Frugality in Path Auctions

EDITH ELKIND *   AMIT SAHAI *   KEN STEIGLITZ *

Preliminary version – April 2003

Abstract

We consider the problem of picking (buying) an inexpensive $s - t$ path in a graph where edges are owned by independent (selfish) agents, and the cost of an edge is known to its owner only. We study the problem of finding *frugal mechanisms* for this task, *i.e.* we investigate the payments the buyer must make in order to buy a path.

First, we show that any mechanism with (weakly) dominant strategies (or, equivalently, any truthful mechansim) for the agents can force the buyer to make very large payments. Namely, for every such mechanism, the buyer can be forced to pay $c(P) + \frac{1}{2}n(c(Q) - c(P))$, where $c(P)$ is the cost of the shortest path, $c(Q)$ is the cost of the second-shortest path, and n is the number of edges in P. This extends the previous work of Archer and Tardos [1], who showed a similar lower bound for a subclass of truthful mechanisms called *min-function* mechanisms. Our lower bounds have no such limitations on the mechanism.

Motivated by this lower bound, we study mechanisms for this problem providing Nash equilibrium strategies for the agents. In this class, we identify the *optimal* mechanism with regard to total payment. We then demonstrate a separation in terms of average overpayments between the classical VCG mechanism and the optimal mechanism showing that under various natural distributions of edge costs, the optimal mechanism pays at most logarithmic factor more than the actual cost, whereas VCG pays $\sqrt{n}$ times the actual cost. On the other hand, we also show that the optimal mechanism does incur at least a constant factor overpayment in natural distributions of edge costs. Since our mechanism is optimal, this gives a lower bound on all mechanisms with Nash equilibria.

*Department of Computer Science, Princeton University. Princeton, NJ 08544.
E-mail:`elkind, sahai, ken@cs.princeton.edu`.

1 Introduction

Internet protocols often involve interaction among multiple participants. Most often, these participants, or *agents*, are selfishly motivated, and as such cannot be expected to follow the rules of a protocol if deviating from the protocol allows the agents to achieve some gain. In contexts in computer science where privacy is important, cryptographic tools have been developed and used to protect parties against dishonest behavior. However, where the implication of dishonest behavior is a quantifiable gain or loss (in terms of some resource, whether it be money or time), looking at this problem from an economic perspective is useful. The field of *mechanism design* concerns the design of protocols which discourage certain "bad" behaviors (*i.e.* behaviors which have quantifiable consequences on resources that the protocol designer wishes to avoid).

Mechanisms are protocols that collect information privately known to a group of agents and use it to determine some *outcome*. (For instance, a standard auction mechanism will collect bids, and based on these bids selecet a winner for the auctioned item and determine how much that winner should pay.) The main concern here is that an agent may not honestly report its private information if it believes that lying can increase its expected profit. A mechanism is called *truthful* if it is in every agent's "best interest" to simply report its private information honestly.[1]

A critical aspect of this definition is in what sense do we interpret "best interest." A very strong interpretation of this condition requires that no matter what (potentially foolish) strategy any other agent follows, each agent is individually best off if it reports its private information honestly to the mechanism. In other words, truth-telling is a (weakly) *dominant* strategy for each agent. A less stringent interpretation is possible if one assumes rational selfish behavior for all participants. Here, we are satisfied if the strategies of each agent honestly reporting their private information form a *Nash equilibrium*. This means that if any particular agent believes that all the other agents will tell the truth, then truth-telling is the best option for that agent also.

Frugality. To "convince" each agent to act honestly, typically a mechanism will have to pay agents in ways that encourage honest behavior. Unfortunately, this can often lead to mechanisms with payment rules that require potentially enormous payments to agents. In this paper, we study this critical problem, and concern ourselves with the question of designing mechanisms that avoid large "overpayments."

We study this phenomenon in a situation which appears to suffer most acutely from this problem. We consider mechanisms which must assemble a team of agents to perform a larger task. Each agent performs a specific fixed service, but only certain combinations of these services will suffice to accomplish the larger objective. A special case of considerable interest (c.f. [9, 10, 6, 1]) is where each agents controls an edge in a network, and the mechanism must purchase a path between some nodes s and t. In this case, allowed teams are sets of agents which control s—t paths. The private information of each agent is its *cost* in selling the use of its edge. Note that in fact, conversely, many team-organizing problems can be modeled as finding paths in a graph.

This setting differs critically from that of standard auctions in that the buyer gains no value unless it acquires a particular allowable *combination* of goods. To see why this situation lends itself to the problem of overpayment, consider the situation from the perspective of an agent in the winning team. (For rough intuition, we will consider the "first-price auction" protocol in which each agent places a bid, and the team with the smallest total bids wins, and each agent is paid its bid.) Suppose each agent truthfully reports its private cost (we will argue that this is *not* in the best interest of the agents). If the total cost C of the agents on the winning team is significantly lower than the total cost C' of the second least expensive team (a likely situation if the sizes of needed teams are large), then *each* agent on the winning team is tempted to have

[1] One might wonder why we should care so much about insuring honesty, when our overall objective might be something different – like maximizing revenue or minimizing cost. In fact, concentrating on truth-telling equilibria is without loss of generality by the "revelation principle" of game theory: for every mechanism with some equilibrium, there exists another essentially equivalent mechanism with truth-telling as its equilibrium.

increased its bid by $\Delta = C' - C$. While the difference in *total* cost is Δ, the overpayment in this case could be $n\Delta$, where n is the number of agents in the winning team.

1.1 Our Results

We focus on the path problem as described above. We first study the problem of worst-case overpayment in the context of mechanisms which have truth-telling as a (weakly) dominant strategy for each agent. Here, we prove a strong negative result for *all* mechanisms, essentially matching the rough intuition given above: We show that for every mechanism, if truth-telling is a dominant strategy for each player, then the mechanism can be forced to pay $c(P) + \frac{1}{2}n(c(Q) - c(P))$, where $c(P)$ is the cost of the shortest path, $c(Q)$ is the cost of the second-shortest path, and n is the number of edges in P. We extend this bound to hold for randomized mechanisms, and furthermore, we show that for almost all cases, the factor of $\frac{1}{2}$ can be eliminated as well. This extends the previous work of Archer and Tardos, who showed a similar lower bound for a subclass of such mechanisms called *min-function* mechanisms. Our lower bounds have no limitations of this kind on the mechanism.

