

Approximation-Variance Tradeoffs in Facility Location Games

Ariel D. Procaccia

Computer Science Department
Carnegie Mellon University

David Wajc

Computer Science Department
Carnegie Mellon University

Hanrui Zhang

Department of Computer Science
Duke University

Abstract

We revisit the well-studied problem of constructing strategyproof approximation mechanisms for facility location games, but offer a fundamentally new perspective by considering *risk averse* designers. Specifically, we are interested in the tradeoff between a randomized strategyproof mechanism’s approximation ratio, and its variance (which has long served as a proxy for risk). When there is just one facility, we observe that the social cost objective is trivial, and derive the optimal tradeoff with respect to the maximum cost objective. When there are multiple facilities, the main challenge is the social cost objective, and we establish a surprising impossibility result: under mild assumptions, no smooth approximation-variance tradeoff exists. We also discuss the implications of our work for computational mechanism design at large.

1 Introduction

A *facility location game* consists of n players who are located on the real line; x_i denotes the location of player i . A mechanism f takes the vector of player locations $\mathbf{x} \in \mathbb{R}^n$ as input, and outputs a vector of k facility locations $\mathbf{y} \in \mathbb{R}^k$. The facilities are usually thought of as public goods, such as libraries or police stations, but the facility location setting can be interpreted in many other ways, e.g., player locations can represent opinions on a (quantitative) political spectrum, and a facility can be a policy choice. The *cost* of player i is her distance from the nearest facility, that is, $\min_{\ell \in [k]} |x_i - y_\ell|$. We wish to minimize one of two natural objectives: the utilitarian objective of *social cost*, which is the sum of individual costs; and Rawlsian objective of *maximum cost*, which is, obviously, the maximum individual cost.

However, naive optimization of these objectives may lead to undesirable strategic behavior on the part of the players. For example, the optimal solution for the case of $k = 1$ (a single facility), and the maximum cost objective, is to place the facility at the average of the leftmost and rightmost player locations, that is, at $(\min_i x_i + \max_i x_i)/2$. The problem is that, say, the rightmost player can drag the facility towards her true location by reporting a location that is further to the right, thereby decreasing her cost. The goal is,

therefore, to design facility location mechanisms that optimize the foregoing objectives, and are also *strategyproof*, in the sense that no player can decrease her cost by misreporting her location.

This challenge is the original and paradigmatic instance of *approximate mechanism design without money* (Procaccia and Tennenholtz 2013), an agenda that focuses on problems where monetary transfers are not allowed, which is why the need for approximation typically stems from strategic considerations (the optimal solution is not strategyproof) rather than computational complexity. Procaccia and Tennenholtz advocate using the approximation ratio of a strategyproof mechanism (the worst-case ratio between the objective value of the mechanism’s solution and the optimal solution) to quantify the solution quality that must inevitably be sacrificed in order to achieve strategyproofness. The design of strategyproof approximation mechanisms for facility location has been extensively studied (Procaccia and Tennenholtz 2013; Alon et al. 2010; Lu et al. 2010; Nissim, Smorodinsky, and Tennenholtz 2012; Thang 2010; Fotakis and Tzamos 2010; 2013a; 2013b; Cheng, Yu, and Zhang 2013; Wilf and Feldman 2013; Feldman, Fiat, and Golomb 2016; Golomb and Tzamos 2017), and, in particular, has been a topic of significant interest in recent AI conferences (Todo, Iwasaki, and Yokoo 2011; Zou and Li 2015; Serafino and Ventre 2015; Filos-Ratsikas et al. 2015; Cai, Filos-Ratsikas, and Tang 2016).

Our point of departure from this dense literature is that we re-examine the assumptions underlying randomized strategyproof mechanisms, which are known to provide better guarantees than their deterministic counterparts (Procaccia and Tennenholtz 2013). Specifically, in line with the literature on randomized approximation algorithms in general, previous work measures the *expected* objective value of a randomized mechanism, and disregards its *variance*. However, a *risk-averse* designer would be concerned with both. In fact, expectation-variance analysis has long been viewed as one of the fundamental approaches to reasoning about risk aversion, and nowadays it is ubiquitous in economics and finance (Markowitz 1952). In our case, given two distributions over facility locations with the same expected objective value, the designer should prefer the one with lower risk (variance); and may prefer a distribution with higher risk only if that risk is offset by sufficiently lower expected

objective value (for a minimization objective). The optimal distribution depends on the designer’s individual level of risk aversion, as well as on the optimal tradeoff between expected objective value and risk.

We therefore aim to characterize the optimal tradeoff between approximation (equivalently, expectation) and variance in facility location games. Formally, our research question is:

Given $\gamma \in \mathbb{R}^+$, what is the optimal approximation ratio achievable by a strategyproof (randomized) facility location mechanism whose variance is at most γ ?

We believe this question is important for two reasons. First, it provides a fundamentally new viewpoint on facility location games. Second, it can serve as a starting point for a broader investigation of expectation-variance tradeoffs in mechanism design, as we discuss in §4.

1.1 Our Results

In §2, we study the case of a single facility. For the social cost objective, placing the facility on the median reported location is strategyproof, optimal, and deterministic (so the variance of the social cost is 0). We focus, therefore, on the maximum cost objective.

We define a family of mechanisms, parameterized by $\alpha \in [0, 1/2]$, which includes the LEFT-RIGHT-MIDDLE (LRM) Mechanism of Procaccia and Tennenholtz (2013) as a special case. Informally, given a location profile $\mathbf{x} \in \mathbb{R}^n$, the GENERALIZED-LRM $_{\alpha}$ Mechanism chooses uniformly at random among four potential facility locations: leftmost player location, rightmost location, and two locations whose distance from the optimal solution depends on α . We prove:

Theorem 2.3 (informally stated). *For all $\alpha \in [0, 1/2]$, GENERALIZED-LRM $_{\alpha}$ is a (group) strategyproof mechanism for the 1-facility location problem. Moreover, on location profile $\mathbf{x} \in \mathbb{R}^n$, the expectation of its maximum cost is $(3/2 + \alpha) \cdot \text{opt}(\mathbf{x})$ (that is, its approximation ratio is $3/2 + \alpha$), and the standard deviation of its maximum cost is $(1/2 - \alpha) \cdot \text{opt}(\mathbf{x})$.*

Theorem 2.3 is especially satisfying in light of the next theorem — our first major technical result — which implies that GENERALIZED-LRM(α) gives the *optimal* approximation-variance tradeoff for the maximum cost objective.

