

How Pervasive Is the Myerson-Satterthwaite Impossibility?*

Abraham Othman and Tuomas Sandholm

Computer Science Department

Carnegie Mellon University

{aothman, sandholm}@cs.cmu.edu

Abstract

The Myerson-Satterthwaite theorem is a foundational impossibility result in mechanism design which states that no mechanism can be Bayes-Nash incentive compatible, individually rational, and not run a deficit. It holds universally for priors that are continuous, gapless, and overlapping. Using automated mechanism design, we investigate how often the impossibility occurs over discrete valuation domains. While the impossibility appears to hold generally for settings with large numbers of possible valuations (approaching the continuous case), domains with realistic valuation structure circumvent the impossibility with surprising frequency. Even if the impossibility applies, the amount of subsidy required to achieve individual rationality and incentive compatibility is relatively small, even over large unstructured domains.

1 Introduction

The Myerson-Satterthwaite theorem [1983] belongs to a seminal line of impossibility results in mechanism design. Its relatives include the results of Arrow [1970], Gibbard-Satterthwaite [1973; 1975], and Green-Laffont [1977]. These theorems begin by positing a set of *prima facie* reasonable desiderata, and conclude by proving the impossibility of satisfying those desiderata together.

The Myerson-Satterthwaite theorem states that no mechanisms exist that do not run a deficit, are (*ex post*) efficient, Bayes-Nash incentive compatible, and (*ex interim*) individually rational. It is one of the most important results in mechanism design. It is also important to the field of political economics, because it serves as a negative counterweight to the famous Coase theorem [1960]. First proposed by Ronald Coase in 1937, before the advent of game theory, the “theorem” claims that, with zero-cost access to lawsuits and bargaining parties will establish socially efficient outcomes among themselves. Thus, a free market will lead to an efficient allocation of goods. In contrast, the Myerson-Satterthwaite theorem shows that this ceases to be the case

*This material is based upon work supported by the National Science Foundation under grant IIS-0427858. We also acknowledge Intel Corporation for a machine gift.

with two-sided private information. The impossibility holds even with just two players, a buyer and a seller, who hold private valuations of an item. This impossibility was one of the main citations in Myerson’s 2007 Nobel prize in economics.

However, the assumptions of the theorem are suspicious. The impossibility result holds only when the buyer and seller hold gapless continuous priors over overlapping sets of values. Obviously, we cannot actually represent real numbers exactly in a finite-sized field. Normally, we simply brush over this distinction, because an arbitrary approximation to a real number is good enough. But because the Myerson-Satterthwaite theorem is an *impossibility* result, rather than a constructive proof, we are right to be suspicious about its meaningfulness in a discrete real-world setting. Moreover, the proof of the theorem relies heavily on specific properties of real numbers such as density and differentiability and is therefore not immediately convertible to a proof over a discrete finite set (e.g., by replacing integrals with sums and density functions with mass functions). It is well-known that there exist simple discrete settings in which the Myerson-Satterthwaite impossibility does not hold (see Section 3). However, to our knowledge, this paper is the first to examine how frequently the impossibility holds.

Our exploration can be regarded as somewhat analogous to work examining how often voting profiles are manipulable—that is, how often the Gibbard-Satterthwaite impossibility pertains [Blair, 1981; Friedgut *et al.*, 2008]. Though both that line of work and this one consider manipulation over a discrete, finite set of reports, the key difference is that our setting has money (quasi-linear quantitative utility functions as preferences). While that work on voting has shown that manipulable instances occur with high frequency, we find that considerable possibility exists, even over large domains.

Certainly we should expect differences when moving from continuous to discrete cases. But the violation of impossibility results is the most jarring of these differences, particularly when building real systems. As our study shows, it may be entirely possible to construct a mechanism that subverts the Myerson-Satterthwaite impossibility, opening up efficient bilateral trades even for rational, self-interested agents.