This lower bound motivates us to study the problem of overpayment with regard to mechanisms using Nash equilibria. In one of the main technical contributions of this work, we present an analysis which identifies the *optimal* mechanism with regard to total payment – *i.e.* the Nash equilibrium of this mechanism minimizes the total expected payment for the buyer, over all mechanisms. It turns out that while the optimal mechanism has dominant strategies for all players (and hence the bound established in the first part of the paper still applies), the average overpayment can be reduced dramatically. Namely, we show that under various natural distributions of edge costs, the optimal mechanism pays at most logarithmic factors more than the actual cost, whereas the classical VCG mechanism must pay $\sqrt{n}$ times the actual cost. On the other hand, we also show that the optimal mechanism does incur at least a constant factor overpayment in natural distributions of edge costs. Since our mechanism is optimal, this gives a lower bound on all mechanisms with Nash equilibria.

1.2 Related Work

Recently, a number of results have been obtained by applying ideas and methods from theoretical computer science to problems of economics and game theory [10]. Many of them concentrate on the field of mechanism design (also known as implementation theory, or theory of incentives). One of the most important results here is the celebrated Vickrey–Clarke–Groves (VCG) mechanism [12, 2, 5], which guarantees *efficient* allocation, i.e., the bidder with the highest valuation, or, in the case of procurement auctions, the lowest cost wins the auction. In our context, efficiency corresponds to picking the cheapest path; while this maximizes social welfare, as we will see, it does not always minimize the total overpayment.

The first paper to explicitly use this approach for the shortest path problem is [9], in which the authors formulate the problem, describe the solution given by VCG, and discuss the associated computational difficulties. However, [9] does not address the issue of minimizing the payments to the agents, but rather treats the payments as a tool for eliciting truthful responses.

On the other hand, revenue maximizing (or, equivalently, payment minimizing) auctions in other settings received a great deal of attention from both economists and computer scientists. One of the most prominent papers in this area is [8], which constructs the optimal auction for selling an item to one out of n buyers. While Myerson's results do not apply directly in our situation, we use his techniques to derive the results in the second part of the paper. In [11], Ronen and Saberi generalize the results of [8] in a different direction, namely, that of interdependent valuations. The recent papers [4, 3] investigate the properties of revenue-maximizing auctions for digital goods.

These two lines of research are combined in [1], which raises the issue of frugality in shortest path auctions. The authors address the same question as we do, namely, that of designing a mechanism that selects a reasonably short path and induces truthful bidding without paying an unacceptably high premium. They provide a partial negative solution to the problem. Namely, they describe a general class of *min function* mechanisms and show that every mechanism from this class can be forced to pay $c(P) + k(c(Q) - c(P))$, where $c(P)$ is the cost of the chosen path, $c(Q)$ is the cost of an alternate path, and k is the number of edges in the chosen path. Furthermore, they list a number of properties (other than truthfulness) that can be desired from an auction mechanism, and show that on two large classes of graphs (namely, ones containing an $s - t$ arc and ones that consist of some connected graph between s and t and two extra $s - t$ paths that are disjoint from the rest of the graph) every mechanism that enjoys these properties is a min-function mechanism. Together, these results imply that any truthful mechanism with these properties on a graph with three node-disjoint $s - t$ paths cannot avoid paying an unacceptably high premium.

To compare this with the main result of the first part of this paper, note that our approach works for *any* individually rational truthful mechanism on a graph with at least two node disjoint $s - t$ paths. However, the premium that the mechanism is forced to pay is $c(P) + \frac{1}{2}k(c(Q) - c(P))$, while in the setting of [1] the payment can be as high as $c(P) + k(c(Q) - c(P))$. Furthermore, our bound applies to randomized mechanisms as well.

2 Preliminaries

We model the network by a graph $G = (V, E)$ with two distinguished vertices s and t. Each edge $e_i \in E$ has an associated *cost* $c(e_i)$, which is private, that is, known to the owner of e_i only. A *cost* of a path in G is the sum of the costs of the edges on the path. By the *shortest path* we mean the path that has the smallest cost.

The costs that the edges announce for themselves are called *bids*. We denote the bid of edge e_i by $b(e_i)$; b_P stands for the vector of bids along a path P. The auction mechanism is supposed to select a path between s and t; we refer to this path as the *winning* path, and say that edges on the selected path *win*, while edges not on the path *lose*.

We assume that the buyer has no alternative to picking some path in the graph, i.e., his costs when no path is chosen are infinite. There is no loss of generality in this assumption: if this is not the case, and the buyer has an alternative of cost c_0, we can model this by introducing an arc of cost c_0 from s to t; if this arc wins, the buyer reverts to his back-up plan.

A *mechanism* is a triple $(\mathcal{B}, Q(\mathbf{b}), M(\mathbf{b}))$, where $\mathcal{B} = \prod_i B_i$ is the set of possible bids, $Q : \mathcal{B} \mapsto [0, 1]^E$ is an *allocation rule*, and $M : \mathcal{B} \mapsto \mathbb{R}^E$ is a *payment rule*, where $Q_i(\mathbf{b})$ is the probability that e_i is on the chosen path given that the bid vector is $\mathbf{b}$, and $M_i(\mathbf{b})$ is the corresponding payment to agent i.

There are two basic solution concepts considered in the game-theoretic literature, namely, an equilibrium in dominant strategies and Nash equilibrium, the former being much stronger than the latter.

Definition 2.1 A strategy in a game is (weakly) *dominant* if, regardless of what any other players do, the strategy earns a player a payoff that is at least as large as that earned by any other strategy.

Definition 2.2 A (weak) *Nash equilibrium* is a set of strategies, one for each player, such that a change in strategies by any player would lead that player to earn no more than if she remained with her current strategy.

Clearly, if for a given game each player has a dominant strategy, then this set of strategies constitutes a Nash equilibrium. The converse is not always true, and, moreover, for some games there exists a Nash equilibrium but none of the players has a dominant strategy (for example, consider the game of Rock, Paper, Scissors: it is not difficult to show that in the one-round version of the game picking each move with probability $1/3$ is an equilibrium strategy; however, if you know that your opponent always plays Rock, this strategy is no longer optimal).

In the first part of the paper, we concentrate on mechanisms that possess dominant strategies; the analysis of the second part applies to Nash equilibria as well.