Theorem 2.4 (informally stated). *For any strategyproof mechanism for the 1-facility location problem with the maximum cost objective, given a location profile $\mathbf{x} \in \mathbb{R}^n$, if the mechanism’s maximum cost has standard deviation at most $(1/2 - \alpha) \cdot \text{opt}(\mathbf{x})$, then its expected maximum cost is at least $(3/2 + \alpha) \cdot \text{opt}(\mathbf{x})$. In other words, the sum of expectation and standard deviation is at least $2 \cdot \text{opt}(\mathbf{x})$.*

In §3, we explore the case of multiple facilities. This time it is the maximum cost objective that is less challenging: We observe that the best known approximation ratio for any number of facilities $k \geq 2$ is given by a randomized mechanism of Fotakis and Tzamos (2013b), which (miraculously) happens to have zero variance.

Next we consider the social cost objective, and things take a turn for the strange: Our second major result asserts that, in this setting, a “reasonable” approximation-variance tradeoff simply does not exist, even when there are just two facilities.

Theorem 3.1 (very informally stated). *For the 2-facility location problem with the social cost objective, there is no family of mechanisms f_{θ} for every $\theta \in [0, 1]$ that satisfies two mild technical conditions, and smoothly interpolates between zero variance and constant approximation ratio, i.e., which satisfies the following properties: (i) f_0 has a constant approximation ratio, (ii) the variance of the social cost decreases monotonically with θ , down to zero variance at f_1 , and (iii) f_{θ} changes continuously with θ .*

Importantly, for the case of 2 facilities, deterministic strategyproof mechanisms are severely limited (Fotakis and Tzamos 2013a), but a randomized strategyproof 4-approximation mechanism is known (Lu et al. 2010). Our initial goal was to provide an approximation-variance tradeoff with this mechanism on one end, and a bounded deterministic mechanism on the other, but surprisingly, Theorem 3.1 rules this out.

1.2 Related Work

We are aware of only a single paper in computational mechanism design that directly studies variance (Esfandiari and Kortsarz 2016), in the context of kidney exchange. In contrast to our paper, it does not investigate the tradeoff between variance and approximation. Rather, the main result is a mechanism whose approximation ratio matches that of a mechanism of Ashlagi et al. (2015), but has lower variance.

Bhalgat, Chakraborty, and Khanna (2012) study multi-unit auctions with risk-averse sellers, where risk aversion is modeled as a concave utility function. They design polynomial-time strategyproof mechanisms that approximate the seller’s utility under the best strategyproof mechanism. The results depend on the notion of strategyproofness in question, and whether the buyers are also risk averse; in one case Eso and Futó (1999) have previously shown how to achieve the maximum utility. This work is different from ours in many ways, but one fundamental difference is especially important to point out: The goal of Bhalgat, Chakraborty, and Khanna (2012) is to achieve utility as close as possible to that of the optimal strategyproof mechanism; in principle it is possible to achieve an approximation ratio of 1 by running the optimal mechanism itself (which incorporates the concave utility function of the seller) — the obstacle is computational efficiency. Crucially, there is no *tradeoff* in their setting. In contrast, in our setting the benchmark is the unconstrained optimum, and the smaller the allowed variance, the worse our approximation becomes; our goal is to quantify this tradeoff. Relatedly, Sundararajan and Yan (2017) also endow a risk-averse seller with a concave utility function, and seek to simultaneously provide an approximation to the optimal utility of any possible seller, independently of her specific utility function.

Further afield, there is a body of work in auction theory that studies optimal auctions for risk averse *buyers* (Maskin and Riley 1984; Bhalgat, Chakraborty, and Khanna 2012;

Fu, Hartline, and Hoy 2013; Dughmi and Peres 2012). See §4 for a discussion of our problem with risk-averse players.

2 One Facility: The Optimal Tradeoff

In this section we consider the one-facility game. Let us first briefly discuss the social cost objective. As observed by Procaccia and Tennenholtz (2013), selecting the median¹ is an optimal GSP (group strategyproof) mechanism for this objective. (The proof of optimality and group strategyproofness is left as a very easy exercise for the reader.) As the median is a deterministic mechanism, the variance of its social cost is zero. It follows that the approximation-variance tradeoff is a nonissue in one-facility games with the social cost objective. We therefore focus in this section on the maximum cost objective.

2.1 Upper Bound

Our starting point is the optimal SP mechanism for the maximum cost, without variance constraints: the LEFT-RIGHT-MIDDLE (LRM) Mechanism of Procaccia and Tennenholtz (2013). This simple mechanism selects $lt(\mathbf{x})$ with probability $1/4$, $rt(\mathbf{x})$ with probability $1/4$, and the optimal solution $mid(\mathbf{x})$ — whose maximum cost is $opt(\mathbf{x}) = diam(\mathbf{x})/2$ — with probability $1/2$ (see Figure 1). The approximation ratio of the mechanism is clearly $3/2$: with probability $1/2$ it selects one of the extreme locations, which have maximum cost $diam(\mathbf{x}) = 2opt(\mathbf{x})$; and with probability $1/2$ it selects the optimal solution. Why is this mechanism SP? In a nutshell, consider a player $i \in N$; she can only affect the outcome by changing the position of $lt(\mathbf{x})$ or $rt(\mathbf{x})$. Assume without loss of generality that i reports a location x'_i to the left of $lt(\mathbf{x})$, such that $lt(\mathbf{x}) - x'_i = \delta > 0$. Then the leftmost location moves away from x_i by exactly δ , and that location is selected with probability $1/4$. On the other hand, the midpoint might move towards x_i , but it moves half as fast, that is, i might gain at most $\delta/2$ with probability $1/2$ — and the two terms cancel out. This argument is easily extended to show that LRM is GSP (in fact, the proof of Theorem 2.3 rigorously establishes a more general claim). Furthermore, even if we just impose strategyproofness, no mechanism can give an approximation ratio better than $3/2$ for the maximum cost (Procaccia and Tennenholtz 2013).

