

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, naïve 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, Wang, and Zhou 2009; Lu et al. 2010; Nissim, Smorodinsky, and Tennenholtz 2012; Fotakis and Tzamos 2010; 2013a; 2013b; Thang 2010; 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 Appendix A. For the specific problem in question, this mechanism yields the following result.

Corollary 2.1 (of Theorem A.2). *Let X be the maximum cost of $CONVEX_p$ on input \mathbf{x} . Then,*

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

In particular, if $p \neq 0, 1$ then $\mathbb{E}[X] + \text{std}(X) > 2 \cdot opt(\mathbf{x})$. 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 $\mathbb{E}[X] = (3/2 + \alpha) \cdot opt(\mathbf{x})$ and $\text{std}(X) = (1/2 - \alpha) \cdot opt(\mathbf{x})$.*

Proof. Table 1 summarizes the maximum cost for each possible y that $GENERALIZED-LRM_\alpha$ outputs (recall that $opt(\mathbf{x}) = diam(\mathbf{x})/2$). Inspecting this table we find that indeed the expectation satisfies

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

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

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 $
$\text{mid}(\mathbf{x}) - \alpha \cdot \text{diam}(\mathbf{x})$	$\text{rt}(\mathbf{x})$	$(1 + 2\alpha) \cdot \text{opt}(\mathbf{x})$
$\text{mid}(\mathbf{x}) + \alpha \cdot \text{diam}(\mathbf{x})$	$\text{lt}(\mathbf{x})$	$(1 + 2\alpha) \cdot \text{opt}(\mathbf{x})$
$\text{lt}(\mathbf{x})$	$\text{rt}(\mathbf{x})$	$2 \cdot \text{opt}(\mathbf{x})$
$\text{rt}(\mathbf{x})$	$\text{lt}(\mathbf{x})$	$2 \cdot \text{opt}(\mathbf{x})$

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.² Throughout the remainder of this section, we denote

²The extension to more than two players is almost immediate, as we can identify more than one player with either extreme location, using Lemma C.2.

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

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

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 Appendix B.1.

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 $X = Z + 1$.

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 $\mathbb{E}[Z|Z \geq 1-t] \cdot \Pr[Z \geq 1-t] \geq \frac{1}{2} - \delta$. 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

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

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

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

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

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

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

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

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

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

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 Appendix B.2.

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 Appendix B.3.

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 consider the first two assumptions as rather mild. In particular, they are both satisfied by every “useful” SP mecha-

nism for minimizing the social cost in two-facility games,³ including the best known SP approximation mechanism of Lu et al. (2010), all the mechanisms characterized by Miyagawa (2001),⁴ 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 Appendix C.1, 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 in Appendix C.2 — 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 tradeoff in facility location games, we study the simplest

³Unlike the maximum cost objective, for which “useful” mechanisms such as LRM and $\text{GENERALIZED-LRM}_\alpha$ are known to make use of the freedom to choose facilities outside the player locations.

⁴Miyagawa (2001) assumes Pareto efficiency, which implies our Assumption 2

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 (see Appendix A) 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.

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.
- Lu, P.; Wang, Y.; and Zhou, Y. 2009. Tighter bounds for facility games. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, 137–148.
- 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.

Appendix: Approximation-Variance Tradeoffs in Facility Location Games

A General Mechanism: Convex Combinations of Mechanisms

In this section we analyze a natural and general mechanism to obtain approximate-variance tradeoffs.

Definition A.1. Let \mathcal{M}_a and \mathcal{M}_b be two approximate mechanisms for some (common) optimization problem. Then mechanism $\text{CONVEX}_p(\mathcal{M}_a, \mathcal{M}_b)$ is defined as follows: this mechanism emulates \mathcal{M}_a with probability p and emulates \mathcal{M}_b with probability $1 - p$.

Linearity of expectation ensures that if \mathcal{M}_a and \mathcal{M}_b are both SP, then so is the derived mechanism $\text{CONVEX}_p(\mathcal{M}_a, \mathcal{M}_b)$. Moreover, also by linearity of expectation, \mathcal{M}_p obtains an approximation ratio of $p \cdot \alpha_a + (1 - p) \cdot \alpha_b$; that is, its approximation ratio varies linearly with p . Unfortunately, standard deviation does not degrade linearly, as we shall see.