2 Testing Feasibility

Using the approach of *automated mechanism design* [Conitzer and Sandholm, 2002], we can establish a linear

program (LP) which encodes the set of constraints that define a mechanism with properties relevant to the Myerson-Satterthwaite impossibility. Let the buyer have valuations drawn from a discrete set $B = \{b_1, \dots, b_k\}$ and the seller have valuations drawn from a discrete set $S = \{s_1, \dots, s_k\}$. We will call the elements of these sets *report points*. We use $\mathbb{E}_X(f)$ to represent the expectation of f over the random variable X (e.g., $\mathbb{E}_B(b)$ represents the expected value of the bidder).

The goal of the LP (given below) is to determine whether, given a set of report points and common-knowledge priors about the probabilities of agents having the valuations of those report points, it is feasible to create a mechanism that satisfies the desiderata in the Myerson-Satterthwaite impossibility result. Here our decision variables are $2k^2$ terms corresponding to *payments*, $p(\cdot, \cdot)$, by the buyer, and *receipts*, $r(\cdot, \cdot)$, of the seller.

$$\begin{aligned} \min \quad & \sum_{b \in B, s \in S} p(b, s) - r(b, s) \\ \text{Subject to:} \quad & t(b, s) = (b \geq s) \\ & p(b, s), r(b, s) \geq 0 \\ & p(b, s) \geq r(b, s) \quad \forall b, s \\ \mathbb{E}_S(t(b, s)b - p(b, s)) & \geq 0 \quad \forall b \\ \mathbb{E}_S(t(b, s)b - p(b, s)) & \geq \mathbb{E}_S(t(b', s)b - p(b', s)) \quad \forall b' \neq b \\ \mathbb{E}_B(r(b, s) - t(b, s)s) & \geq 0 \quad \forall s \\ \mathbb{E}_B(r(b, s) - t(b, s)s) & \geq \mathbb{E}_B(r(b, s') - t(b, s')s) \quad \forall s' \neq s \end{aligned}$$

Our objective function is to minimize the amount of *money burnt* by the system, that is, the difference between the payments made by the buyer and the payments received by the seller. We refer to this money being burnt because it must be used in a way that does not impact the utility of either the buyer or the seller. Though it might be a surprise, money burning is an essential feature of the feasible mechanisms generated in many of our trials. This LP will find a solution that does not involve money burning, if one exists. If the constraints of the LP cannot be satisfied, the instance is *infeasible*, i.e., the impossibility applies. On the other hand, if the LP returns a positive objective value, then the instance is feasible, but only by burning money. Finally, if the LP returns a zero objective value, the instance is feasible and no money needs to be burned.

The quantities $t(b, s)$ in the LP are constants that describe whether or not the trade occurs. Because we are interested in efficient mechanisms, the constants are set as follows. The constant has a value of 1 if the buyer has a higher valuation than the seller (i.e., it is present when taking expectations). Otherwise, it has a value of 0 (i.e., it is not present when taking expectations). As we will describe later, we stagger the report points of the buyer and the seller to avoid the theoretical complications that arise as to what should happen if the buyer and seller were to report identical valuations.

The constraint $p(b, s) \geq r(b, s)$ represents the requirement that the mechanism be *ex post no-deficit*. Regardless of the actual values realized by the buyer and the seller, the seller will never receive more than the buyer pays out.

The constraint

$$\mathbb{E}_S(t(b, s)b - p(b, s)) \geq 0 \quad \forall b$$

and

$$\mathbb{E}_B(r(b, s) - t(b, s)s) \geq 0 \quad \forall s$$

represent *ex interim individual rationality (IR)* for the buyer and the seller, respectively. After the buyer or seller learns their own valuations, *ex interim IR* requires that they have a non-negative expected value from participating, where the expectation is taken over the reports of the other party. This condition is weaker than *ex post individual rationality*, which would imply:

$$t(b, s)b - p(b, s) \geq 0 \quad \forall b, s$$

for the buyer, and

$$r(b, s) - t(b, s)s \geq 0 \quad \forall b, s$$

for the seller. *Ex post IR* would insure that, for instance, the buyer is never compelled to make a payment without receiving the item, something which could theoretically occur under *ex interim IR*.