A first attempt: The $CONVEX_p$ Mechanism. On a location vector $\mathbf{x} \in \mathbb{R}^n$, the LRM Mechanism has variance $opt(\mathbf{x})^2/4$, or, equivalently, standard deviation $opt(\mathbf{x})/2$. Given a smaller variance “budget”, how would the approximation ratio change? The most natural approach to reducing the variance of the LRM Mechanism is to randomize between it and the optimal deterministic (G)SP mechanism, which gives a 2-approximation for the maximum cost by simply selecting $lt(\mathbf{x})$. Specifically, we select $lt(\mathbf{x})$ with probability $1 - p \geq 0$, and with probability p follow LRM (see Figure 1). This is a special case of a general mechanism, which randomizes between the optimal deterministic mechanism and the optimal randomized mechanism. We call this mechanism $CONVEX_p$, and analyze it in some generality in

the paper’s full version.² For the specific problem in question, this mechanism yields the following result.

Proposition 2.1. *Let X be the maximum cost of $CONVEX_p$ on input \mathbf{x} . Then,*

$$\mathbb{E}[X] + \text{std}(X) = \left(2 - \frac{p}{2} + \sqrt{\left(1 - \frac{p}{2}\right) \cdot \frac{p}{2}}\right) \cdot \text{opt}(\mathbf{x}).$$

In particular, if $p \neq 0, 1$ then $\mathbb{E}[X] + \text{std}(X) > 2 \cdot \text{opt}(\mathbf{x})$. As we shall see in Section 2.1, this approximation to standard deviation tradeoff is suboptimal.

It is worth noting that another natural approach — modifying LRM by increasing the probability of each of the two extreme points to $q \in [1/4, 1/2]$, and decreasing the probability of the midpoint to $1 - 2q$ — turns out to be equivalent to $CONVEX_p$ for $p = 4q - 1$. Indeed, the former mechanism is just a symmetrized version of $CONVEX_p$.

The optimal mechanism. In retrospect, the extension of LRM that does achieve the optimal approximation-variance tradeoff is no less intuitive than the ones discussed earlier. The idea is to think of $mid(\mathbf{x})$, which is selected by LRM with probability $1/2$, as two points, each selected with probability $1/4$. These two points can then be continuously moved at equal pace towards the extremes (see Figure 1). In what follows, this mechanism is defined formally.

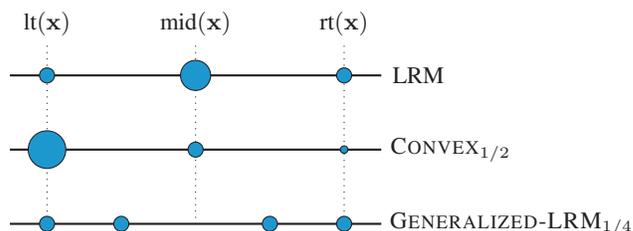


Figure 1: Illustration of the three randomized mechanisms. The balls’ radii correspond to their points’ probabilities of being selected.

Definition 2.2. *The $GENERALIZED-LRM_\alpha$ Mechanism is parameterized by $\alpha \in [0, 1/2]$; on location vector \mathbf{x} , $GENERALIZED-LRM_\alpha$ outputs a point y chosen uniformly at random from the set $\{lt(\mathbf{x}), mid(\mathbf{x}) - \alpha \cdot diam(\mathbf{x}), mid(\mathbf{x}) + \alpha \cdot diam(\mathbf{x}), rt(\mathbf{x})\}$.*

The next theorem presents the properties satisfied by $GENERALIZED-LRM_\alpha$.

Theorem 2.3. *For all $\alpha \in [0, 1/2]$, $GENERALIZED-LRM_\alpha$ is a GSP mechanism for one-facility games. Moreover, if X is the random variable corresponding to the maximum cost of the mechanism on input \mathbf{x} , then $\mathbb{E}[X] = (3/2 + \alpha) \cdot \text{opt}(\mathbf{x})$ and $\text{std}(X) = (1/2 - \alpha) \cdot \text{opt}(\mathbf{x})$.*

Proof. Table 1 summarizes the maximum cost for each possible y that $GENERALIZED-LRM_\alpha$ outputs (recall that

¹Take the left median when the number of players is even.

²Available at <http://procaccia.info/research>.

Table 1: Maximum cost of GENERALIZED-LRM $_{\alpha}$ for its different choices of y .

y	$\arg \max_{x_i \in \mathbf{x}} y - x_i $	$X = \max_{x_i \in \mathbf{x}} y - x_i $
$\text{mid}(\mathbf{x}) - \alpha \cdot \text{diam}(\mathbf{x})$	$\text{rt}(\mathbf{x})$	$(1 + 2\alpha) \cdot \text{opt}(\mathbf{x})$
$\text{mid}(\mathbf{x}) + \alpha \cdot \text{diam}(\mathbf{x})$	$\text{lt}(\mathbf{x})$	$(1 + 2\alpha) \cdot \text{opt}(\mathbf{x})$
$\text{lt}(\mathbf{x})$	$\text{rt}(\mathbf{x})$	$2 \cdot \text{opt}(\mathbf{x})$
$\text{rt}(\mathbf{x})$	$\text{lt}(\mathbf{x})$	$2 \cdot \text{opt}(\mathbf{x})$

$\text{opt}(\mathbf{x}) = \text{diam}(\mathbf{x})/2$). Inspecting this table we find that indeed the expectation satisfies

$$\mathbb{E}[X] = \left(\frac{3}{2} + \alpha\right) \cdot \text{opt}(\mathbf{x}).$$

Given $\mathbb{E}[X]$ and our table of X given y , we see that the variance is

$$\text{Var}(X) = \left(\frac{1}{2} - \alpha\right)^2 \cdot \text{opt}(\mathbf{x})^2,$$

and so

$$\text{std}(X) = \left(\frac{1}{2} - \alpha\right) \cdot \text{opt}(\mathbf{x}),$$

as claimed.

