demand to that k -flow. We further assume that we have a deterministic tie-breaking rule so that if there is more than one cheapest k -flow, we take the lexicographically first integral one.

We consider two specific path auction games. In the *deterministic path auction game*, the user first announces k , his total demand. Then the edges announce bids and the user runs a first price auction to buy the necessary flow. It is easy to imagine that the assumptions of this model might be unrealistic in practice. Does a user really know his total demand at the time he begins the auction? In our second model, the *random path auction game*, the user announces a probability distribution on k . Then the edges announce bids. Finally, the user draws k according to this distribution and buys flow accordingly. In the rest of this paper, we analyze upper and lower bounds on the overpayments in strong ϵ -Nash equilibria for each of these games.

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 minimum weight integral k -flow² in G with respect to edge weights \mathbf{w} (if there is more than one minimum weight k -flow in G with respect to \mathbf{w} , let $F_{\mathbf{w}}(k, G)$ denote the set of edges in the unique k -flow that wins the deterministic tie-breaking rule of the mechanism). 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

²The weight of this flow is equal to the weight of the minimum weight k -flow, i.e., requiring integrality doesn't change the value of the optimal solution.

vector \mathbf{b} . Finally, for any flow or edge set F , we define $W_{\mathbf{w}}(F)$ to be the weight of F with respect to edge weights \mathbf{w} . We say $W_{\mathbf{c}}(F)$ is the cost of flow F and $W_{\mathbf{b}}(F)$ to be the price of flow F 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 $C(F)$ instead of $W_{\mathbf{c}}(F)$, as well as $C(k)$ for the (cost of the) lowest cost k -flow. Finally, we denote the number of agents, or edges in G , by n .

3. DETERMINISTIC PATH AUCTION GAME

Recall that in the *deterministic path auction game*, the user first announces k , his total demand. Then the edges announce bids and the user runs a first price auction to buy the necessary flow. We would like to analyze the payment properties of this mechanism. First, we prove that this mechanism has a strong ϵ -Nash equilibrium.

THEOREM 1. *Any deterministic k -unit first price auction has a strong ϵ -Nash equilibrium.*

PROOF. We construct a strong ϵ -Nash equilibrium as follows. Set the initial bid vector $\mathbf{b}^i = \mathbf{c}$, i.e. each edge bids its true cost initially. Order the edges in the graph in an arbitrary way. For each edge e in this order, if e is part of the current lowest price k -unit flow $F_{\mathbf{b}}(k, G)$, let e raise its bid until $W_{\mathbf{b}^f}(F_{\mathbf{b}^f}(k, G)) \geq W_{\mathbf{b}}(F_{\mathbf{b}}(k, G - \{e\})) - \epsilon/2$ (where $G - \{e\}$ denotes the graph G with edge e removed). Otherwise let e 's bid remain unchanged. Call the final bid vector \mathbf{b}^f .

We claim \mathbf{b}^f is a strong ϵ -Nash equilibrium for the deterministic k -unit first price auction. To show this, suppose the contrary, i.e., there is a coalition S of edges in which each edge can improve its payoff by at least ϵ by changing its bid. Note that for any bid vector constructed during this process, the auction always selects the same k -flow. Therefore, the edges which are not on the winning flow in \mathbf{b}^f are bidding their true cost and cannot bid lower. Furthermore, the edges which are on the winning flow will get smaller payoff if they decrease their bid. Therefore no edge can benefit from lowering its bid. Thus, the edges in the coalition S can only raise their bids. Suppose the edges in $S \cap F_{\mathbf{b}^f}$ increase their bids by a total of x units and the remaining edges in the coalition increase their bids by a total of y units (note $x, y > 0$). Call the new bid vector \mathbf{b} . In order for all edges in S to increase their payoff, $S \subseteq F_{\mathbf{b}}$. Thus $W_{\mathbf{b}}(F_{\mathbf{b}^f}) = W_{\mathbf{b}^f}(F_{\mathbf{b}^f}) + x$ while $W_{\mathbf{b}}(F_{\mathbf{b}}) = W_{\mathbf{b}^f}(F_{\mathbf{b}}) + x + y$. But then $W_{\mathbf{b}}(F_{\mathbf{b}}) > W_{\mathbf{b}}(F_{\mathbf{b}^f})$ since $W_{\mathbf{b}^f}(F_{\mathbf{b}^f}) \leq W_{\mathbf{b}^f}(F_{\mathbf{b}})$ by optimality of $F_{\mathbf{b}^f}$. This contradicts the optimality of $F_{\mathbf{b}}$. \square

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, we could show that without loss of generality the bid of this path should be close to its true cost. This proof idea easily extends to k -flows in general graphs as can be seen below.