The constraints

$$\mathbb{E}_S(t(b, s)b - p(b, s)) \geq \mathbb{E}_S(t(b', s)b - p(b', s)) \quad \forall b' \neq b$$

and

$$\mathbb{E}_B(r(b, s) - t(b, s)s) \geq \mathbb{E}_B(r(b, s') - t(b, s')s) \quad \forall s' \neq s$$

represent *ex interim incentive compatibility (IC)* constraints for the buyer and the seller, respectively. *Ex interim IC* ensures that once an agent learns her own valuation, she will *in expectation* do no worse by reporting it truthfully than by reporting some other value.

As usual, we assume that all valuation distributions are common knowledge: the buyer and seller (and center) know the true distributions over the buyer's and seller's valuations.

3 A Simple Example

The following simple example shows why this problem is interesting and subtle. Consider a single buyer and a single seller negotiating over a single item. The buyer can have a valuation $b \in \{1, 3\}$, while the seller can have a valuation $s \in \{0, 2\}$. When $\mathbb{P}(b = 3) = 1$ and $\mathbb{P}(s = 0) = 1$, the Myerson-Satterthwaite impossibility does not hold. For example, a mechanism where the buyer pays a fixed amount $x \in [0, 3]$ and the seller receives that amount x satisfies all of the Myerson-Satterthwaite desiderata and also does not burn money.

Now consider implementation in dominant-strategy equilibrium, as opposed to Bayes-Nash equilibrium. Dominant strategy equilibrium would make it a best response for each agent to report her valuation truthfully without regard to the report (or distribution over potential reports) of the other agent. If all possible reports have positive probability, it is impossible to construct an efficient, no-deficit, dominant-strategy mechanism. Such a mechanism would, by incentive compatibility and individual rationality, imply that $p(1, 0) = p(3, 0) < 1$, and that $r(3, 0) = r(3, 2) > 2$. A contradiction

follows from noting that $p(3, 0) \not\geq r(3, 0)$. As a result, feasibility in dominant strategies occurs only in a space of measure 0.

To measure feasibility in Bayes-Nash equilibrium, we drew 30,000 random samples from the square defined by $(x, y) = (\mathbb{P}(b = 1), \mathbb{P}(s = 0))$ and plotted a dot where a draw was feasible. Figure 1 shows the results of that experiment. As the prevalence of the dots indicates, approximately 85% of the instances showed feasibility. The infeasible region is characterized by the buyer having a high probability of having a high valuation (small value on the x axis) while the seller has a high probability of having the lower valuation (large value on the y axis). Note that when the seller has the low value or the buyer has the high value, trade must occur.




Figure 1: Dots indicate feasibility in a simple two-valuation model. The x axis represents the probability the buyer has his low valuation, and the y axis represents the probability the seller has his low valuation.

To consider why infeasibility might occur in these circumstances, assume that the seller has a high likelihood of having the lower value. Speaking broadly, one must have the payments of both high- and low-value buyers be approximately equal and smaller than 1. Now, assume that the buyer has a high likelihood of having the higher value. The sales price for the seller reporting both the low and high value must be approximately the same and higher than 2. One can see that these two situations—where the buyer has a high likelihood of having a high value, and the seller has a high likelihood of having a low value—create mutually infeasible constraints on the no-deficit condition, given that the seller needs to be compensated more than the buyer is willing to pay. Consequently, the northwest corner of the plot is barren, indicating no feasible instances.

4 Experimental Results for Feasibility

We now move beyond that simple example, and into a world where the buyer and seller draw their valuations from distributions with multiple support points. We consider $b \in \{0/(2k - 1), 2/(2k - 1), \dots, 2(k - 1)/(2k - 1)\}$, and $s \in \{1/(2k - 1), 3/(2k - 1), \dots, 2k - 1/(2k - 1)\}$. There are k report points available to the buyer, and k separate report points available to the seller. These report points are the only places at which agents have non-zero probability of having valuations.

Since we cannot make any blanket statements about feasibility—we could arrange probability masses to create both feasible and infeasible instances—we use experimental simulation to determine how pervasive the Myerson-Satterthwaite impossibility is.