To establish that GENERALIZED-LRM $_{\alpha}$ is GSP, suppose a group of players $S \subseteq [n]$ misreport their locations, resulting in a different location vector \mathbf{x}' . Denote $\Delta_L \triangleq \text{lt}(\mathbf{x}) - \text{lt}(\mathbf{x}')$ and $\Delta_R \triangleq \text{rt}(\mathbf{x}') - \text{rt}(\mathbf{x})$. Note that Δ_L and Δ_R may be positive for any misreporting group $S \subseteq [n]$, but for Δ_L (or Δ_R) to be negative requires the leftmost (respectively, the rightmost) player in $[n]$ to be in S . By considering the cases given by the signs of Δ_L and Δ_R , we show that for any values of Δ_L, Δ_R , there is some misreporting player $i \in S$ whose cost does not decrease.

Case 1: $\Delta_L, \Delta_R \geq 0$. Let $z_L \triangleq \text{mid}(\mathbf{x}) - \alpha \cdot \text{diam}(\mathbf{x})$ and $z_R \triangleq \text{mid}(\mathbf{x}) + \alpha \cdot \text{diam}(\mathbf{x})$ and let z'_L, z'_R be defined analogously for the misreported location vector \mathbf{x}' . Then, for any player location x_i (clearly $x_i \in [\text{lt}(\mathbf{x}), \text{rt}(\mathbf{x})]$) we have

$$\begin{aligned} \text{cost}(f(\mathbf{x}), x_i) &= \frac{1}{4} \cdot ((x_i - \text{lt}(\mathbf{x})) + (\text{rt}(\mathbf{x}) - x_i) \\ &\quad + |z_L - x_i| + |z_R - x_i|), \end{aligned}$$

and

$$\begin{aligned} \text{cost}(f(\mathbf{x}'), x_i) &= \frac{1}{4} \cdot ((x_i - \text{lt}(\mathbf{x}) + \Delta_L) \\ &\quad + (\text{rt}(\mathbf{x}) - x_i + \Delta_R) \\ &\quad + |z'_L - x_i| + |z'_R - x_i|). \end{aligned}$$

But by the triangle inequality, we find that

$$\begin{aligned} |z'_L - x_i| &\geq |z_L - x_i| - \left| \frac{\Delta_R - \Delta_L}{2} - \alpha(\Delta_L + \Delta_R) \right|, \\ |z'_R - x_i| &\geq |z_R - x_i| - \left| \frac{\Delta_R - \Delta_L}{2} + \alpha(\Delta_L + \Delta_R) \right|. \end{aligned}$$

For

$$\alpha \in \left\{ 0, \frac{|\Delta_R - \Delta_L|}{2(\Delta_L + \Delta_R)}, \frac{1}{2} \right\},$$

it is easily verified that the implied lower bound on $|z'_L - x_i| + |z'_R - x_i|$ is at least $|z_L - x_i| + |z_R - x_i| - (\Delta_L + \Delta_R)$. Furthermore, as this lower bound is linear in α in the two ranges defined by these values, the same holds for all $\alpha \in [0, \frac{1}{2}]$. Putting the above together we get

$$\begin{aligned} \text{cost}(f(\mathbf{x}'), x_i) &\geq \text{cost}(f(\mathbf{x}), x_i) + \frac{1}{4} \cdot (\Delta_L + \Delta_R - (\Delta_L + \Delta_R)) \\ &\geq \text{cost}(f(\mathbf{x}), x_i). \end{aligned}$$

Case 2(a): $\Delta_L < 0$ and $\Delta_R \geq 0$. As observed above, for Δ_L to be negative the leftmost player must be in the deviating set S , but this player cannot gain from this change, and in fact only stands to lose from such a change, as all four points in the support of the mechanism's output move further away from the leftmost player's location.

Case 2(b): $\Delta_L \geq 0$ and $\Delta_R < 0$. This is symmetric to case 2(a) above.

Case 3: $\Delta_L, \Delta_R < 0$. In this case the mechanism outputs a location $y \in [\text{lt}(\mathbf{x}'), \text{rt}(\mathbf{x}')] \subseteq [\text{lt}(\mathbf{x}), \text{rt}(\mathbf{x})]$ with probability one, and by the triangle inequality $|\text{rt}(\mathbf{x}) - y| + |y - \text{lt}(\mathbf{x})| = \text{diam}(\mathbf{x})$. Thus, by linearity of expectation, $\text{cost}(f(\mathbf{x}'), \text{lt}(\mathbf{x})) + \text{cost}(f(\mathbf{x}'), \text{rt}(\mathbf{x})) = \text{diam}(\mathbf{x})$. By the same argument $\text{cost}(f(\mathbf{x}), \text{lt}(\mathbf{x})) + \text{cost}(f(\mathbf{x}), \text{rt}(\mathbf{x})) = \text{diam}(\mathbf{x})$. Consequently, either

$$\text{cost}(f(\mathbf{x}'), \text{lt}(\mathbf{x})) \geq \text{cost}(f(\mathbf{x}), \text{lt}(\mathbf{x}))$$

or

$$\text{cost}(f(\mathbf{x}'), \text{rt}(\mathbf{x})) \geq \text{cost}(f(\mathbf{x}), \text{rt}(\mathbf{x})).$$

But for Δ_L and Δ_R to both be negative, both the leftmost and rightmost players must be in the deviating set S , and so some player in S does not gain from S misreporting their locations. \square

2.2 Matching Lower Bound

We are now ready to present our main technical result for the single-facility location problem: a lower bound for the expectation-variance tradeoff matching the upper bound of Theorem 2.3.