Specifically, our analysis focuses on *minimization* problems. We show that CONVEX_p yields a super-linear approximation to standard deviation tradeoff. Consequently, for 1-facility games with the maximum cost objective, this mechanism is suboptimal.

Theorem A.2. Let \mathcal{M}_a and \mathcal{M}_b be approximate mechanisms for some minimization problem. Consider an input \mathbf{x} which (up to scaling) has optimal value $\text{opt}(\mathbf{x}) = 1$. Suppose that on this input these mechanisms' respective approximation ratios and variances are α_a, α_b and σ_a^2, σ_b^2 . If X is the random variable corresponding to the cost of $\text{CONVEX}_p(\mathcal{M}_a, \mathcal{M}_b)$ on input \mathbf{x} , then for all $p \in (0, 1)$, if $\alpha_a \neq \alpha_b$ or $\sigma_a \neq \sigma_b$, then

$$\mathbb{E}[X] + \text{std}(X) > p \cdot (\alpha_a + \sigma_a) + (1 - p) \cdot (\alpha_b + \sigma_b).$$

Generally, $\mathbb{E}[X] = p \cdot \alpha_a + (1 - p) \cdot \alpha_b$ and

$$\begin{aligned} \text{std}(X)^2 &= (p \cdot \sigma_a + (1 - p) \cdot \sigma_b)^2 \\ &\quad + p \cdot (1 - p) \cdot ((\alpha_a - \alpha_b)^2 + (\sigma_a^2 - \sigma_b^2)). \end{aligned}$$

Proof. By linearity of expectation we have that indeed $\mathbb{E}[X] = p \cdot \alpha_a + (1 - p) \cdot \alpha_b$. Next, denote by X_a and X_b the cost of mechanisms \mathcal{M}_a and \mathcal{M}_b respectively. By definition, we have that

$$\sigma_a^2 = \text{Var}(X_a) = \mathbb{E}[X_a^2] - \mathbb{E}[X_a]^2,$$

or equivalently $\mathbb{E}[X_a^2] = \alpha_a^2 + \sigma_a^2$. Likewise, $\mathbb{E}[X_b^2] = \alpha_b^2 + \sigma_b^2$. Conditioning on whether or not mechanism $\text{CONVEX}_p(\mathcal{M}_a, \mathcal{M}_b)$ follows \mathcal{M}_a , we find that

$$\begin{aligned} \text{Var}(X) &= \mathbb{E}[X^2] - \mathbb{E}[X]^2 \\ &= p \cdot \mathbb{E}[X_a^2] + (1 - p) \cdot \mathbb{E}[X_b^2] \\ &\quad - (p \cdot \mathbb{E}[X_a] + (1 - p) \cdot \mathbb{E}[X_b])^2 \\ &= p \cdot (\alpha_a^2 + \sigma_a^2) + (1 - p) \cdot (\alpha_b^2 + \sigma_b^2) \\ &\quad - (p \cdot \alpha_a + (1 - p) \cdot \alpha_b)^2 \\ &= (p \cdot \sigma_a + (1 - p) \cdot \sigma_b)^2 \\ &\quad + p \cdot (1 - p) \cdot ((\alpha_a - \alpha_b)^2 + (\sigma_a^2 - \sigma_b^2)). \end{aligned}$$

$$+ p \cdot (1 - p) \cdot ((\alpha_a - \alpha_b)^2 + (\sigma_a^2 - \sigma_b^2)).$$

The term $p \cdot (1 - p) \cdot ((\alpha_a - \alpha_b)^2 + (\sigma_a^2 - \sigma_b^2))$ above is strictly greater than zero, provided $p \neq 0, 1$ and $\alpha_a \neq \alpha_b$ or $\sigma_a \neq \sigma_b$, in which case we have that indeed

$$\mathbb{E}[X] + \text{std}(X) > p \cdot (\alpha_a + \sigma_a) + (1 - p) \cdot (\alpha_b + \sigma_b). \quad \square$$