4.1 Selection of Priors

In our study, we created five different schemes for generating the probabilities over report points. One scheme, which we dub *arbitrary*, is based around straightforward independently uniform generation of probabilities at report points. For the other schemes, we model valuations according to underlying continuous distributions, where we generate probabilities at report points by rounding *down* the buyer's valuation to the nearest report point, and rounding *up* the seller's valuation. (This was inspired by the optimal mechanism for restricted-communication agents [Blumrosen *et al.*, 2007].) We explore the following schemes:

Uniform The underlying distribution for both the buyer and the seller is the uniform distribution on $[0, 1]$. Note that there are no random draws associated with a uniform instance.

Arbitrary Probabilities at each report point are drawn uniformly from $[0, 1]$. Then the probability tuples for both the buyer and seller are scaled so that they sum to one.

Independent Normal The buyer's and the seller's underlying distributions are normal distributions with mean for each selected uniformly in $[0, 1]$ and standard deviation .25.

Identical Normal The buyer's and the seller's underlying distributions are identical normal distributions, with mean selected uniformly on $[0, 1]$ and standard deviation .25.

Opposite Normal The buyer's underlying valuations are a normal distribution with mean x uniform on $[0, 1]$ and standard deviation .25. The seller's underlying distributions are normal with mean $1 - x$ and standard deviation .25.

Since the normal distribution has unbounded domain, we scale the report probability tuples for both agents so that they sum to one.

4.2 Results

In all cases with repeated experiments (that is, the non-uniform distributions that involve random draws), we ran trials over 5,000 instances. Those feasible instances with positive objectives (indicating a payment higher than a receipt,

that is, money being burned) are tallied under “Instances with Burns”. Our results are discussed in the following subsections and plotted in Figures 2 and 3.

Uniform Prior

Since both buyers’ and sellers’ valuations are drawn uniform on $[0, 1]$, and buyers round down while sellers round up, we have:

$$\mathbb{P}(b_i) = \mathbb{P}(s_j) = 2/(2k - 1)$$

for $i \neq k - 1$ and $j \neq 0$. Also,

$$\mathbb{P}(b_{k-1}) = \mathbb{P}(s_0) = 1/(2k - 1)$$

We found that the LP induced by uniform valuations was feasible for $k < 5$ and infeasible for $k \geq 5$. There were no money burns associated with the uniform valuation instances.

Arbitrary Priors

k	Fraction Feasible	Instances with Burns
3	.974	102
4	.652	68
5	.255	41
6	.063	17
7	.022	5
8	.007	1
9	.004	1
10	.001	1

Independent Normal Priors

k	Fraction Feasible	Instances with Burns
3	.994	139
4	.874	124
5	.702	149
6	.580	158
7	.472	114
8	.431	103
9	.375	140
10	.317	155

Identical Normal Priors

k	Fraction Feasible	Instances with Burns
3	1.000	29
4	1.000	57
5	1.000	212
6	.801	218
7	.479	88
8	.307	62
9	.206	55
10	.141	39

Opposite Normal Priors

k	Fraction Feasible	Instances with Burns
3	.973	127
4	.734	132
5	.565	134
6	.514	158
7	.464	175
8	.452	186
9	.447	161
10	.337	141

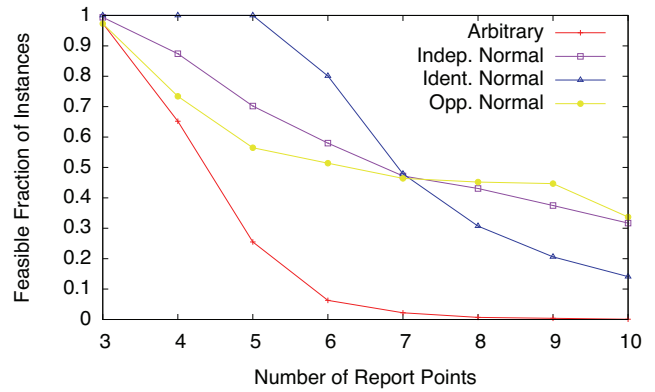


Figure 2: The fraction of feasible instances decreases in k .