Theorem 2.4. *For all $\alpha \in [0, 1/2]$, no SP mechanism for one-facility location games which is $(3/2 + \alpha)$ -approximate for maximum cost minimization has standard deviation of maximum cost less than $(1/2 - \alpha) \cdot \text{opt}(\mathbf{x})$ on every location vector \mathbf{x} .*

In our proof we fix some SP mechanism f . We will consider inputs of the form $\mathbf{x} = (l, r)$, where $l \leq r$, that is, two-player inputs; this is without loss of generality as the two extreme player locations always define the maximum

cost.³ Throughout the remainder of this section, we denote by $Y(\mathbf{x}) \sim f(\mathbf{x})$ the random variable corresponding to the location of the facility output by the mechanism f on input \mathbf{x} . We write $Y = Y(\mathbf{x})$, whenever the input \mathbf{x} is clear from context. The following two definitions will prove useful in our proof of Theorem 2.4.

Definition 2.5. Given an instance $\mathbf{x} = (l, r)$ and a “gap” t , the normalized leakage of (l, r) with relaxation parameter t is

$$\Lambda(l, r, t) \triangleq \mathbb{E} \left[\left| Y - \frac{l+r}{2} \right| \mid Y \notin (l+t, r-t) \right] \cdot \Pr[Y \notin (l+t, r-t)] \cdot \left(\frac{r-l}{2} \right)^{-1}.$$

Intuitively, $\Lambda(l, r, t)$ is the contribution of probabilities outside $(l+t, r-t)$ to the expected distance from the facility to $\text{mid}(\mathbf{x}) = \frac{l+r}{2}$, normalized by $\text{opt}(\mathbf{x}) = \frac{r-l}{2}$.

Definition 2.6. The left- and right-normalized distance of an instance (l, r) are defined to be

$$d_L(l, r) \triangleq \mathbb{E}[|Y - l|] \cdot \left(\frac{r-l}{2} \right)^{-1},$$

$$d_R(l, r) \triangleq \mathbb{E}[|Y - r|] \cdot \left(\frac{r-l}{2} \right)^{-1}.$$

By the triangle inequality, f satisfies $d_L(l, r) + d_R(l, r) \geq 2$. Moreover, as we may safely assume that f is at worst 2-approximate, we also have $d_L(l, r), d_R(l, r) \leq 2$, and so $d_L(l, r) + d_R(l, r) \leq 4$.

The next result is the core lemma underlying the proof of Theorem 2.4; its rather intricate proof is relegated to the paper’s full version.

Lemma 2.7. For all $\delta > 0$ and $t \in (0, 1/2)$ there exists some input $\mathbf{x} = (l, r)$, such that

$$\Lambda(l, r, t(r-l)) \geq \frac{1}{2} - \delta.$$

We proceed to inspect the variance of bounded SP approximate single-facility mechanisms for maximum cost minimization. For the remainder of the section we assume f is an SP mechanism with expected approximation ratio at most $\frac{3}{2} + \alpha$ for all inputs (with $\alpha < \frac{1}{2}$, as Theorem 2.4 is trivial for $\alpha \geq \frac{1}{2}$.)

By Lemma 2.7, for any (δ, t) , there exists an instance $\mathbf{x} = \mathbf{x}_{\delta, t}$ satisfying $\Lambda(\mathbf{x}, t) \geq \frac{1}{2} - \delta$. Without loss of generality we shift and scale \mathbf{x} to be $(-1, 1)$. Let $Y(\delta, t) \sim f(\mathbf{x}_{\delta, t})$ denote the output of the mechanism on the instance $\mathbf{x}_{\delta, t}$. We omit the parameters δ and t when the context is clear. Let $Z = |Y|$. The following lemma, due to Procaccia and Tennenholtz (2013), relates Z to X , the maximum cost of f on \mathbf{x} .

Lemma 2.8 (Procaccia and Tennenholtz 2013). Let X be the maximum cost of f on input $(-1, 1)$. Then $X = Z + 1$.

³The extension to more than two players is almost immediate, as we can identify more than one player with either extreme location, using (Lu et al. 2010, Lemma 2.1).

Consequently, the maximum cost X has variance $\text{Var}(X) = \text{Var}(Z)$ and so we turn our attention to lower bounding the variance of Z . Moreover, as mechanism f is $(\frac{3}{2} + \alpha)$ -approximate and clearly $\text{opt}(-1, 1) = 1$, Lemma 2.8 implies that $\mathbb{E}[Z] = \frac{1}{2} + \alpha'$ for some $\alpha' \leq \alpha$. By our choice of $\mathbf{x} = (-1, 1)$ satisfying $\Lambda(-1, 1, t) \geq \frac{1}{2} - \delta$, we have $\mathbb{E}[Z \mid Z \geq 1-t] \cdot \Pr[Z \geq 1-t] \geq \frac{1}{2} - \delta$. In order to lower bound $\text{Var}(Z)$ we first consider a simpler distribution, defined below.

Definition 2.9. The concentrated version $Z_c(\delta, t) \triangleq \{(x_c, p_c), (y_c, 1-p_c)\}$ of $Z(\delta, t)$ is a two-point distribution, where

$$y_c = \mathbb{E}[Z \mid Z \in [0, 1-t]],$$

$$x_c = \mathbb{E}[Z \mid Z \in [1-t, \infty)],$$

$$p_c = \Pr[Z \in [1-t, \infty)].$$

In words, Z_c is obtained from Z by concentrating probabilities in the intervals $[1-t, \infty)$ and $[0, 1-t)$ respectively to points x_c and y_c . Note that concentrating probabilities in both intervals to points yields the same expectation as Z and can only decrease the variance. That is, $\mathbb{E}[Z_c] = \mathbb{E}[Z] = \frac{1}{2} + \alpha'$ and $\text{Var}(Z_c) \leq \text{Var}(Z)$. Moreover, the contribution to $\mathbb{E}[Z]$ of Z conditioned on $Z \notin [0, 1-t)$ and the equivalent contribution to $\mathbb{E}[Z_c]$ are the same. That is,

$$p_c x_c = \Lambda(-1, 1, t) \geq \frac{1}{2} - \delta.$$

Revisiting the variance of Z_c , it is easy to see that

$$\begin{aligned} \text{Var}(Z_c) &= \mathbb{E}[Z_c^2] - \mathbb{E}[Z_c]^2 \\ &= p_c x_c^2 + \frac{(\frac{1}{2} + \alpha' - p_c x_c)^2}{1 - p_c} - \left(\frac{1}{2} + \alpha' \right)^2. \end{aligned}$$

Extracting the form of $\text{Var}(Z_c)$, we obtain the following definition.