A mechanism can, in principle, be quite complicated, since so far we have made no restrictions on what the participants' bids can be. Fortunately, it turns out that we can restrict our attention to *truthful*, or *incentive compatible* mechanisms, in which each player's best strategy is to announce his true value. This result is usually referred to as the *revelation principle* and can be stated as follows:

Proposition 2.3 *For any mechanism $M = (\mathcal{B}, Q(\mathbf{b}), M(\mathbf{b}))$ that has a Nash equilibrium, there is a corresponding mechanism M' such that for any set of players' values, the outcomes and payments of M' are the same as in a given equilibrium of M, and*

- *under M', it is an equilibrium for each agent to report her value truthfully;*
- *If a player i had a dominant strategy under M, then truth-telling is a dominant strategy under M'.*

It is well known (see, e.g. [9, 4]) that if truth-telling is a dominant strategy, the payment to an edge e_i depends only on the bids of other edges and whether this edge wins or loses. It follows immediately that there must exist a *threshold bid*, i.e., the highest bid of this edge that still wins the auction, given that the bids of other participants remain the same.

We require that the mechanism has the property of *voluntary participation* or *individual rationality*, that is, the payment an edge receives is no less than its bid. Furthermore, since we are interested in minimizing the total payment, there is no loss of generality in restricting ourselves to mechanisms that pay zero to losing edges. In this case, the payments to edges are completely determined by the path selection rule: each winning edge gets its threshold bid T, each losing edge gets 0.

The VCG mechanism is a truthful mechanism that maximizes the "social welfare". In our case, this simply means picking the shortest available path and paying each agent his threshold bid.

3 Costly Example

In this section, we show that for any graph that contains two node-disjoint $s-t$ paths, any auction mechanism for which truthful bidding is a dominant strategy can be forced to pay $b(P) + \frac{1}{2}k(b(Q) - b(P))$, where P is the shortest path, Q is the shortest alternative path, $b(T) = \sum_{e \in T} b(e)$ is the total cost of path T, and $k = \min(|P|, |Q|)$.

Consider a truthful mechanism on a graph that consists of two node-disjoint paths, P and Q, $|P| = n_1$, $|Q| = n_2$.

Theorem 3.1 *For any $L, \epsilon > 0$, there are bid vectors b_P, b_Q such that $b(P) = L$, $b(Q) = L + \epsilon$, and the total payment is at least $L + \frac{\epsilon}{2}\min(n_1, n_2)$.*

PROOF: Fix arbitrary positive L, ϵ. Let b_P^i, $i = 1, \ldots, n_1$, denote the vector of bids along P where each edge bids $\frac{L}{n_1}$ except for the ith edge that bids $\frac{L}{n_1} + \epsilon$. Similarly, let b_Q^j, $j = 1, \ldots, n_2$, denote the vector of bids along Q where each edge bids $\frac{L}{n_2}$ except for the jth edge that bids $\frac{L}{n_2} + \epsilon$. By b_P^0 (respectively, b_Q^0) denote the vector of bids along P (respectively, Q) where each edge bids $\frac{L}{n_1}$ (respectively, $\frac{L}{n_2}$).

Consider a directed bipartite graph G whose vertices are b_P^i, b_Q^j, $i = 1, \ldots, n_1$, $j = 1, \ldots, n_2$, and there is an edge from b_P^i to b_Q^j if whenever the edges along P bid according to b_P^i and the edges along Q bid according to b_Q^j, path Q wins.

This graph has $n_1 n_2$ edges (there is exactly one edge for each pair (b_P^i, b_Q^j)) and $n_1 + n_2$ vertices, so there is a vertex that has at least $\frac{n_1 n_2}{n_1+n_2}$ edges leaving it. Without loss of generality, suppose that this vertex is b_Q^1, and the endpoints of the edges that leave it are $b_P^{i_1}, \ldots, b_P^{i_t}$, $t = \frac{n_1 n_2}{n_1+n_2}$.

Consider the situation when edges along Q bid according to b_Q^1 and edges along P bid according to b_P^0. It is known that any truthful auction is monotone, that is, a losing edge cannot cause itself to win by raising its bid. Since the mechanism chooses $b_P^{i_1}$ over b_Q^1, $i_1 > 0$, this implies that in our setting P is going to win.

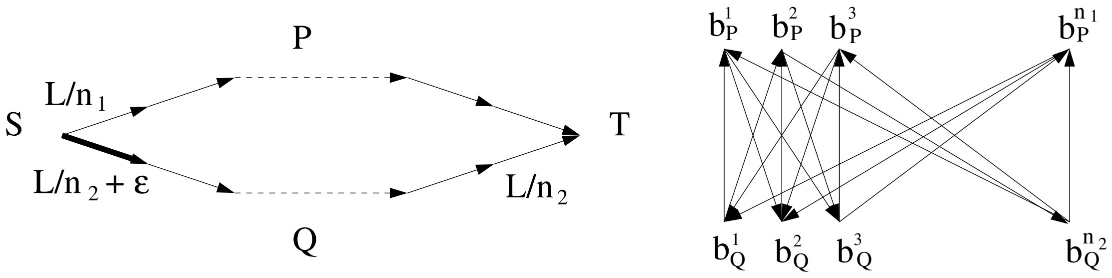

Figure 1: Costly example.

The payments to edges on P are determined by their threshold bids. Obviously, each edge will be paid at least its bid, that is, $\frac{L}{n_1}$. Furthermore, we know that if any of the edges $i_1, \ldots, i_t$ raises its bid by ϵ, path P still wins, so for these $\frac{n_1 n_2}{n_1+n_2}$ edges the threshold bid is at least $\frac{L}{n_1} + \epsilon$. Hence, the total payment is at least $L + \epsilon\frac{n_1 n_2}{n_1+n_2}$. Note that the total cost of the winning path is L, while the total cost of the shortest alternative path is $L + \epsilon$. As

$$\frac{1}{2}\min(n_1, n_2) \leq \frac{n_1 n_2}{n_1 + n_2} \leq \min(n_1, n_2),$$

it follows that the total payment can be much greater than the actual cost, or even the cost of the shortest alternative path. □

Given this result, it is natural to ask whether we can construct a truthful mechanism $\mathcal{M}^0$ such that for any L and ϵ and any bid vectors b_P, b_Q satisfying $b(P) = L$, $b(Q) = L + \epsilon$, the mechanism pays at most $L + \epsilon k/2$.