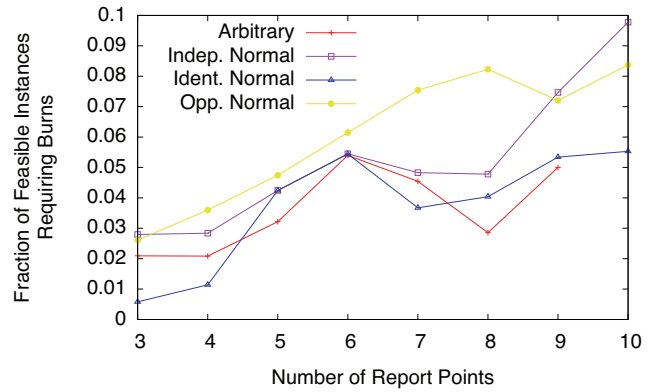


Figure 3: The fraction of feasible instances requiring money to be burnt increases in k .

4.3 Discussion

Our most obvious result is that feasibility falls with k . For some distribution schemes, feasibility drops off sharply. For others, like the opposite normal and the independent normal schemes, feasibility falls off more gradually. With k report points there are $\Theta(k^2)$ constraints that must be satisfied in the LP and $\Theta(k^2)$ variables. From a high-level perspective, this helps explain the decrease in feasibility with k . Still, some distributions showed surprisingly high numbers of feasible instances even with $k = 10$.

The arbitrary distribution had a particularly steep decline in feasibility in k . To explain why, consider the toy example of the previous section. Infeasibility was prevalent when there was a high chance the seller would have a low valuation, and a high chance the buyer would have a high valuation. The expected value for the largest of k uniformly distributed numbers is $k/k + 1$, the expected value for the smallest of k uniformly distributed numbers is $1/k + 1$, and the expected value of the j -th smallest is $j/k + 1$. As a result, as k increases there are not only more high values (relative to the other probabilities), but there are also more slots at which placing those values will cause infeasibility.

Perhaps the most realistic scheme, with underlying independent normal distributions, had a high fraction of feasible

instances, even with $k = 10$. However, many instances were feasible because they involved limited possibilities for trade. Recall again from our toy example that infeasibility was associated with a trade almost certainly taking place—the seller having a low valuation or the buyer having a high valuation. In contrast, situations where trade was unlikely to take place—the buyer having a low valuation and the seller having a high valuation—were associated with feasibility. Limiting the valuation points that yield trades is the simplest way to get a feasible instance.

We believe this explains the behavior seen in the opposite normal distributions. That scheme had a rapid decrease in feasible instances up to $k = 5$, but then the number of feasible instances decreased very slowly, to the point that at $k = 10$ the opposite normal distribution had the most feasible instances of any of the distribution schemes we tested. Note that the opposite normal scheme either consists of instances in which trade is to be expected (when the seller has the smaller of the two distributional means) to instances in which trade is not to be expected (when the seller has the larger of the two distributional means). Those situations in which trade is certain to occur lose feasibility very quickly in k , while a surfeit of report points does not adversely affect feasibility if trade is unlikely to take place. We explore the opposite normal distribution in more detail in the next section.

Money burning was most prevalent when the supports of the priors had many point (i.e., large k). That is also where the probability of feasibility was the lowest. Our LP was designed to minimize money burning, and therefore the only instances tallied in the “Instances with Burns” columns are those which would not have been feasible without burning money. Since money burning is normally associated with the setting of payment redistribution (e.g., [Cavallo, 2006; Guo and Conitzer, 2007]), it is perhaps surprising that a not-insignificant fraction of feasible instances were feasible only through burning money. Money burning provides additional flexibility in variable assignment in the LP: without money burning, $r(b, s) = p(b, s)$ but with money burning $r(b, s) \in [0, p(b, s)]$. When the supports of the priors have many points, this extra flexibility makes the difference for an increasing number of instances.