At this point, we should note a delicate point, namely that $\alpha_a, \alpha_b, \sigma_a^2$ and σ_b^2 in Theorem A.2's statement need not be the worst-case approximation ratios and variances of both mechanisms. In particular, if the "hard inputs" for mechanism \mathcal{M}_a and \mathcal{M}_b do not coincide, the above expression parameterized by the worst-case approximation ratios and variances of the mechanisms serves as an upper bound for the approximation-variance tradeoff achieved by Mechanism $\text{CONVEX}_p(\mathcal{M}_a, \mathcal{M}_b)$. However, for 1-facility location games, the hard instances for the best-known optimal deterministic and randomized mechanisms are one and the same, as the distributions of these mechanisms' approximation ratios are invariant under shifting and scaling. Therefore, for this problem, we may replace $\alpha_a, \alpha_b, \sigma_a^2$ and σ_b^2 with the worst-case approximation ratios and variances. In particular, by Theorem A.2 and our lower bound of Theorem 2.4, we obtain the corollary stated in §2.

Corollary 2.1 (reformulated). For one-facility maximum cost minimization, using an optimal (2-approximate) deterministic mechanism and the optimal ($\frac{3}{2}$ -approximate and $\frac{1}{4}$ -variance) randomized mechanism LRM to play the roles of \mathcal{M}_a and \mathcal{M}_b , in $\text{CONVEX}_p(\mathcal{M}_a, \mathcal{M}_b)$ yields a randomized mechanism whose approximation ratio X satisfies $\mathbb{E}[X] = 2 - \frac{p}{2}$ and $\text{std}(X) = \sqrt{\frac{p}{2} \cdot (1 - \frac{p}{2})}$.

This corollary coupled with our upper bound of Theorem 2.3 implies that the approximation-variance tradeoff achieved by Mechanism CONVEX_p for single-facility maximum cost minimization,

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


is suboptimal for this particular game, as Mechanism GENERALIZED-LRM $_{\alpha}$ has approximation ratio X satisfying $\mathbb{E}[X] + \text{std}(X) \leq 2$. For reference, Figure 2 contains a comparison of the standard deviation to expectation curve obtained by CONVEX_p compared to the optimal mechanism, GENERALIZED-LRM $_{\alpha}$, and the "error term" (their difference) as a function of $\mathbb{E}[X]$. Note that the standard deviation of CONVEX_p decreases monotonically with its expectation, though not linearly.

B Proof of Theorem 2.4: Omitted Lemmas


This section contains proofs of lemmas that were omitted from the body of the paper. The lemmas themselves are stated in §2.2.

B.1 Proof of Lemma 2.7

Assume for the sake of contradiction that the lemma does not hold; then (throughout the proof) we can fix some $\delta > 0$



(a) Relation between $\mathbb{E}[X]$ and $\text{std}(X)$.



(b) Error term.

Figure 2: CONVEX_p contrasted with $\text{GENERALIZED-LRM}_\alpha$.

and $0 < t < \frac{1}{2}$ such that for all (l, r) ,

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

We begin by studying local properties of normalized leakage. The inputs of interest are given in the following definition.

Definition B.1. A gadget G with parameters l, r and offset $x \leq r-l$ is a set of three 2-player instances,

$$G(l, r, x) \triangleq \{(l, r), (l+x, r), (l, r-x)\}.$$

Lemma B.2. For a gadget $G(l, r, x)$ where $x \leq t(r-l)$,

$$\begin{aligned} & d_L(l, r) + d_R(l, r) \\ & \geq (d_L(l+x, r) + d_R(l, r-x)) \cdot \left(1 + \frac{x}{r-l-2x}\right) \\ & \quad - \frac{x}{r-l-2x} \cdot (2-4\delta). \end{aligned}$$

Proof. Let $Y \sim f(l, r)$ be the location output by mechanism f on input (l, r) . By strategyproofness of f , the left player in $(l+x, r)$ will not deviate to (l, r) , nor will the right player in $(l, r-x)$. Thus,

$$\begin{aligned} \frac{r-l-x}{2} \cdot d_L(l+x, r) & \leq \mathbb{E}[|Y-l-x|], \\ \frac{r-l-x}{2} \cdot d_R(l, r-x) & \leq \mathbb{E}[|Y-r+x|]. \end{aligned}$$