The following theorem shows that this is not true; namely, for any truthful mechanism, there are some values of L, ϵ, b_P, and b_Q that force the mechanism to pay almost $L + \epsilon k$. Furthermore, these "bad cases" are, loosely speaking, everywhere dense: for any $L_1 < L_2$, there are two bid vectors b_P, b_Q, $L_1 \leq b(P) < b(Q) \leq L_2$ such that P wins and receives almost $b(P) + k(b(Q) - b(P))$; moreover, we can guarantee that $b(Q) - b(P)$ is not too small. More precisely, the following theorem holds:

Theorem 3.2 *For any truthful mechanism $\mathcal{M}$ on G, any $L_1 < L_2$, and any $\epsilon' \in [0, L_2 - L_1]$, there are two bid vectors b_P, b_Q, $L_1 \leq b(P) < b(Q) \leq L_2$, such that P wins and either*

- *$b(P) = L_1$, $b(Q) = L_2$, and P is paid at least $b(P) + k(b(Q) - b(P) - \epsilon') = b(P) + k(L_2 - L_1 - \epsilon')$, or*
- *$b(Q) - b(P) \geq \epsilon'/k$ and P is paid at least $b(P) + k(b(Q) - b(P))$.*

Remark 3.3 Note that ϵ' can be arbitrarily small, which means that the amount of overpayment can be made arbitrarily close to $k(b(Q) - b(P))$; however, in this case $b(Q) - b(P)$ becomes small, too.

For the proof, the reader is referred to Appendix A.

3.1 Embedding into a Bigger Graph

We would like to embed this costly example into a bigger graph. That is, given an arbitrary graph that contains two node-disjoint paths P and Q between S and T, we would like to assign weights to edges that do not belong to P or Q so that the mechanism never chooses any of these edges. If this can be done, the argument goes through without change.

It turns out that while this is not necessarily the case, it is still true that any mechanism on a graph that contains such two paths can be forced to pay at least $b(P) + \frac{1}{2}\min(|P|, |Q|)(b(Q) - b(P))$. The proof, as well as the counterexample, can be found in Appendix B.

3.2 Extension to Randomized Case

We are able to show that our lower bound applies to randomized auction mechanisms as well. Following [9], we define a randomized mechanism as a probability distribution over deterministic mechanisms and say that the resulting mechanism is truthful if all the underlying deterministic mechanisms are truthful.

Definition 3.4 (cf. [9]) A *randomized mechanism* is a probability distribution over a family $\{m_r \mid r \in I\}$ of mechanisms, all sharing the same sets of strategies and possible outputs.

The outcome of such a mechanism is a probability distribution over outputs and payments; the problem specification must specify what output distributions are required. For the case of optimization problems, the *objective function* on such a distribution is taken to be the expectation over the choice of $r \in I$.

Definition 3.5 A strategy a^i is called *universally dominant* (in short, dominant) for agent i if it is a dominant strategy for every mechanism in the support of the randomized mechanism. A randomized mechanism is called *universally truthful* (in short, truthful) if truth-telling is a dominant strategy, and *strongly truthful* if it is the only one.

Now, consider a randomized auction mechanism on the graph of Fig. 1. Fix arbitrary L, ϵ and construct the above-described bipartite graph for each deterministic mechanism in its support. By linearity of expectation, there is a vertex of the bipartite graph whose expected outdegree is at least $\frac{n_1 n_2}{n_1+n_2}$. Suppose again that this vertex is b_Q^1 and consider the auction in which Q bids according to b_Q^1 and P bids according to b_P^0. As we have seen, for any deterministic mechanism m_r, the payment to the edges of P is at least $L + \mathrm{outdegree}_r(b_Q^1) \times \epsilon$, where $\mathrm{outdegree}_r(b_Q^1)$ is the outdegree of b_Q^1 in the bipartite graph that corresponds to m_r. Taking the expectation over all m_r, we see that the expected payment is at least $L + \frac{1}{2}k\epsilon$.

4 Optimal Auctions

In this section, we consider a wider class of mechanisms, namely, ones that have a Nash equilibrium (but not necessarily a dominant strategy for each player). We assume that the mechanism designer knows the distributions of costs on each edge. Under some weak regularity conditions on the cost distributions, we find the optimal mechanism for our problem, that is, the one that minimizes the total payment to the edges.

Our argument extends the reasoning in the seminal paper of Myerson [8], as presented in [7]. However, Myerson's original results are derived for standard auctions only, that is, the ones in which a single object is to be allocated to one out of n buyers, while in our case the buyer only gains utility when he acquires a path, i.e., a *set* of edges. There are several generalizations of Myerson's work, but none of them applies directly to our situation.

We start by introducing some new notation. Suppose that the cost x_i of each edge $e_i \in E$ is randomly chosen from a probability space $X_i = [0, \omega_i]$ according to a distribution F_i with density function f_i.

Set $X = \prod_{e_i \in E} X_i$. Define the joint density function $f(\mathbf{x}) = \prod_{e_i \in E} f_i(x_i)$. For any vector $\mathbf{v} \in \mathbb{R}^n$, set $\mathbf{v}_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$.

By revelation principle, we can restrict ourselves to truthful mechanisms. In this case, a mechanism is completely defined by an allocation rule $Q : X \mapsto [0,1]^E$ and a payment rule $M : X \mapsto \mathbb{R}^E$.

Define

$$q_i(z_i) = \int_{X_{-i}} Q_i(z_i, \mathbf{x}_{-i}) f_{-i}(\mathbf{x}_{-i}) d\mathbf{x}_{-i}$$

and

$$m_i(z_i) = \int_{X_{-i}} M_i(z_i, \mathbf{x}_{-i}) f_{-i}(\mathbf{x}_{-i}) d\mathbf{x}_{-i}.$$

That is, $q_i(z_i)$ is the probability that i wins if he reports his costs as z_i and everyone else reports their costs truthfully, and similarly for $m_i(z_i)$.

The *expected utility* of a truthful player i whose costs are x_i is $U_i(x_i) = m_i(x_i) - x_i q_i(x_i)$. Note that the payment received by the agent is not conditional on his winning the auction; rather, the probability of this event is incorporated into the payment rule.