THEOREM 2. *The total payment of the deterministic k -unit first price auction in a strong ϵ -Nash equilibrium is at most*

$$k[C(F_c(k+1)) - C(F_c(k))] + kn\epsilon,$$

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

PROOF. Fix a strong ϵ -Nash equilibrium vector of bids \mathbf{b} and define edge sets

$$\begin{aligned} E_+ &= \{e \in F_c(k+1) - F_b(k)\} \\ E_o &= \{e \in F_c(k+1) \cap F_b(k)\} \\ E_- &= \{e \in F_b(k) - F_c(k+1)\} \end{aligned}$$

E_+ is the subset of edges on an augmenting path that are not in the original flow $F_b(k)$. We show that without loss of generality we may assume that these edges are bidding close to their true cost. To show this, consider a bid vector \mathbf{b}' such that

$$b'_i = \begin{cases} \min\{b_i, c_i + \epsilon\} & \text{for } i \in E_+, \\ b_i & \text{for } i \notin E_+. \end{cases}$$

We want to argue that $W_{\mathbf{b}'}(F_{\mathbf{b}'}(k)) = W_{\mathbf{b}}(F_{\mathbf{b}}(k))$. First we show $F_{\mathbf{b}}(k) = F_{\mathbf{b}'}(k)$. Suppose not. Let $E'_+ = E_+ \cap F_{\mathbf{b}'}(k)$ be the set of edges in the new lowest price flow that are also in E_+ . We have only changed the bids of the edges in E_+ , so if E'_+ is empty then $F_{\mathbf{b}}(k) = F_{\mathbf{b}'}(k)$ (this assumes some consistency properties of the tie-breaking rule). If E'_+ is nonempty, then we can consider a bid vector \mathbf{b}'' constructed from \mathbf{b} in which we only decrease bids of edges in E'_+ :

$$b''_i = \begin{cases} \min\{b_i, c_i + \epsilon\} & \text{for } i \in E'_+, \\ b_i & \text{for } i \notin E'_+. \end{cases}$$

Since by our assumption the winning flow has changed, we must have $b''_i = c_i + \epsilon < b_i$ for a non-empty subset E''_+ of E'_+ . Under this new bid vector, $W_{\mathbf{b}''}(F) \geq W_{\mathbf{b}'}(F)$ for any flow F since $\mathbf{b}''_i \leq \mathbf{b}'_i$ for all edges i . By construction, $W_{\mathbf{b}''}(F_{\mathbf{b}'}(k)) = W_{\mathbf{b}'}(F_{\mathbf{b}'}(k))$ and so, by the consistency of the tie-breaking rule, $F_{\mathbf{b}''}(k) = F_{\mathbf{b}'}(k)$. Thus, under the bid vector \mathbf{b} the set of edges E''_+ can form a coalition in which each member bids ϵ above its true cost and all members profit by ϵ . This contradicts the fact the \mathbf{b} was a strong ϵ -Nash equilibrium.

Now, noting that $W_{\mathbf{b}'}(F_{\mathbf{b}'}(k)) = W_{\mathbf{b}}(F_{\mathbf{b}}(k)) = W_{\mathbf{b}}(F_{\mathbf{b}}(k))$, it suffices to bound $W_{\mathbf{b}'}(F_{\mathbf{b}'}(k))$. Consider the (non-integral) flow $(k/(k+1))F_c(k+1)$, i.e. the flow which sends $k/(k+1)$ units of flow along the flow paths determined by $F_c(k+1)$. Since $F_{\mathbf{b}'}(k)$ is a lowest price k -flow,

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

This reduces to

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

which, solving for $W_{\mathbf{b}'}(E_o) + W_{\mathbf{b}'}(E_-)$, gives

$$W_{\mathbf{b}'}(F_{\mathbf{b}'}(k)) = W_{\mathbf{b}'}(E_o) + W_{\mathbf{b}'}(E_-) \quad (1)$$

$$\leq k(W_{\mathbf{b}'}(E_+) - W_{\mathbf{b}'}(E_-)) \quad (2)$$

$$\leq k(W_c(E_+) + n\epsilon - W_c(E_-)) \quad (3)$$

$$\leq k(W_c(F_c(k+1)) - W_c(F_{\mathbf{b}'}(k)) + n\epsilon) \quad (4)$$

$$\leq k(W_c(F_c(k+1)) - W_c(F_c(k)) + n\epsilon) \quad (5)$$

where 3 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 5 follows from the optimality of $F_c(k)$ with respect to \mathbf{c} . \square

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 $k(C(F_c(k+1)) - C(F_c(k)))$.

Finally, we emphasize two properties of our mechanism. The first property states 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. The second property states that the social welfare of the resulting solution is an additive approximation to the optimum social welfare. The proof of Theorem 3 is deferred to an extended version of this paper, due to space limitation.

THEOREM 3. *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 $k(C(F_c(k+1)) - C(F_c(k)))$.*

THEOREM 4. *In a strong ϵ -Nash equilibrium \mathbf{b} , $C(F_{\mathbf{b}}(k)) \leq C(F_c(k)) + \epsilon n$ (i.e. the strong ϵ -Nash equilibria of the first price auction are approximately efficient).*

4. IMPLEMENTATION IN ϵ -NASH

The simple first-price auction may have costly ϵ -Nash equilibria, as shown in the example in Figure 2. In Section 3 we used the ϵ -strong Nash solution concept to get around this problem. However, assuming that the bidders will reach an ϵ -strong Nash equilibrium is perhaps too strong an assumption: it requires extensive coordination between agents. In this section, we present a variant of the mechanism in which every ϵ -Nash equilibrium results in a low price.