Definition 2.10. The formal variance $v(p, x, \varepsilon)$ is the expression

$$v(p, x, \varepsilon) \triangleq p x^2 + \frac{(\frac{1}{2} + \varepsilon - p x)^2}{1 - p} - \left(\frac{1}{2} + \varepsilon \right)^2,$$

and the simplified formal variance is $v(p, x) \triangleq v(p, x, \alpha)$.

We aim to bound $v(p, x, \varepsilon)$ and $v(p, x)$ with some constraints on (p, x, ε) , instead of bounding $\text{Var}(Z_c)$ or $\text{Var}(Z)$ directly.

Definition 2.11. The feasible domain $\Omega(\delta, t)$ is defined to be

$$\Omega(\delta, t) \triangleq \left\{ (p, x) \mid p \in [0, 1], x \in [1-t, \infty), \frac{1}{2} - \delta \leq p x \right\}$$

and the relaxed variance bound $V(\delta, t)$ is defined to be

$$V(\delta, t) \triangleq \inf \{ v(p, x) \mid (p, x) \in \Omega(\delta, t) \}.$$

In words, $\Omega(\delta, t)$ is a domain of simplified formal variance $v(p, x)$ containing all possible concentrated versions of $Z(\delta, t)$, and $V(\delta, t)$ is the tightest lower bound on the simplified formal variance $v(p, x)$ in this domain.

The next lemma establishes that the relaxed variance bound serves as a lower bound for $\text{Var}(Z(\delta, t))$; its first inequality was observed earlier, and the proof of the second inequality appears in the paper’s full version.

Lemma 2.12. *For any δ and $t \leq \frac{1}{2} - \alpha$,*

$$\text{Var}(Z(\delta, t)) \geq \text{Var}(Z_c(\delta, t)) \geq V(\delta, t).$$

By Lemma 2.12, it suffices to derive a lower bound on $V(\delta, t)$. The final lemma helps us do that, by giving a formula for the relaxed variance bound; its proof is relegated to the paper’s full version.

Lemma 2.13. *For $t \leq \frac{1}{2} - \alpha$, the relaxed variance bound $V(\delta, t)$ satisfies*

$$V(\delta, t) = v\left(\frac{\frac{1}{2} - \delta}{1 - t}, 1 - t\right).$$

With Lemma 2.13 in hand, we are finally ready to prove this section’s main result.

Proof of Theorem 2.4. Consider a sequence of (δ, t) values $\{(\frac{1}{i}, \frac{1}{i}) \mid i \in \mathbb{N}\}$. By Lemmas 2.12 and 2.13, for i large enough, i.e., $\frac{1}{i} \leq \frac{1}{2} - \alpha$ (recall that $\alpha < \frac{1}{2}$, so such an i exists), we have

$$\text{Var}\left(Z\left(\frac{1}{i}, \frac{1}{i}\right)\right) \geq V\left(\frac{1}{i}, \frac{1}{i}\right) = v\left(\frac{\frac{1}{2} - \frac{1}{i}}{1 - \frac{1}{i}}, 1 - \frac{1}{i}\right).$$

Note that $v\left(\frac{\frac{1}{2} - \tau}{1 - \tau}, 1 - \tau\right)$, a function of τ , is continuous at 0. Therefore

$$\begin{aligned} \sup_{\mathbf{x}} \text{Var}(Z(\mathbf{x})) &\geq \sup_{\frac{1}{i} \leq \frac{1}{2} - \alpha} \text{Var}\left(Z\left(\frac{1}{i}, \frac{1}{i}\right)\right) \\ &\geq \lim_{i \rightarrow \infty} v\left(\frac{\frac{1}{2} - \frac{1}{i}}{1 - \frac{1}{i}}, 1 - \frac{1}{i}\right) \\ &= v\left(\frac{1}{2}, 1\right) \\ &= \left(\frac{1}{2} - \alpha\right)^2, \end{aligned}$$

completing the proof. \square

3 The Curious Case of Multiple Facilities

Having fully characterized the optimal approximation-variance tradeoff for the case of a single facility in Section 2, we turn our attention to multiple facilities. Our first observation is that now the tables are turned: the maximum cost objective is relatively straightforward (given previous work), whereas the social cost objective turns out to be quite convoluted.

In more detail, the best known SP mechanism for the maximum cost objective, and any number of facilities $k \geq 2$, is the EQUAL COST (EC) Mechanism of Fotakis and Tzamos (2013b). The mechanism first covers the player locations with k disjoint intervals $[\alpha_i, \alpha_i + \ell]$, in a way that minimizes the interval length ℓ . Then, with probability $1/2$,

the mechanism places a facility at each α_i if i is odd, and at $\alpha_i + \ell$ if i is even; and, with probability $1/2$, the mechanism places a facility at each α_i if i is even, and at $\alpha_i + \ell$ if i is odd.

It is easy to see that the EC Mechanism is 2-approximate. Moreover (if not as obvious), it is GSP. The crucial observation in our context is that the maximum cost under the EC Mechanism is always exactly ℓ , that is, its maximum cost has zero variance — even though it relies strongly on randomization!

We conclude that, in order to establish any kind of approximation-variance tradeoff for the maximum cost objective, we would need to improve the best known SP approximation mechanism *without variance constraints*, which is not our focus. In the remainder of this section, therefore, we study the social cost objective. Moreover, we restrict ourselves to the case of two facilities; the reason is twofold. First, very little is known about SP mechanisms for social cost minimization with $k \geq 3$ facilities — not for lack of trying. Second, and more importantly, we establish an impossibility result, that holds *even* for the case of two facilities.