Claim 4.1 *For any incentive compatible mechanism* (Q, M)*, we have*

$$U_i(x_i) = U_i(0) - \int_0^{x_i} q_i(t_i) dt_i,$$

$$m_i(x) = m_i(0) - \int_0^{x_i} q_i(t_i) dt_i + x_i q_i(x_i).$$

PROOF: From incentive compatibility,

$$U_i(z_i) \geq m_i(x_i) - z_i q_i(x_i) = U_i(x_i) + (x_i - z_i) q_i(x_i).$$

Since $U_i(x_i) = \max_{z_i}\{m_i(z_i) - x_i q_i(z_i)\}$, $U_i(x_i)$ is a convex function, so it is absolutely continuous and differentiable almost everywhere. Whenever $U_i(x_i)$ is differentiable, $U_i'(x_i) = -q_i(x_i)$, and, consequently, $U_i(x_i) = U_i(0) - \int_0^{x_i} q_i(t_i) dt_i$.

The second line follows from the definition of $U_i(x_i)$ and the fact that $m_i(0) = U_i(0)$. □

Claim 4.2 *A mechanism is incentive compatible if and only if all* q_i*s are nonincreasing.*

PROOF: It has been shown that $U_i'(x_i) = -q_i(x_i)$ almost everywhere. Since $U_i(x_i)$ is convex, $q_i(x_i)$ is nonincreasing.

For the converse direction, note that $U_i(z_i) \geq m_i(x_i) - q_i(x_i) z_i$ if and only if

$$U_i(z_i) \geq U_i(x_i) + (x_i - z_i) q_i(x_i),$$

or, using the expression for $U_i(t)$,

$$(x_i - z_i) q_i(x_i) \leq \int_{z_i}^{x_i} q_i(t_i) dt_i.$$

Clearly, if $q_i(x_i)$ is nonincreasing, this condition always holds. □

Now we can derive the requirements for the optimal mechanism, that is, the one that minimizes the total payment $\sum_{e_i \in E} \mathbf{E}[m_i(x_i)]$.

We say that the mechanism design problem is *regular* if for all i the function $x_i + \frac{F_i(x)}{f_i(x)}$ is nondecreasing. This requirement holds for a wide class of distributions, e.g., any distribution whose density function is nonincreasing. In what follows, we refer to $c_i(x_i) = x_i + \frac{F_i(x_i)}{f_i(x_i)}$ as the *virtual cost* of edge e_i.

Theorem 4.3 *If the mechanism design problem is regular, then the optimal mechanism is given by a pair of functions $(Q(\mathbf{x}), M(\mathbf{x}))$, where $Q(\mathbf{x})$ always picks the path with the smallest virtual cost, and*

$$M_i(\mathbf{x}) = Q_i(\mathbf{x})x_i + \int_x^{\omega_i} Q_i(t_i, \mathbf{x}_{-i})dt_i.$$

PROOF: By definition,

$$\mathbf{E}[m_i(x_i)] = \int_0^{\omega_i} m_i(x_i)f_i(x_i)dx_i = m_i(0) + \int_0^{\omega_i} x_i q_i(x_i)f_i(x_i)dx_i - \int_0^{\omega_i}\int_0^{x_i} q_i(t_i)dt_i f_i(x_i)dx_i.$$

Changing the order of integration,

$$\mathbf{E}[m_i(x_i)] = m_i(0) + \int_0^{\omega_i} x_i q_i(x_i)f_i(x_i)dx_i - \int_0^{\omega_i} (1 - F_i(t_i))q_i(t_i)dt_i,$$

or, rewriting,

$$m_i(0) + \int_0^{\omega_i} \left(x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}\right) q_i(x_i)f_i(x_i)dx_i = m_i(0) + \int_X \left(x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}\right) Q_i(\mathbf{x})f(\mathbf{x})d\mathbf{x}.$$

The mechanism designer is allowed to choose $Q_i(\mathbf{x})$ and $m_i(0)$ (by choosing $M_i(0)$) so as to minimize $\sum_{e_i \in E} E[m_i(x_i)]$, subject to the following constraints:

- For any $\mathbf{x}$, $Q(\mathbf{x})$ must specify a path in the graph, that is, for some path P from s to t, it must be that $Q_i(\mathbf{x}) = 1$ for all $e_i \in P$.
- Incentive compatibility, which is equivalent to $q_i(x_i)$ being a nonincreasing function of x_i for all i.
- Individual rationality, i.e. $U_i(x_i) \geq 0$ for all i, x_i.

Since $U_i(x_i) = U_i(0) - \int_0^{x_i} q_i(t_i)dt_i$, from the last condition we get $m_i(0) = U_i(0) \geq \int_0^{x_i} q_i(t_i)dt_i$ for all x_i. As $q_i(t_i) \geq 0$ for all values of t_i, this is equivalent to stipulating that $m_i(0) \geq \int_0^{\omega_i} q_i(t_i)dt_i$. Since our goal is to minimize the total payment, we can set $m_i(0) = \int_0^{\omega_i} q_i(t_i)dt_i$. Hence, the buyer's objective function can be rewritten as

$$\sum_i \int_0^{\omega_i} \left(x_i + \frac{F_i(x_i)}{f_i(x_i)}\right) q_i(x_i)f_i(x_i)dx_i = \int_X \left(\sum_i \left(x_i + \frac{F_i(x_i)}{f_i(x_i)}\right) Q_i(\mathbf{x})\right) f(\mathbf{x})d\mathbf{x}.$$

Since the first constraint requires the mechanism to pick a path in the graph, the best it can do for any fixed $\mathbf{x}$ is to choose a path P with the smallest virtual cost, i.e., the one that minimizes $\sum_{e_i \in P}(x_i + \frac{F_i(x_i)}{f_i(x_i)})$. Picking the optimal path at every point $\mathbf{x}$ will also minimize the average cost with respect to f, that is, the integral we consider.

Formally, fix a shortest path algorithm A, run A on the original graph using the values $x_i + \frac{F_i(x_i)}{f_i(x_i)}$ as the corresponding edge weights, and set $Q_i(\mathbf{x}) = 1$ if e_i is on the path chosen by A and $Q_i(\mathbf{x}) = 0$ otherwise.

It remains to verify that this path selection rule is incentive compatible. To see that, note that for any fixed $\mathbf{x}_{-i}$, increasing x_i does not increase the ith player's chances of winning, since the cost of all paths including this edge goes up (here we use the assumption that $x_i + \frac{F_i(x_i)}{f_i(x_i)}$ is nondecreasing). That is, as x_i grows, $Q_i(\mathbf{x})$ can go from 1 to 0, but not the other way around. Averaging over all possible values of $\mathbf{x}_{-i}$, we see that $q_i(x_i)$ is nonincreasing.