Adding the two inequalities we obtain

$$\begin{aligned} & \frac{r-l-x}{2} \cdot (d_L(l+x, r) + d_R(l, r-x)) \\ & \leq \mathbb{E}[|Y-l-x| + |Y-r+x|]. \end{aligned} \quad (2)$$

We focus on the right-hand side of the above expression, $\mathbb{E}[|Y-l-x| + |Y-r+x|]$, conditioned on the events

I and O , corresponding to $Y \in (l+x, r-x)$ and $Y \notin (l+x, r-x)$. That is, whether or not Y is inside or outside the range $(l+x, r-x)$. For the latter case, we rewrite the definition of normalized leakage,

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

By the triangle inequality, this yields

$$\begin{aligned} & \mathbb{E}[|Y-l-x| + |Y-r+x| \mid O] \cdot \Pr[O] \\ & = 2 \cdot \mathbb{E}\left[\left|Y - \frac{r+l}{2}\right| \mid O\right] \cdot \Pr[O] \\ & = (r-l) \cdot \Lambda(l, r, x). \end{aligned} \quad (3)$$

For the former case (i.e., $Y \in (l+x, r-x)$), again by the triangle inequality we have that

$$\mathbb{E}[|Y-l-x| + |Y-r+x| \mid I] = r-l-2x,$$

and similarly $\mathbb{E}[|Y-l| + |Y-r| \mid I] = r-l$. We therefore have

$$\begin{aligned} & \mathbb{E}[|Y-l-x| + |Y-r+x| \mid I] \\ & = \frac{r-l-2x}{r-l} \cdot \mathbb{E}[|Y-l| + |Y-r| \mid I]. \end{aligned} \quad (4)$$

In order to bound the above expectation conditioned on I , we consider the same expectation conditioned on I 's complement, O . Now, for $Y \in [l, r]$ we have $2 \cdot |Y - \frac{l+r}{2}| \leq r-l = |Y-l| + |Y-r|$. On the other hand, for $Y \notin [l, r]$ we have that $2 \cdot |Y - \frac{l+r}{2}| = |Y-l| + |Y-r|$. Therefore we find that

$$2 \cdot \mathbb{E}\left[\left|Y - \frac{l+r}{2}\right| \mid O\right] \leq \mathbb{E}[|Y-l| + |Y-r| \mid O]. \quad (5)$$

Relating the above expressions to normalized distances, we note that by the law of total expectation,

$$\begin{aligned} & (d_L(l, r) + d_R(l, r)) \cdot \frac{r-l}{2} \\ & = \sum_{E \in \{I, O\}} \mathbb{E}[|Y-l| + |Y-r| \mid E] \cdot \Pr[E]. \end{aligned} \quad (6)$$

Therefore, using Equations (5) and (6), and again relying on the definition of $\Lambda(l, r, x)$, we obtain

$$\begin{aligned} & \mathbb{E}[|Y-l-x| + |Y-r+x| \mid I] \cdot \Pr[I] \\ & \leq \left(\frac{r-l}{2}\right) \cdot (d_L(l, r) + d_R(l, r)) \\ & \quad - 2 \cdot \mathbb{E}\left[\left|Y - \frac{r+l}{2}\right| \mid O\right] \cdot \Pr[O] \\ & = \left(\frac{r-l}{2}\right) \cdot (d_L(l, r) + d_R(l, r) - 2 \cdot \Lambda(l, r, x)). \end{aligned} \quad (7)$$

Concluding the above discussion,

$$\begin{aligned} & \mathbb{E}[|Y-l-x| + |Y-r+x|] \\ & = \sum_{E \in \{I, O\}} \mathbb{E}[|Y-l-x| + |Y-r+x| \mid E] \cdot \Pr[E] \end{aligned}$$

$$\begin{aligned}
&\leq (r-l) \cdot \Lambda(l, r, x) \\
&\quad + \frac{r-l-2x}{2} \cdot (d_L(l, r) + d_R(l, r) - 2 \cdot \Lambda(l, r, x)) \\
&= \frac{r-l-2x}{2} \cdot (d_L(l, r) + d_R(l, r)) + 2x \cdot \Lambda(l, r, x) \\
&< \frac{r-l-2x}{2} \cdot (d_L(l, r) + d_R(l, r)) + 2x \cdot \left(\frac{1}{2} - \delta\right),
\end{aligned}$$

where second transition follows from Equations (3), (4), and (7), and the last transition follows from $\Lambda(l, r, x) \leq \Lambda(l, r, t(r-l))$,⁵ and from Equation (1).