One idea to achieve this is to pay a bonus to edges that increases as their bid decreases. This encourages edges to submit low bids. However, this has the side-effect of incentivizing edges to bid even below their true cost, as long as they remain off the winning path. This would make the bargaining problem that links must solve much more complex, as it would include bargains between off-path and on-path links. 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 be incentivized 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. If the bonuses and probabilities are small enough, then these bonus payments will not be very large, and we can prove a bound on the total payment of the mechanism similar to that in Theorem 2.

We achieve this result by making the mechanism outcome a *lottery* over paths instead of a single path: Every edge is on a selected path with at least a small probability, and edges off the shortest path are given an incentive to bid their true cost. This is known as *virtual implementation* in the economics literature (see, eg. Jackson [13]). We 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 is given in Figure 3. The mechanism starts by computing a collection of paths $\{P_e\}$. We discuss the computation of this collection in Section 4.1. The mechanism then invites 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

1. For each edge e , find P_e , a 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. Invite bids $\mathbf{b} = (b_1, \dots, b_e, \dots, b_n)$ 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 lowest price path. Call the selected path P^* . Pay each edge $e \in P^*$ its bid b_e .

5. 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

$$f_e^P(\mathbf{b}) = \alpha(B - b_e) + \tau b_e \sum_{j \in P} b_j - \tau \frac{b_e^2}{2}$$

(This payment is in addition to any payment edge e may get in step 4.)

Figure 3: Mechanism FP2. The parameters α and τ are selected to be small positive constants such that $\alpha < n^{-2}B^{-1}$ and $\tau < \alpha n^{-1}B^{-1}$.

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, which is *not* on the lowest-price path to bid its true cost. For simplicity, we present the mechanism and analysis for a single unit flow; the results can easily be extended to any constant $k > 1$ units of flow. First we note that ϵ -Nash equilibria exist in this mechanism; indeed the same construction as in Theorem 1 yields an ϵ -Nash equilibrium.

LEMMA 1. *For any cost vector \mathbf{c} and any $\epsilon > 0$, an ϵ -Nash equilibrium always exists.*

Given existence of ϵ -Nash equilibria, we can bound the bid of the edges not on the lowest true-cost path by examining their optimal bid given their total payoff function (sum of bonus and expected selection payoff). 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 $b_i = c_i$.

By evaluating the expected payoff of an off-path link in mechanism FP2, we can show:

LEMMA 2. *Let \mathbf{b} be an ϵ -Nash equilibrium bid vector in the mechanism FP2. Then, for any edge e not on the lowest-price path with bids \mathbf{b} , $b_e \in [c_e - \sqrt{2\epsilon/\tau}, c_e + \sqrt{2\epsilon/\tau}]$.*

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 mechanism FP2 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 $\epsilon \rightarrow 0$, the maximum expected payment in any Nash equilibrium is bounded by the same constant as before.

Observing that as $\epsilon \rightarrow 0$, $\sqrt{2\epsilon/\tau} \rightarrow 0$, we get the following result:

THEOREM 5. *Choose any $\alpha < n^{-2}B^{-1}, \tau < \alpha n^{-1}B^{-1}$. For these values of α and τ ,*

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

4.1 Computing the set of covering flows $\{P_e\}$

Recall that the mechanism FP2 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 [10].

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 [10]. 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. DISTRIBUTION ON DEMANDS

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 mechanism so that the equilibria chosen by the user 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. In particular, we assume that the edges know only of some *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 3 that a first price auction for k units of demand led to a payment of $P_k = k[C(F_c(k+1)) - C(F_c(k))]$. It is thus natural to hope that the same mechanism operating over *random* k is also stable, with expected payment $E_k[P_k]$. This turns out to be false—in fact, we show in Section 5.1 that the simple first-price auction mechanism described previously has *no* ϵ -Nash equilibria. 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. So we turn to more complex mechanisms.

We exhibit a mechanism involving *two parameter bids* that, unlike the single-parameter first-price mechanism, *does* have ϵ -Nash equilibria. Intuitively, a two-parameter mechanism gets around the problem of a single-parameter mechanism by letting the edges express their preferences over the entire price-probability space. It allows to an edge to bid a “price” that depends on its winning probability; this prevents the bidding cycles that occur with single-parameter bids. Furthermore, using an indifference-breaking technique similar to that of Section 4, we are able to restrict the set of equilibria to ones with bounded user payments. The bound is not quite the $E_k[P_k]$ we hoped to achieve, but does bear a clear resemblance to it.

5.1 No equilibrium with 1-parameter bidding

In this section, we analyze the scenario in which the demand is random with a known distribution, and the bidders (links) have to commit to a price before the demand is revealed, and there is deterministic tie-breaking. Subsequently, the demand is revealed to be k , and the k lowest-priced paths are picked. We show that there is in general no ϵ -Nash equilibrium in pure strategies for this game.

Consider a graph with four parallel links 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 $w + 50\epsilon < x + 42\epsilon = y + 12\epsilon = z$.