5 Dealing with Infeasibility

We further investigate the nature of infeasibility by relaxing some constraints in the LP. Of the three classes of constraints for efficient mechanisms (incentive compatibility, individual rationality, and no-deficit), only relaxing no-deficit still encourages participation and truth-telling. The center giving *ex post* subsidies ensures predictable agent action and provides a simple metric (i.e., the amount of subsidy) for measuring just how infeasible instances are. This measure of infeasibility is actually identical to the one used in the original proof of the Myerson-Satterthwaite impossibility for continuous supports.¹

¹Another relaxation which we do not consider here would be to relax the requirement of efficiency (perhaps by treating the t 's in the LP as a probability). However, retaining efficiency is more in keeping with the proof of the Myerson-Satterthwaite result, and also

5.1 An LP for Minimum Subsidies

To measure how much subsidy is needed on various problem instances, we used the following LP:

$$\begin{aligned}
 & \min \mathbb{E}_{B,S}(\zeta_{ND}) \\
 \text{Subject to: } & t(b, s) = (b \geq s) \\
 & p(b, s), r(b, s) \geq 0 \\
 & p(b, s) + \zeta_{ND}(b, s) \geq r(b, s) \quad \forall b, s \\
 & \mathbb{E}_S(t(b, s)b - p(b, s)) \geq 0 \quad \forall b \\
 & \mathbb{E}_S(t(b, s)b - p(b, s)) \geq \mathbb{E}_S(t(b', s)b - p(b', s)) \quad \forall b' \neq b \\
 & \mathbb{E}_B(r(b, s) - t(b, s)s) \geq 0 \quad \forall s \\
 & \mathbb{E}_B(r(b, s) - t(b, s)s) \geq \mathbb{E}_B(r(b, s') - t(b, s')s) \quad \forall s' \neq s
 \end{aligned}$$

Now in addition to the $p(\cdot, \cdot)$ and $r(\cdot, \cdot)$ variables, we also have the (linear) $\zeta_{ND}(\cdot, \cdot)$ variables, which represent subsidies from the center to the agents. Our objective is to minimize the expected subsidy given to the agents.

The objective value can sometimes be negative, so in expectation the center can actually *make money*. To see how this could be, consider the phenomenon of money burning, where the buyer pays more than the seller receives. The center can abscond with the difference.

Note also that the objective value being non-positive does not necessarily imply a mechanism that satisfies the desiderata of the Myerson-Satterthwaite impossibility. The Myerson-Satterthwaite theorem posits *ex post* no-deficit, which in this context implies

$$\zeta_{ND}(b, s) \leq 0 \quad \forall b, s$$

A probability distribution over report points that induces a program with non-positive objective instead corresponds to a mechanism that is *ex ante* no-deficit, so that only in *expectation* does the center expect to not lose money.

5.2 Results

We ran 5,000 trials for distributions corresponding to the five schemes from the previous section, for number of report points $k \in \{5, 10, 15, 20\}$. We averaged the expected subsidies from the trials, and these averages are reported in the table below and plotted in Figure 4.

k	Uniform	Arbitrary	Independent Normal	Identical Normal	Opposite Normal
5	.017	.043	.030	-.007	.071
10	.079	.104	.106	.046	.150
15	.105	.123	.136	.074	.191
20	.119	.134	.150	.089	.203

Table 1: *Expected subsidies for each prior scheme.*

5.3 Discussion

Just as the fraction of feasible instances decreased with k , here we see subsidies increasing in k . However, for some valuation distributions, even with a fairly fine support of valuations, the expected subsidy was relatively small. This was particularly the case for the identical normal scheme, where more relevant for its impact on the Coase theorem.