The best known SP mechanism for social cost minimization in two-facility games is due to Lu et al. (2010). It selects the first facility from the player locations uniformly at random. Then, it selects the second facility also from the player locations with each location selected to be the second facility with probability proportional to its distance from the first selected facility. Lu et al. show that this mechanism is an SP 4-approximate mechanism. The best *deterministic* approximation is given by the GSP mechanism which simply selects $\text{lt}(\mathbf{x})$ and $\text{rt}(\mathbf{x})$ — its approximation ratio is $\Theta(n)$.

It is natural to think that it should at least be possible to obtain some (possibly suboptimal) approximation-variance tradeoff by randomizing between the two foregoing mechanisms, via the CONVEX_p Mechanism. Strangely enough, the following theorem — our second major technical result — essentially rules this out.

Theorem 3.1. *Let $\{f_\theta\}_{\theta \in [0,1]}$ be a family of SP mechanisms for two-facility games that satisfy the following technical assumptions:*

1. *For any $\theta \in [0, 1]$ and location vector \mathbf{x} , $f_\theta(\mathbf{x})$ places facilities only on locations in \mathbf{x} .*
2. *For any $\theta \in [0, 1]$, if the location vector \mathbf{x} contains at least two different locations, then $f_\theta(\mathbf{x})$ always selects two different locations.*

Define the random variable $C(f_\theta, \mathbf{x})$ to be the social cost of mechanism f_θ on location vector \mathbf{x} . Then the following conditions are mutually exclusive:

3. *f_θ is constant-approximate; i.e., there is a constant $\alpha \geq 1$ such that $\mathbb{E}[C(f_\theta, \mathbf{x})] \leq \alpha \cdot \text{opt}(\mathbf{x})$.*
4. *For any location vector $\mathbf{x} \in \mathbb{R}^n$, $\text{Var}(C(f_\theta, \mathbf{x}))$ decreases monotonically with θ , down to $\text{Var}(C(f_1, \mathbf{x})) = 0$.*
5. *For any location vector $\mathbf{x} \in \mathbb{R}^n$, $\mathbb{E}[C(f_\theta, \mathbf{x})]$ is continuous in θ .*

We think of Conditions 3–5 as the basic requirements that any “reasonable” tradeoff must satisfy. We also con-

sider the first two assumptions as rather mild. In particular, they are both satisfied by every “useful” SP mechanism for minimizing the social cost in two-facility games,⁴ including the best known SP approximation mechanism of Lu et al. (2010), all the mechanisms characterized by Miyagawa (2001),⁵ and the winner-imposing mechanism of Fotakis and Tzamos (2010).

Let us now revisit CONVEX_p in this setting; why is it not a counterexample to the theorem? To be clear, we are thinking of f_0 as the 4-approximation mechanism of Lu et al. (2010), and of f_1 as the rule that deterministically selects $\text{lt}(\mathbf{x})$ and $\text{rt}(\mathbf{x})$ (and has a bounded, though not constant, approximation ratio). It is easy to see that this mechanism satisfies Conditions 1, 2, 3, and 5. Therefore, the theorem implies that CONVEX_p (surprisingly) violates Condition 4: the variance does not decrease monotonically with θ . This stands in contrast to Section 2.1, where the variance of CONVEX_p (as well as $\text{GENERALIZED-LRM}_\alpha$) is monotonic.

The proof of Theorem 3.1 relies on establishing the following, clearly contradictory lemmas.

Lemma 3.2. *If $\{f_\theta\}_{\theta \in [0,1]}$ is a family of SP mechanisms for 2-facility location which satisfies the conditions of Theorem 3.1, then mechanism f_1 has unbounded approximation ratio for the social cost, (even) when restricted to 3-location instances.*

In the proof of the lemma, which can be found in the full version, we first show that the zero-variance mechanism f_1 must, in fact, be deterministic. We can therefore leverage a characterization of deterministic bounded SP mechanisms for 2-facility location (Fotakis and Tzamos 2013a) to establish that f_1 has unbounded approximation ratio, by proving that it cannot belong to this family. We then prove the opposite statement — and Theorem 3.1 follows.

Lemma 3.3. *If $\{f_\theta\}_{\theta \in [0,1]}$ is a family of SP mechanisms for 2-facility location which satisfies the conditions of Theorem 3.1, then mechanism f_1 has bounded approximation ratio for the social cost, (even) when restricted to 3-location instances.*

4 Discussion

We wrap up with a brief discussion of a few salient points.

Possible criticism: Why is the designer risk averse and the players risk neutral? One may wonder why we are studying approximation-variance tradeoffs for the designer, yet the players care only about expected cost. But the two issues are orthogonal. For example, the papers we discuss in Section 1.2 consider sellers and buyers, and typically assume that one side is risk averse while the other is risk neutral. Our model is even more asymmetric as we have two completely different types of objective functions (individual distance from the facility versus an aggregate cost function). Furthermore, to be able to distill the approximation-variance

⁴Unlike the maximum cost objective, for which “useful” mechanisms such as LRM and $\text{GENERALIZED-LRM}_\alpha$ are known to use the freedom to choose facilities outside the player locations.

⁵Miyagawa (2001) assumes Pareto efficiency, which implies our Assumption 2

tradeoff in facility location games, we study the simplest version of the problem, which includes risk-neutral players, in addition to several other strong assumptions made by Procaccia and Tennenholtz (2013), e.g., the cost of a player is exactly her distance to the nearest facility, and players and facilities are located on the real line.

That said, it is worth discussing whether our results can be extended to the case of risk-averse players. If we modeled the players’ risk aversion by changing their utility functions, we would change the set of strategyproof mechanisms. Nevertheless, it might be the case that the optimal approximation-variance tradeoff — for the social cost or maximum cost objective — is *independent* of the players’ individual utility functions. It is somewhat encouraging that the EQUAL COST Mechanism (see §3) of Fotakis and Tzamos (2013b) gives the same approximation guarantees (the best known for the maximum cost) for players with any concave cost function. But risk aversion corresponds to a *convex* cost function (or a concave utility function), for which Fotakis and Tzamos establish negative results.