THEOREM 6. *There is no pure-strategy ϵ -Nash equilibrium for this game.*

The proof repeatedly uses the ϵ -Nash conditions to show that, at any bid vector, 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. The full proof will appear in an extended version of this paper.

5.2 Equilibrium with 2-parameter bidding

In section 5.1, we saw that when the demand is a random variable with a known distribution, a simple first-price auction may not have an ϵ -Nash equilibrium. In this section, we present a different auction model, in which agents’ bids are pairs of values, and show that it has a nonempty set of ϵ -Nash equilibrium.

To prove that the bidding game induced by this auction has a strong ϵ -Nash equilibrium, we construct a cooperative game model

of the auction. We show that the cooperative game has a nonempty *core*. We then connect the cooperative game to the actual bidding game, and show that the path auction has an ϵ -Nash equilibrium corresponding to any core element.

The model is 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 $k = 1, 2, \dots, r$. 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 $k \in \{1, \dots, r\}$. The agents’ bids are pairs of numbers: each agent i bids a pair $\tilde{a} = (\tilde{c}_i, \tilde{u}_i)$, where \tilde{c}_i is interpreted as i ’s reported cost, and \tilde{u}_i is interpreted as i ’s demanded profit.

The mechanism receives the bids, and announces flows F_1, F_2, \dots, F_r for each possible demand value. We call the collection $\mathcal{F} = \{F_1, F_2, \dots, F_r\}$ a candidate solution or simply a flow. We also identify a flow \mathcal{F} with the set of links in the union $F_1 \cup F_2 \cup \dots \cup F_r$, and say that $i \in \mathcal{F}$ if $i \in F_k$ for some k .

For each $i \in \mathcal{F}$, the mechanism calculates the probability that i is in the lowest-price flow, $\rho_i = \sum_{\{k|i \in F_k\}} p_k$. Later, the actual demand transpires; suppose that the demand turns out to be k . The mechanism uses the links in F_k to route the flow, and pays each link $i \in F_k$ a sum of $\tilde{c}_i + \frac{\tilde{u}_i}{\rho_i}$. Consider any link i selected in some flow. If $\tilde{c}_i = c_i$ (*i.e.*, if i bid its true cost), her expected profit would be \tilde{u}_i . Given the input bid pairs, the mechanism selects a set of flows F_1, F_2, \dots, F_r that minimizes the total expected payments. This can be expressed in terms of solving an integer program.

As before, we use $W_c(\mathcal{F})$ or $C(\mathcal{F})$ for short, to denote the total expected cost of a solution $\mathcal{F} = (F_1, \dots, F_r)$ when the individual link costs are c , and $W_{\tilde{a}}(\mathcal{F})$ to denote the price of the flow \mathcal{F} when the bids are \tilde{a} . We denote the mechanism output (*i.e.*, the min-price flow) by $\hat{\mathcal{F}}(\tilde{a})$.

5.2.1 The cooperative game \mathcal{G}

In this section, we define a cooperative game based on any specific instance of this mechanism, and prove that it has a non-empty core. This cooperative game is introduced only for strategic analysis of the mechanism. It is not explicitly played by the agents, but helps to shed light on the agents’ strategies in the two-parameter auction.

DEFINITION 2. *Given a set of players P , a cooperative game is defined by a characteristic function $v : 2^P \rightarrow \mathbb{R}_{\geq 0}$, with $v(\emptyset) = 0$. For any $S \subseteq P$, $v(S)$ is called the value of the set S .*

Given a directed graph G with distinguished source and sink, and a true cost c_i for each link $i \in G$, we define the cooperative game \mathcal{G} as follows:

The set of players in the game \mathcal{G} is $P = \{0, 1, \dots, n\}$, where each $i > 0$ is the player corresponding to link i , and 0 is a special player corresponding to the customer. Let Z be the customer’s budget, and assume that Z is large enough to be irrelevant; $Z > r \times \text{cost of minimum-cost } (r+1)\text{-flow}$ is sufficient. For each set $S \subseteq P$, $S \neq \emptyset$, define the value $v(S)$ of S in \mathcal{G} as follows:

If S does not contain a r -unit flow from s to t , $v(S) = 0$. If S contains the customer 0 as well as all edges on a k -unit flow from the source to the sink, $v(S)$ is defined to be the optimal value of the linear program given below:

Define $\delta_{i,S}$ to be the indicator of i in S , *i.e.*, $\delta_{i,S} = 0$ if $i \notin S$ and $\delta_{i,S} = 1$ if $i \in S$. Also, for any node α in the network, we use the notation $\text{In}(\alpha)$ to denote the set of incoming edges, and $\text{Out}(\alpha)$

to denote the set of outgoing edges. Then,

$$\begin{aligned}
v(S) &= \max \left\{ Z - \sum_{k=1}^r [p_k \sum_{i>0} c_i x_{ik}] \right\} \\
\text{Subject to:} & \\
\sum_{i \in \text{Out}(\alpha)} x_{ik} - \sum_{i \in \text{In}(\alpha)} x_{ik} &= 0 & \forall k \forall \alpha \neq s, t \\
\sum_{i \in \text{Out}(s)} x_{ik} - \sum_{i \in \text{In}(s)} x_{ik} - k &= 0 & \forall k \\
x_{ik} &\leq \delta_{i,S} & \forall k, \forall i \geq 0 \\
x_{ik} &\geq 0 & \forall k, \forall i \geq 0
\end{aligned} \tag{6}$$

This linear program is interpreted as follows: For any link i , and any demand value k , the variable x_{ik} indicates the flow along i in F_k . Intuitively, the value of a set S is related to the net surplus that is created when only the agents in set S are involved in the flow. If S does not contain the customer and a r -unit flow, $v(S)$ is defined to be 0. Thus, only sets that contain at least one candidate solution are assigned a positive value.