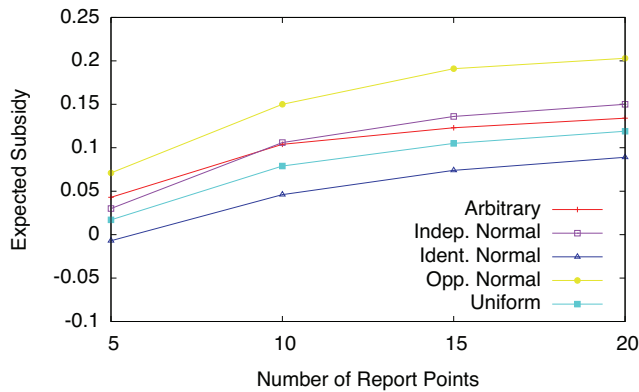


Figure 4: *The expected subsidy increases in k .*

with 20 report points for each agent the subsidy was only .089. Additionally, with five report points the identical normal scheme had a *negative* subsidy, indicating that the center could expect to actually generate revenue by skimming from the payment.

Subsidies are not only a measure of how often instances are feasible (feasible instances imply solutions with a non-positive objective value) but also a measure of just how infeasible infeasible instances are. This is why the opposite normal distribution needed such a significant subsidy. Recall from our results above that the opposite normal distribution had the highest fraction of feasible instances for large k , which we suggested was due to instances in which buyers had a lower mean than sellers and so trade was unlikely, making the instances more likely to be feasible. The other side of that argument is that the opposite normal scheme will have a large number of instances in which the buyer will have a higher value than the seller and trade will be expected. As a result, much more subsidy is required in these situations.

When there is a large gap in valuation between the seller and the buyer, potentially a large subsidy is needed. Of course, when there is a small gap in valuations, less subsidy is required. We suggest that this is why the identical normal distribution scheme required less subsidy than the other schemes. In other words, if the buyer and seller have similar distributions over valuations, smaller subsidies sufficed.

6 Conclusion and Future Directions

Though the Myerson-Satterthwaite theorem holds no sway over situations with discrete values, its presence was a specter haunting our experimental inquiry. As we increased the number of points at which agents could have valuations, we experienced a falloff in the number of feasible instances. In some schemes this falloff was dramatic, in others, slower. Regardless, the Myerson-Satterthwaite impossibility did hold more than 65% of the time with ten possible valuation points for all of the schemes of generating distributions we tried. When agents have large numbers of possible valuations, our work suggests it is unlikely that a mechanism satisfying the Myerson-Satterthwaite desiderata can be crafted.

When we loosened our feasibility LP to allow payments

from the center to the buyer, we found that the expected subsidy paid out was generally small, and that in one case, where buyers and sellers have identical normal distributions over five report points, the center could actually be expected to make a profit from the exchange. But on some instances—particularly when buyers and sellers had divergent valuations where buyers had much higher expected valuations—the subsidies required were quite large.

It might be of interest to investigate subsidies further in more automated mechanism design problems, particularly with the insight that skimming money from good inputs could compensate, in expectation, for subsidies required from bad inputs, while still maintaining incentive compatibility and individual rationality. In any case, using automated mechanism design to incorporate prior information can ameliorate, or even eliminate, the pain entailed by impossibility results.

References

- [Arrow, 1970] K.J. Arrow. *Social Choice and Individual Values*. Yale University Press, 1970.
- [Blair, 1981] D.H. Blair. On the Ubiquity of Strategic Voting Opportunities. *International Economic Review*, 22(3):649–655, 1981.
- [Blumrosen *et al.*, 2007] Liad Blumrosen, Noam Nisan, and Ilya Segal. Auctions with severely bounded communication. *J. Artif. Intell. Res. (JAIR)*, 28:233–266, 2007.
- [Cavallo, 2006] R. Cavallo. Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments. In *AAMAS*, 2006.
- [Coase, 1960] RH Coase. The Problem of Social Cost. *The Journal of Law and Economics*, 3(1):1, 1960.
- [Conitzer and Sandholm, 2002] Vincent Conitzer and Thomas Sandholm. Complexity of mechanism design. In *UAI*, 2002.
- [Friedgut *et al.*, 2008] Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *FOCS*, 2008.
- [Gibbard, 1973] Allan Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- [Green and Laffont, 1977] J Green and J-J Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
- [Guo and Conitzer, 2007] Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments. In *EC*, 2007.
- [Myerson and Satterthwaite, 1983] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- [Satterthwaite, 1975] Mark Satterthwaite. Strategy-proofness and Arrow’s conditions: existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.