By Claim 9.1, the expected payment $m_i(x_i)$ is completely determined by $m_i(0)$ and $q_i(x)$. As we have seen, the optimal value of $m_i(0)$ is $\int_0^{\omega_i} q_i(t_i)dt_i$, so

$$m_i(x_i) = \int_0^{\omega_i} q_i(t_i)dt_i - \int_0^{x_i} q_i(t_i)dt_i + x_i q_i(x_i) = \int_{x_i}^{\omega_i} q_i(t_i)dt_i + x_i q_i(x_i).$$

Choosing

$$M_i(\mathbf{x}) = Q_i(\mathbf{x})x_i + \int_{x_i}^{\omega_i} Q_i(t_i, \mathbf{x}_{-i})dt_i$$

satisfies this condition. □

5 VCG vs. Optimal Auction

First, observe that the amount paid by the optimal mechanism to each agent is exactly this agent's threshold bid. That is, to compute this payment, it suffices to identify the cheapest alternate path in terms of the virtual costs and then perform some simple calculations. This implies that the payments to the winners can be computed very efficiently using the approach of [6], provided that it is easy to compute the virtual cost $c(x) = x + \frac{F(x)}{f(x)}$ and its inverse.

The virtual cost of an edge can be much larger than its actual cost. This is the main reason why for many natural distributions the optimal mechanism performs considerably better than VCG in terms of total overpayment. In particular, if $f(x)$ is a sharply declining function, $c(x)$ grows very quickly. Hence, the edges with even moderately high actual costs are going to have extremely high virtual costs. Now, the calculation of threshold bids involves raising the bid of just one edge by the total difference between the costs of two paths; this is exactly the kind of configuration that is 'penalized' by $c(x)$, so the payments to winning edges are going to be relatively low.

Consider, for example, the graph G_n that consists of two disjoint paths P and Q of length n, and suppose that all edge costs are distributed according to $f_{a,\epsilon}(x)$, where

$$f_{a,\epsilon}(x) = \begin{cases} 1-\epsilon & \text{if } x \in [0,1], \\ \epsilon/(a-1) & \text{if } x \in [1,a], \\ 0 & \text{if } x > a. \end{cases}$$

Set $a = O(n)$, $\epsilon = O(\frac{1}{n^2})$. Then with probability at least $1 - n\epsilon$ all agents bid at most 1 (and their corresponding virtual costs are at most 2). Whenever this happens, the threshold bid of each agent is at most 1, since the virtual cost that corresponds to bidding $x > 1$ is $\Omega(n^3)$, so the total overpayment is at most n. Even if this is not the case, the threshold bid of any agent is at most a, so the expected total overpayment is bounded by $(1 - n\epsilon)(n * 1) + n\epsilon(n * a) = O(n)$.

On the other hand, since for large enough n the variance of this distribution is at least C, where C is some constant independent of n, for the original costs $c(P)$, $c(Q)$, we have

$$\mathbf{E}[|c(P) - c(Q)|] = \sqrt{2\mathbf{Var}(P)} \geq \sqrt{2nC} = \Theta(\sqrt{n}).$$

Since the overhead paid by the original VCG mechanism to each agent on the winning path is precisely $|c(P) - c(Q)|$, the total overpayment is $\Theta(n\sqrt{n})$. Moreover, it is clear that this is true not just for this distribution, but for any distribution with finite variance that does not depend on n.

This example provides a separation of $\Theta(\sqrt{n})$ between the optimal mechansim and VCG. We can obtain a similar result for unbounded distributions (i.e., ones that allow arbitrarily high edge costs), most notably, the *halfnormal distribution* $HD(a)$ with density

$$f_a(x) = \frac{2a}{\pi} \exp(\frac{-a^2x^2}{\pi})$$

or the *exponential distribution* $ED(a)$ with density $f_\lambda(x) = \lambda e^{-\lambda x}$.

We ran a series of simulations on the graph G_n with $n = 2^k$, $k = 1, \dots, 9$, and the edge costs distributed according to $HD(1)$ and $ED(1)$. In both cases, we conducted both the respective optimal auction and the VCG auction and computed the average payments to an edge on the winning path. The results are shown in Fig. 2.

One can see that for both distributions the payment to each agent by the optimal auction is roughly $\log n$, and the payment by VCG is $\Theta(\sqrt{n})$. (The latter result is to be expected, of course, as these distributions have constant variance.) We present the proof for the case of the exponential distribution; the argument for the halfnormal distribution is similar.

Theorem 5.1 *In the optimal auction on the graph G_n that consists of two disjoint paths of length n, where the edge costs are distributed according to the exponential distribution with density $f(x) = e^{-x}$, the overpayment to each edge is $O(\log n)$.*

PROOF: For $f(x) = e^{-x}$, the corresponding virtual cost function $c(x) = x + \frac{F(x)}{f(x)}$ is equal to $x + \frac{1-e^{-x}}{e^{-x}} = e^x + x - 1$. Note that $e^x - 1 \leq c(x) \leq e^{x+1}$, and $c(x)$ is monotone increasing, so $c^{-1}(y) \leq \ln(y+1)$.

Consider an edge e on the winning path. Suppose that the edge costs on the losing path are $x_1, \dots, x_n$, and the edge costs on the winning path are $x'_1, \dots, x'_n$. Set $z_{2n} = \max\{x_1, \dots, x_n, x'_1, \dots, x'_n\}$. Clearly, the payment to e is at most

$$c^{-1}(c(x_1) + \ldots + c(x_n)) \leq c^{-1}(nc(z_{2n})) \leq \ln(nc(z_{2n}) + 1) \leq \ln n + z_{2n} + 2.$$

Hence, the expected payment to e is bounded from above by $\mathbf{E}[z_{2n}] + O(\log n)$. Let us compute $\mathbf{E}[z_{2n}]$. Since the edge costs are independent, we have

$$F(t) = \mathbf{P}r[z_{2n} < t] = (1 - e^{-t})^{2n}$$

$$f(t) = 2n(1 - e^{-t})^{2n-1} e^{-t},$$

and

$$\mathbf{E}[z_{2n}] = \int_0^\infty 2n(1-e^{-t})^{2n-1} e^{-t} t dt = -\int_0^1 2ny^{2n-1} \ln(1-y) dy =$$

$$\int_0^1 2ny^{2n-1}(y + \frac{y^2}{2} + \frac{y^3}{3} + \ldots) dy = 2n \int_0^1 (\sum_{k=1}^\infty \frac{y^{2n+k-1}}{k}) dy =$$

$$2n \sum_{k=1}^\infty \frac{1}{k(2n+k)} = \sum_{k=1}^\infty (\frac{1}{k} - \frac{1}{2n+k}) = O(\ln n).$$