We also note that if $S = P$, then the linear program has an integral optimal solution, corresponding to an integral min-cost k -flow for each k . In other words, there is a solution in which x_{ik} is either 0 or 1 for all i and k . It is also clear that $v(S) \leq v(P)$ for all $S \subseteq P$.

Thus, the function $v(S)$ defines a finite, nonnegative value for each coalition set S , and hence it is the characteristic function of a valid cooperative game \mathcal{G} .

Our analysis is centered on the concept of the *core* of a cooperative game. Loosely speaking, the core of a cooperative game consists of all ways to divide up the overall value $v(P)$ among the agents such that no group S has reason to be unhappy – *i.e.*, S attains a combined utility of at least $v(S)$. Formally, the core is defined as follows:

DEFINITION 3. A vector $u = (u_0, u_1, \dots, u_n)$ is in the core of the game \mathcal{G} iff it satisfies all of the following:

$$\begin{aligned}
\forall i \quad u_i &\geq 0 & , \text{ and} \\
\sum_{i \in P} u_i &= v(P) & , \text{ and} \\
\forall S \subseteq P \quad \sum_{i \in S} u_i &\geq v(S).
\end{aligned}$$

In general, the core of a cooperative game might be the empty set. However, we can prove that this is not the case for the game \mathcal{G} :

LEMMA 3. The game \mathcal{G} has a nonempty core.

PROOF. Consider any division of $v(P)$ among the players. We show that there is at least one such division that satisfies all the core constraints. For any set S with $v(S) = 0$, the core constraint is trivially satisfied. Now, consider a set S with $v(S) > 0$. The linear program defining $v(S)$ can be summarized in the form $\max\{x \cdot l\}$ subject to $xH = 0$, $xA \leq b^S$, and $x \geq 0$, where x is a vector of all the variables, H and A are matrices independent of S , and b^S is a 0-1 vector representing the capacity constraints for set S . Then, the dual of the linear program (6) is the following linear program:

$$\begin{aligned}
\text{Subject to:} \quad v(S) &= \min \{ b^S \cdot y \} \\
Ay + Hz &\geq l \\
y &\geq 0
\end{aligned} \tag{7}$$

Now, consider the dual program that defines $v(P)$, *i.e.*, the value of the set containing the customer and all the links. Let (\hat{y}, \hat{z}) denote an optimal solution to this problem. Now, define $u_i = b^{\{i\}} \cdot \hat{y}$ for all i . Recall that b^S is a 0-1 vector, with 1s in precisely those equations that involve some $i \in S$; thus, $b^S = \sum_{i \in S} b^{\{i\}}$. Then,

as $\hat{y} \geq 0$, we have $u_i \geq 0$, and

$$\sum_{i \in P} u_i = b^P \cdot \hat{y} = v(P).$$

Next, observe that for any set $S \subseteq P$, the solution (\hat{y}, \hat{z}) is also feasible in the dual of the program (7) defining $v(S)$. Thus, we have

$$\sum_{i \in S} u_i = b^S \cdot \hat{y} \geq v(S).$$

Thus, the vector u is in the core of the game \mathcal{G} . \square

5.2.2 Existence of an ϵ -Nash equilibrium

We now show that given any point u in the core of this game, we can perturb it slightly to get a vector of bid pairs \tilde{a} that is an ϵ -Nash equilibrium of the bidding game. We use the game \mathcal{G} to draw conclusions about the bidding game induced by the mechanism.

THEOREM 7. Let u be any vector in the core of \mathcal{G} that minimizes the value of u_0 . Then, for any $\epsilon > 0$, the bid profile defined by

$$a_i^- = (c_i, u_i^- = \max\{0, u_i - \frac{\epsilon}{2n}\})$$

for each link i is an ϵ -Nash equilibrium.

PROOF. (Sketch) Suppose a^- is not an ϵ -Nash equilibrium. Then, there is some i such that i can change her bid to increase her payoff by ϵ . Let (c', u') be i 's successful strategy, and let a' denote the bid profile given by $a'_i = (c', u')$ and $a'_j = a_j^-$ for all $j \neq i$. Let $\mathcal{F}' = \hat{\mathcal{F}}(a')$; it must be the case that $i \in \mathcal{F}'$.