Possible criticism: Is $\text{GENERALIZED-LRM}_\alpha$ actually a good mechanism? A curious — perhaps even troubling — property of the $\text{GENERALIZED-LRM}_\alpha$ mechanism is that for $\alpha < \alpha'$, $\text{GENERALIZED-LRM}_\alpha$ has higher variance than $\text{GENERALIZED-LRM}_{\alpha'}$, yet the outcome of the former mechanism stochastically dominates that of the latter, in the sense that for every t , the probability that the former mechanism has maximum cost at most t is at least as high as that probability under the latter mechanism. However, this is not an inherent property of our model: There are certainly examples of mechanisms such that one has higher variance than the other, yet neither one stochastically dominates the other. We therefore view Theorem 2.3, and the $\text{GENERALIZED-LRM}_\alpha$ mechanism itself, mainly as a tight upper bound on the optimal approximation-variance tradeoff, rather than as a mechanism that a risk-averse designer would necessarily want to employ.

A broader agenda. As briefly mentioned in §1, we believe that our paper potentially introduces a new research agenda. Just to give one example, the problem of *impartial selection* (Alon et al. 2011; Fischer and Klimm 2014; Holzman and Moulin 2013) exhibits an *easy* separation between the approximation ratio achieved by deterministic and randomized SP mechanisms (much like facility location); what is the optimal approximation-variance tradeoff? Even more exciting are general results that apply to a range of problems in mechanism design. And, while our work mainly applies to facility location, it does tease out general insights and questions: Can we build on the ideas behind the CONVEX_p mechanism to obtain “good” (albeit suboptimal, see §2.1), general approximation-variance tradeoffs? Is a “linear” upper bound of the form $c \cdot \text{opt}$ on the sum of expectation and standard deviation (Theorem 2.3) something that we should expect to see more broadly? Can we characterize problems that do not admit approximation-variance tradeoffs satisfying the conditions of Theorem 3.1? These challenges can drive the development of a theory of expectation-variance analysis in computational mechanism design.

Acknowledgments

This work was partially supported by NSF grants CCF-1527110, CCF-1618280, IIS-1350598, IIS-1714140, IIS-1527434, CCF-1525932, and CCF-1733556; by ONR grants N00014-16-1-3075 and N00014-17-1-2428; and by a Sloan Research Fellowship.

References

- Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2010. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research* 35(3):513–526.
- Alon, N.; Fischer, F.; Procaccia, A. D.; and Tennenholtz, M. 2011. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 101–110.
- Ashlagi, I.; Fischer, F.; Kash, I.; and Procaccia, A. D. 2015. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91:284–296.
- Bhalgat, A.; Chakraborty, T.; and Khanna, S. 2012. Mechanism design for a risk averse seller. In *Proceedings of the 8th Conference on Web and Internet Economics (WINE)*, 198–211.
- Cai, Q.; Filos-Ratsikas, A.; and Tang, P. 2016. Facility location with minimax envy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 137–143.
- Cheng, Y.; Yu, W.; and Zhang, G. 2013. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science* 497:154–163.
- Dughmi, S., and Peres, Y. 2012. Mechanisms for risk averse agents, without loss. arXiv:1206.2957.
- Esfandiari, H., and Kortsarz, G. 2016. A bounded-risk mechanism for the kidney exchange game. In *Proceedings of the 12th Latin American Theoretical Informatics Symposium (LATIN)*, 416–428.
- Eso, P., and Futó, G. 1999. Auction design with a risk averse seller. *Economics Letters* 65:71–74.
- Feldman, M.; Fiat, A.; and Golomb, I. 2016. On voting and facility location. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 269–286.
- Filos-Ratsikas, A.; Li, M.; Zhang, J.; and Zhang, Q. 2015. Facility location with double-peaked preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 893–899.
- Fischer, F., and Klimm, M. 2014. Optimal impartial selection. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 803–820.
- Fotakis, D., and Tzamos, C. 2010. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proceedings of the 6th Conference on Web and Internet Economics (WINE)*, 234–245.
- Fotakis, D., and Tzamos, C. 2013a. On the power of deterministic mechanisms for facility location games. In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP)*, 449–460.
- Fotakis, D., and Tzamos, C. 2013b. Strategyproof facility location for concave cost functions. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, 435–452.
- Fu, H.; Hartline, J. D.; and Hoy, D. 2013. Prior-independent auctions for risk-averse agents. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, 471–488.
- Golomb, I., and Tzamos, C. 2017. Truthful facility location with additive errors. *arXiv preprint arXiv:1701.00529*.
- Holzman, R., and Moulin, H. 2013. Impartial nominations for a prize. *Econometrica* 81(1):173–196.
- Lu, P.; Sun, X.; Wang, Y.; and Zhu, Z. A. 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM Conference on Economics and Computation (EC)*, 315–324.
- Markowitz, H. M. 1952. Portfolio selection. *The Journal of Finance* 7(1):77–91.
- Maskin, E. S., and Riley, J. G. 1984. Optimal auctions with risk averse buyers. *Econometrica* 52(6):1473–1518.
- Miyagawa, E. 2001. Locating libraries on a street. *Social Choice and Welfare* 18(3):527–541.
- Nissim, K.; Smorodinsky, R.; and Tennenholtz, M. 2012. Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, 203–213.
- Procaccia, A. D., and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation* 1(4): article 18.
- Serafino, P., and Ventre, C. 2015. Truthful mechanisms without money for non-utilitarian heterogeneous facility location. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 1029–1035.
- Sundararajan, M., and Yan, Q. 2017. Robust mechanisms for risk-averse sellers. *Games and Economic Behavior*. Forthcoming.
- Thang, N. K. 2010. On (group) strategy-proof mechanisms without payment for facility location games. In *Proceedings of the 4th Conference on Web and Internet Economics (WINE)*, 531–538.
- Todo, T.; Iwasaki, A.; and Yokoo, M. 2011. False-name-proof mechanism design without money. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 651–658.
- Wilf, Y., and Feldman, M. 2013. Strategyproof facility location and the least squares objective. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, 873–890.
- Zou, S., and Li, M. 2015. Facility location games with dual preference. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 615–623.