This means that the expected payment to each agent on the winning path is $O(\log n)$ and the total overpayment is $O(n \log n)$. $\square$

5.1 Bounded costs

We have seen that the optimal mechanism can be very different from VCG. However, if we assume that for all i the edge costs are uniformly distributed on some fixed interval of the form $[0, a]$ (i.e., $f_i(x) = 1/a$ if $0 \leq x \leq a$ and 0 otherwise for all i), then a small modification of the VCG mechanism is actually optimal. To see that, note that in this case $c_i(x) = x + \frac{x/a}{1/a} = 2x$, so the mechanism always chooses the shortest path, and the payment to a winning edge is the highest bid at which it would still win the auction. However, when the costs are known to be in $[0, a]$, the bids higher than a are not allowed, so the payment to an agent is capped at a, while in the original model there is no such cap.

To understand the importance of this modification, consider again the graph G_n that consists of two disjoint paths of length n, and suppose that all edge costs are uniformly distributed on $[0, 1]$. As shown above, the total overpayment of the original VCG mechanism is $\Theta(n\sqrt{n})$.

On the other hand, the modified mechanism pays at most 1 to each agent, so the total payment (and the overpayment) is $O(n)$. More generally, if the costs of agent i are at most $\omega_i < \infty$, the expected premium paid by the optimal mechanism to agent i is

$$\int_0^{\omega_i} \int_{x_i}^{\omega_i} q_i(t_i) dt_i f_i(x_i) dx_i = \int_0^{\omega_i} q_i(t_i) F_i(t_i) dt_i.$$

Since $q_i(t_i) \leq 1$, $F_i(t_i) \leq 1$ for all i, t_i, the overpayment to each agent is at most ω_i, i.e., a constant, and the total overpayment is $O(n)$.

On a more pessimistic note, we have to mention that for many distributions, there is a matching lower bound of $\Omega(n)$. We give the proof for the case of the uniform distribution $U[0, a]$. By symmetry, it is easy to show that $q(\frac{a}{2}) = \frac{1}{2}$, and since $q(t)$ is nonincreasing, $q(t) \geq \frac{1}{2}$ if $t < \frac{a}{2}$. This implies

$$\int_0^a q(t) \frac{t}{a} dt \geq \int_0^{a/2} \frac{t}{2a} dt = \frac{a}{16}.$$

6 Conclusion

In the first part of the paper, we have shown that all mechanisms for the shortest path problem that possess dominant strategies are not frugal. We then analyzed the Nash equilibria for this problem and constructed the optimal mechanism for any regular cost distribution. It turns out that for a wide class of distributions, the overpayment by the optimal mechanism is much smaller that that by VCG. On the other hand, for certain distributions over bounded domains, small modifications of VCG can be optimal.

Furthermore, for some natural distributions, even the optimal mechanism may incur a constant factor overpayment. This suggests investigating other models for this problem, such as, for example, repeated games.

References

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

[2] E. Clarke, Multipart pricing of public goods. *Public Choice*, 8:17–33, 1971

[3] A.Fiat, A. Goldberg, J. Hartline, and A. Karlin, Competitive generalized auctions. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computation*, pages 72–81, 2002

[4] A. Goldberg, J. Hartline, and A. Wright, Competitive auctions and digital goods. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 735–744, 2001

[5] T. Groves, Incentives in teams. *Econometrica*, 41(4):617–631, 1973

[6] J. Hershberger, S Suri, Vickrey prices and shortest paths: what is an edge worth? In *Proceedings of the 42nd Symposium on Foundations of Computer Science*, pages 252–259, 2001

[7] V. Krishna, Auction Theory. Academic Press, 2002

[8] R. Myerson, Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981

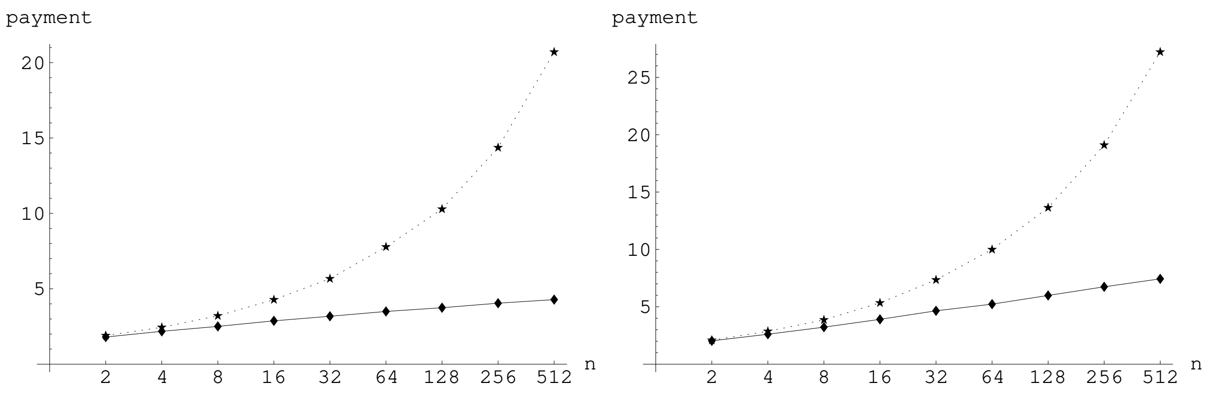

(a) Halfnormal distribution with mean 1

(b) Exponential distribution with mean 1

Figure 2: The lower curve represents the payments by the optimal auction; the upper curve represents the payments by VCG. Note that the x-axis is drawn to logarithmic scale.

[9] N. NISAN AND A. RONEN, Algorithmic mechanism design. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computation*, pages 129–140, 1999

[10] C. PAPADIMITRIOU, Algorithms, games, and the Internet, In *Proceedings of the 33st Annual ACM Symposium on Theory of Computation*, pages 749–753, 2001

[11] A. RONEN AND A. SABERI, On the hardness of optimal auctions. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, pages 396–405, 2002

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

A Proof of Theorem 4.2

PROOF: We start by proving a simple lemma:

Lemma A.1 *For any truthful mechanism on G, if an edge on the winning path lowers its bid, the threshold bids of all edges on that path do not decrease.*

PROOF: Suppose an edge e lowers its bid from $b(e)$ to $b'(e)$. First, by truthfulness, the threshold bid of e itself does not change. Second, note that the path containing e (without loss of generality, assume $e \in P$) still wins (otherwise, e – a losing edge – can cause itself to win by raising its bid).