In the appendix, we show that there is a near-optimal flow \mathcal{F}'' such that \mathcal{F}'' does not use i (Lemmas 5,7). More specifically, $W_{a^-}(\mathcal{F}'') \leq W_{a^-}(\mathcal{F}') + \epsilon/2$. As $i \notin \mathcal{F}''$, we have $W_{a'}(\mathcal{F}'') = W_{a^-}(\mathcal{F}'')$. However, $i \in \mathcal{F}'$, and so $W_{a'}(\mathcal{F}') \geq W_{a^-}(\mathcal{F}') + \epsilon$. Thus, we get $W_{a'}(\mathcal{F}') > W_{a'}(\mathcal{F}'')$, which contradicts the assumption that $\mathcal{F}' = \hat{\mathcal{F}}(a')$. \square

We are working on strengthening Theorem 7 to show that this bid profile is indeed a *strong* ϵ -Nash equilibrium. This seems plausible given the results of Young [19]; however, the strategy space in our 2-parameter game is richer than the strategy space in Young [19].

5.3 Randomized 2-parameter Auction

The mechanism presented in Section 5.2 has an ϵ -Nash equilibrium corresponding to every core allocation, 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 Section 4. We prove that, with this modification, it is possible to bound the total payment.

The *Randomized 2-parameter Auction* is constructed as follows. As earlier, the edges' bids are pairs $\tilde{a}_i = (\tilde{c}_i, \tilde{u}_i)$. The mechanism has two components:

1. **The 2-parameter mechanism.** This mechanism is conducted exactly as described in Section 5.2 with parameters α , τ , and B set as before.
2. **The randomized audit.** For edges on a random path source-destination path, the payoff is based entirely on the \tilde{c}_i component of the bid, and is constructed as in Section 4. The parameters α , τ , and B are as defined in Section 4. To simplify the analysis, we assume that the randomized component results in a payoff function of the following form: 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 2, and has the same shape; the key aspect for us is that this function is maximized at $\tilde{c}_i = c_i$.

We also need to ensure that, for all edges i not in the winning solution, \tilde{u}_i is 0 (or close to zero). We assume that the mechanism simply rejects bid profiles that do not meet this condition. Alternatively, we could impose a small tax on the \tilde{u}_i component of the bid. We can now prove a useful lemma, which shows that *all* edges are nearly truthful about their costs in equilibrium:

LEMMA 4. *Let $\tilde{a} = (\tilde{c}, \tilde{u})$ be an ϵ -Nash equilibrium of the Randomized 2-parameter Auction. Then, for all i ,*

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

Using the fact that the costs are nearly truthful, we can show that their utility values are nearly in the core, and hence, derive the following bound on the total payment.

THEOREM 8. *Let $\tilde{a} = (\tilde{c}, \tilde{u})$ be any ϵ -Nash equilibrium of the Randomized 2-parameter Auction. Let \mathcal{F} be a lowest-cost flow, and let F_{r+1} be a lowest-cost $(r+1)$ -flow. Then, the total price paid by the customer in the randomized 2-parameter auction is at most*

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

The result of Theorem 8 stands in an interesting relation to that of Theorem 2. We do not achieve the intuitively appealing bound of the expectation of the bounds on the deterministic auction in Section 3, *i.e.*, $E_j[P_j] = \sum_{j=1}^r j p_j (C(F_{j+1}) - C(F_j))$ (proving this stronger bound is an interesting problem for future work). Instead we achieve $\sum_{j=1}^r r p_j (C(F_{r+1})(j/r) - C(F_j))$. 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 8 is therefore weaker in two important respects than Theorem 2, but it does have a similar overall structure.

6. CONCLUSION

The results in Section 3 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 example in Section 1 and results in [2, 8] 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 5.1 and 5.2 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 [19]. 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, which could 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 [15], could 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.

Acknowledgment

We would like to thank Eva Tárdoš for suggesting the bound $C(2) - C(1)$ on the payment in a 1-unit auction. We also thank Joan Feigenbaum for helpful comments on an earlier version.

7. REFERENCES

- [1] A. Archer, C. Papadimitriou, K. Talwar and E. Tárdoš. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics* **1**(2): 129-150, 2003.
- [2] A. Archer and E. Tárdoš. Frugal Path Mechanisms. In *Proceeding of the 13th ACM Symposium on Discrete Algorithms* 2002, pp. 991-999.
- [3] R.J. Aumann. Acceptable points in general cooperative n -person games. In: A.W. Tucker, R.D.Luce (Eds.), *Contributions to the Theory of Games, Annals of Mathematical Studies* **40**, Princeton University Press, Princeton, New Jersey, 1959.
- [4] E. Clarke. Multipart pricing of public goods. *Public Choice* **11**:17-33, 1971.

- [5] Imma Curiel. Cooperative Game Theory and Applications. Kluwer Academic Publishers, 1997.
- [6] A. Czumaj, and A. Ronen. On the Expected Payment of Mechanisms for Task Allocation. In *Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing*, 2004, pp. 98–106.
- [7] Rahul Garg, Vijay Kumar, Atri Rudra, Akshat Verma. Coalitional games on graphs: core structure, substitutes and frugality. In Proceedings of the 4th ACM conference on Electronic commerce, 2003, pp. 248–249.
- [8] E. Elkind, A. Sahai and K. Steiglitz. Frugality in Path Auctions. In *Proceedings of 15th ACM Symposium on Discrete Algorithms*, 2004, pp. 701–709.
- [9] J. Feigenbaum, C. Papadimitriou, R. Sami and S. Shenker. A BGP-based Mechanism for Lowest-Cost Routing. In *Proceedings of the 21st Symposium on Principles of Distributed Computing*, 2002, pp. 173–182.
- [10] M. Garey, and D. Johnson. Computers and Intractability. Freeman Press, New York, 1979.
- [11] T. Groves. Incentives in teams. *Econometrica* **41**:617–663, 1973.
- [12] J. Hershberger and S. Suri. Vickrey prices and shortest paths: What is an edge worth? In *Proceedings of the 42nd IEEE Symposium on the Foundations of Computer Science*, 2001, pp. 129–140.
- [13] M. Jackson. A Crash Course in Implementation Theory. *Social Choice and Welfare*, **18**(4):655–708, 2001.
- [14] E. Kalai, A. Postlewaite, and J. Roberts. A group incentive compatible mechanism yielding core allocations. *Journal of Economic Theory* **20**:13–22, 1979.
- [15] V. Krishna and R. Serrano. Multilateral Bargaining. *Review of Economic Studies* **63**:61–80, 1998.
- [16] M. Mihail, C. H. Papadimitriou, A. Saberi. On Certain Connectivity Properties of the Internet Topology. In *Proceedings of the 44th IEEE Symposium on the Foundations of Computer Science*, 2001, pp. 28–35.
- [17] N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior* **35**:166–196, 2001.
- [18] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16**:8–37, 1961.
- [19] H. Peyton Young. Cost Allocation, Demand Revelation, and Core Implementation. *Mathematical Social Sciences* **36**:213–228, 1998.

8. APPENDIX

8.1 Proof of Lemma 2

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

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

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, $\frac{\partial^2 g(b_e)}{\partial b_e^2} = -\tau < 0$. This is true for all paths P containing e , and thus, $b_e = c_e$ is the optimal bid for

player 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, by the condition of ϵ -Nash equilibrium, $\Delta \leq \sqrt{2\epsilon/\tau}$. \square

8.2 Proof of Lemma 1

PROOF. Construct a bid vector \mathbf{b} as in Theorem 1. By this construction we have $b_e = c_e$ for any edge e that is not on the lowest-price path. Then, following the analysis of $g(b_e)$ expected payoff in Lemma 2, 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 its bid. Note that 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 gains $g(b_e)$ from the bonus and probability of off-path selection would also drop. Thus, i cannot possibly gain more than ϵ by raising its bid.

Finally, 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 profits from being on the lowest-price path, and gain at most rBx in $g_e(b_e)$; thus, it could not gain overall. \square

8.3 Proof of Theorem 5

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 $2\alpha n B$.

Finally, using Lemma 2, 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 2, we can bound the total payment to edges on the lowest-price path by

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

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

8.4 Proofs for the 2-parameter mechanism

8.4.1 Positive payoffs in the core

LEMMA 5. Let \mathcal{F} be a lowest-cost solution, and F_{r+1} be a lowest-cost $(r+1)$ -flow. For any vector u in the core of \mathcal{G} , we have $u_0 > Z - [C(F'_{r+1}) \sum_{j=1}^r j p_j - rC(\mathcal{F})]$.

PROOF. Let \mathcal{F} be a minimum-cost flow. Then, $v(P) = Z - C(\mathcal{F})$, and hence, by the core condition, $\sum_{i \in P} u_i = Z - C(\mathcal{F})$. Consider a lowest-cost integral $(r+1)$ -flow F_{r+1} . Then, F_{r+1} consists of $(r+1)$ disjoint paths from s to t ; call them P_1, \dots, P_{r+1} . For each $k \in \{1, 2, \dots, r, r+1\}$, define $F_r^{-k} = F_{r+1} \setminus P_k$, i.e., the r -flow obtained by dropping the k th path. Extend F_r^{-k} to a collection of flows $\mathcal{F}^{-k} = (F_1^{-k}, F_2^{-k}, \dots, F_r^{-k})$, where F_j^{-k} consists of the j lowest-priced paths in F_r^{-k} . Then, as \mathcal{F}^{-k} can

meet the demand, we have: $v(\mathcal{F}^{-k} \cup \{0\}) \geq Z - C(\mathcal{F}^{-k})$. Further, noting that F_j^{-k} has cost at most $\frac{j}{r}$ that of F_r^{-k} , we get:

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

Further, as u is in the core, we have $u_0 + \sum_{i \in \mathcal{F}^{-k}} u_i \geq v(\mathcal{F}^{-k} \cup \{0\})$. Now, adding over all k , we get:

$$\begin{aligned} \sum_{k=1}^{r+1} v(\mathcal{F}^{-k} \cup \{0\}) &\geq \sum_{k=1}^{r+1} [Z - C(\mathcal{F}^{-k})] \\ \sum_{k=1}^{r+1} v(\mathcal{F}^{-k} \cup \{0\}) &\geq (r+1)Z - \sum_{k=1}^{r+1} \left[C(F_r^{-k}) \sum_{j=1}^r p_j \frac{j}{r} \right] \end{aligned}$$