Now, suppose that for some $x \in P$ its threshold bid went down from $t(x)$ to $t'(x)$. This means that if all other edges on P bid as before, e bids $b'(e)$, x bids $(t(x) + t'(x))/2 > t'(x)$, P loses. Raise e's bid to $b(e)$. By monotonicity, P still loses, even though every edge on P bids as in the original setting, except for x, who bids below $t(x)$, its original threshold bid. This contradicts the way $t(x)$ is defined. □

To prove the theorem, consider two possible cases:

- for any bid vectors b_P, b_Q such that $b(P) = L_2 - \epsilon'$, $b(Q) = L_2$, $\mathcal{M}$ chooses P over Q.
- there is a pair of bid vectors b_P, b_Q such that $b(P) = L_2 - \epsilon'$, $b(Q) = L_2$, and Q wins.

In the first case, suppose that P bids $b_P, b(P) = L_1$, and Q bids $b_Q, b(Q) = L_2$. By monotonicity, P wins, and for any $e \in P$ its threshold bid is at least $b(e) + (L_2 - L_1 - \epsilon')$: if e bids this amount, and all other edges on P bid according to b_P, then P's bid is $L_2 - \epsilon'$, so it still wins. Summing over all edges of P, we see that P is paid at least $b(P) + k(L_2 - L_1 - \epsilon')$. Note that $b(Q) - b(P) = L_2 - L_1 \geq \epsilon' > \epsilon'/k$.

In the second case, consider a pair of bid vectors b_P, b'_Q, where b'_Q is obtained from b_Q by decreasing the bids of some edges along Q so that $b'(Q) = L_2 - \epsilon' - \epsilon'/k$. By monotonicity, Q still wins, and the threshold bids of the edges along Q do not decrease. Hence, Q receives at least

$$b(Q) = L_2 = b(P) + \epsilon' \geq b'(Q) + \epsilon' = b'(Q) + k(b(P) - b'(Q)),$$

since we choose $b'(Q)$ so that $b(P) - b'(Q) = \epsilon'/k$. □

B Embedding Into a Bigger Graph - a Proof

It would be easy to embed the costly example of Section 3 into a bigger graph if we could prove the following claim: given an arbitrary graph that contains two node-disjoint paths P and Q between S and T, $|P| = n_1$, $|Q| = n_2$, we can assign weights to edges that do not belong to these paths so that

(1) for all $i = 1, \ldots, n_1$, $j = 1, \ldots, n_2$, whenever the edges along P bid according to b_P^i and the edges along Q bid according to b_Q^j, either P or Q wins;

(2) when the edges along P bid according to b_P^0 and the edges along Q bid according to b_Q^1, P wins;

(3) for all $t = 1, \ldots, l$ and $0 \leq \epsilon' \leq \epsilon$, when the edges along Q bid according to b_Q^1 and the edges along P bid according to b_P^0 except for the edge i_t that bids $\frac{L}{n_1} + \epsilon'$, P wins.

(As in Section 3, we assume that b_Q^1 is the vertex of the bipartite graph with the highest outdegree, and the endpoints of edges that leave it are $i_1, \ldots, i_l$.)

If these conditions hold, the argument of the Section 3 goes through without change, which means that the mechanism will have to pay at least $b(P) + \frac{1}{2}k(b(Q) - b(P))$.

However, if we do not make any extra assumptions about the mechanism, this may not be true. Consider a graph G that consists of four parallel arcs a, b, c, and d between S and T. Suppose that path selection rules are as follows:

- if $b(d) \leq b(c)$, $b(d) \leq b(b)$ or $b(b) > 1$, and $b(d) \leq b(a)$ or $b(a) > 1$, d wins;
- if $b(c) < b(d)$, $b(c) \leq b(b)$ or $b(b) > 1$, and $b(c) \leq b(a)$ or $b(a) > 1$, c wins;
- if $b(b) \leq 1$, $b(b) < b(d)$, $b(b) < b(c)$, and $b(b) \leq b(a)$ or $b(a) > 1$, b wins;
- if $b(a) \leq 1$, $b(a) < b(d)$, $b(a) < b(c)$, and $b(a) \leq b(b)$ or $b(b) > 1$, a wins.

Intuitively, the mechanism chooses the cheapest arc except that if a or b bid higher than 1, they drop out of competition (in case of ties, the mechanism prefers d over c over b over a).

It is not hard to check that this mechanism is truthful; moreover, it has all four desirable properties of truthful mechanisms mentioned in [1] (edge autonomy, path autonomy, independence, and sensitivity). This is still true if a and b are subdivided into n_1 and n_2 arcs, respectively (in this case, by $b(a)$ we mean the sum of all bids along a; same for b).

On the other hand, if we try to embed our high-cost example into G for $L > 1$, we see that no matter how high c and d bid, neither a nor b is going to win.

Nevertheless, we can prove that if any of the conditions (1)–(3) does not hold, then the mechanism can be forced to pay an arbitrary amount L' even if the graph contains a pair of paths of cost L and $L+\epsilon$. Setting $L' > \frac{1}{2}n\epsilon + L$ shows that this mechanism is as bad as any of those that allow to embed our costly example into the graph.

To see this, set the cost of any edge that does not belong to P or Q to L'. Now, suppose that condition (1) does not hold. This means that for some i, j, when P bids b_P^i and Q bids b_P^j, some path other than P or Q wins. As P and Q are node-disjoint, this path must contain an edge that costs L', so by voluntary participation property, the mechanism has to pay at least L'.

If (1) is true, but (2) does not hold, either Q or some path other than P and Q wins. The former case is not possible: otherwise, the losing edge i_1 can cause itself to win by raising its bid from $\frac{L}{n_1}$ to $\frac{L}{n_1}+\epsilon$ (by (1), when P bids $b_P^{i_1}$ and Q bids b_Q^1, P wins), which contradicts monotonicity. In the latter case, the winning path contains an edge of cost L', so the payment is at least L'.

If (1) is true, but (3) does not hold, the argument is identical.