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,

$$(r+1)u_0 + r \sum_{i \in F_{r+1}} u_i \geq (r+1)Z - rC(F_{r+1}) \sum_{j=1}^r p_j \frac{j}{r}$$

Noting that $\sum_{i \in F_{r+1}} u_i \leq \sum_{i \in P} u_i = Z - C(\mathcal{F})$, we get:

$$\begin{aligned} u_0 + r(Z - C(\mathcal{F})) &\geq (r+1)Z - C(F_{r+1}) \sum_{j=1}^r j p_j \\ u_0 &\geq Z - [C(F_{r+1}) \sum_{j=1}^r j p_j - rC(\mathcal{F})] \quad \square \end{aligned}$$

LEMMA 6. *Given a network and a cost vector c , and some element u in the core of \mathcal{G} , define the bid profile \tilde{a} by*

$$\tilde{a}_i = (c_i, u_i) \quad \forall i > 0$$

Then, the lowest-price flow output by the mechanism with input \tilde{a} has a total price of $Z - u_0$. Further, any minimum-cost flow \mathcal{F} is an optimal (minimum-price) flow, and includes all links i with $u_i > 0$.

PROOF. First, let \mathcal{F} be an optimal integral solution to the linear program defining $v(P)$. Then, an examination of the objective function of LP 6 shows that $C(\mathcal{F}) = Z - v(P) = Z - \sum_{i=0}^n u_i$. Now,

$$W_{\tilde{a}}(\mathcal{F}) \leq C(\mathcal{F}) + \sum_{i=1}^n u_i = Z - u_0 \quad (8)$$

We now show that this is also a lower bound on the cost. Suppose there was some flow \mathcal{F}' such that $W_{\tilde{a}}(\mathcal{F}') < Z - u_0$. It follows that $C(\mathcal{F}') < Z - u_0 - \sum_{i \in \mathcal{F}'} u_i$. Now, consider the linear program determining the value of the coalition $S = \mathcal{F}' \cup \{0\}$. The flow \mathcal{F}' is a feasible solution for this set, and hence

$$v(S) \geq Z - C(\mathcal{F}') > u_0 + \sum_{i \in \mathcal{F}'} u_i = \sum_{i \in S} u_i$$

But this contradicts the assumption that u is in the core of G . Hence every flow (including \mathcal{F}) has price at least $Z - u_0$. Thus, $W_{\tilde{a}}(\mathcal{F}) = Z - u_0$ and \mathcal{F} includes all i such that $u_i > 0$. \square

8.4.2 Optimal flow without using player i

LEMMA 7. *Let u be a vector in the core of \mathcal{G} that minimizes the value of u_0 , and define the bid vector \tilde{a} by $\tilde{a}_i = (c_i, u_i) \quad \forall i > 0$. Then, if $u_0 > 0$, for any i there is a flow $\mathcal{F}^{(i)}$ such that $W_{\tilde{a}}(\mathcal{F}^{(i)}) = W_{\tilde{a}}(\tilde{\mathcal{F}}(\tilde{a}))$, i.e., $\mathcal{F}^{(i)}$ is an optimal solution.*

PROOF. Let $\mathcal{F} = \tilde{\mathcal{F}}(\tilde{a})$. Assume there is an i such that the statement is not true. Let $\mathcal{F}^{(i)}$ be the lowest-price flow that does not include i , and assume that $W_{\tilde{a}}(\mathcal{F}^{(i)}) = D > W_{\tilde{a}}(\mathcal{F}) + \delta$ for some $\delta > 0$. Define a vector u' by $u'_0 = u_0 - \delta$, $u'_i = u_i + \delta$, and $u'_j = u_j \forall j \neq 0, i$. We now claim that u' is in the core. If not, there would be some set S such that $v(S) > \sum_{j \in S} u'_j$. We must have $0 \in S$, or else $v(S)$ would be 0. Similarly, S must contain a k -flow, or else its value would be 0. It follows that $i \notin S$, or else we would have $\sum_{j \in S} u'_j = \sum_{j \in S} u_j \geq v(S)$. Let \mathcal{F}'' be the lowest-cost flow in S . Then, $v(S) = Z - W_c(\mathcal{F}'')$ and so we get

$$\begin{aligned} 0 < v(S) - \sum_{i \in S} u'_i &= Z - W_c(\mathcal{F}'') - \sum_{j \in S} u_j + \delta \\ 0 < Z - u_0 + \delta - W_{\tilde{a}}(\mathcal{F}'') &< Z - u_0 - W_{\tilde{a}}(\mathcal{F}) \end{aligned}$$

But, using Lemma 6, $Z - u_0 - W_{\tilde{a}}(\mathcal{F}) = 0$ because \mathcal{F} is a minimum-price flow with bids \tilde{a} , and so this is a contradiction. Thus, u' must be in the core; but this contradicts the assumption that u is an element of the core that minimized u_0 . \square

8.4.3 Proof of Lemma 4

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 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. \square