Combining Equation (2) with the foregoing upper bound on $\mathbb{E}[|Y-l-x| + |Y-r+x|]$, we obtain

$$\begin{aligned}
&(d_L(l+x, r) + d_R(l, r-x)) \cdot \left(1 + \frac{x}{r-l-2x}\right) \\
&\leq d_L(l, r) + d_R(l, r) + \frac{x}{r-l-2x} \cdot (2-4\delta).
\end{aligned}$$

The lemma follows. \square

Next, we study global properties of normalized leakage. We define an *alignment* of instances to be a set of instances with the same lengths and a certain offset. Formally:

Definition B.3. An alignment is defined by

$$\begin{aligned}
A(l, r, x, n) \triangleq \{ &(l, r), (l+x, r+x), \dots, \\ &(l+(n-1)x, r+(n-1)x)\}.
\end{aligned}$$

We let $(l_j, r_j) = (l+(j-1)x, r+(j-1)x)$ denote the j -th instance in alignment $A(l, r, x, n)$ when the context is clear.

Definition B.4. The average distance of an alignment $A = A(l, r, x, n)$ is defined to be

$$d(A) \triangleq \frac{1}{n} \sum_{j=1}^n (d_L(l_j, r_j) + d_R(l_j, r_j)).$$

As we noted before, for any input $\mathbf{x} = (l, r)$ mechanism f satisfies $2 \leq d_L(l, r) + d_R(l, r) \leq 4$. In particular we have that the average distance for any alignment A satisfies $2 \leq d(A) \leq 4$.

Definition B.5. An alignment hierarchy is a set of alignments with the same “starting points”, the same “ending points”, the same offsets and different lengths of instances. To be precise, a hierarchy with parameters x, n, m is defined to be

$$\begin{aligned}
H(x, n, m) \triangleq \{ &A(0, 1+x, x, n), A(0, 1+2x, x, n-1), \\ &\dots, A(0, 1+mx, x, n-m+1)\}.
\end{aligned}$$

We let $A_i = A(0, 1+ix, x, n-i+1)$ denote the i -th alignment in hierarchy $H(x, n, m)$ when the context is clear.

⁵To see why $\Lambda(l, r, x) \leq \Lambda(l, r, t(r-l))$ follows from $x \leq t(r-l)$, recall that $\Lambda(l, r, x)$ is the contribution of Y outside the range $(l-x, r+x) \supseteq (l+t(r-l), r-t(r-l))$ to $\mathbb{E}[|Y - \frac{l+r}{2}|]$. That is, $\Lambda(l, r, x)$ corresponds to the contribution of a smaller range of Y to this expectation than the range which $\Lambda(l, r, t(r-l))$ corresponds to.

Lemma B.6. For any hierarchy $H(x, n, m)$, for any $x \leq t(1+x)$ and $i \in [2, m-1]$,

$$d(A_{i+1}) \geq d(A_i) + \frac{4x\delta}{1+(i-1)x} - \frac{4}{n-i}.$$

Proof. Let (l_j^i, r_j^i) denote the j -th instance in A_i , i.e., let $(l_j^i, r_j^i) = ((j-1)x, 1+(j+i-1)x)$. Note that for all $j \in [n-i]$, the three inputs $\{(l_j^{i+1}, r_j^{i+1}), (l_{j+1}^i, r_{j+1}^i), (l_j^i, r_j^i)\}$ form a gadget $G(l_j^{i+1}, r_j^{i+1}, x)$ with offset x and width $r_j^{i+1} - l_j^{i+1} = 1 + (i+1)x$, so $x \leq t(1+x) = t(r_j^{i+1} - l_j^{i+1})$. Hence by Lemma B.2 we have that $d_L(l_j^{i+1}, r_j^{i+1}) + d_R(l_j^{i+1}, r_j^{i+1})$ is lower bounded by

$$\begin{aligned}
&(d_L(l_{j+1}^i, r_{j+1}^i) + d_R(l_j^i, r_j^i)) \cdot \left(1 + \frac{x}{1+(i-1)x}\right) \\
&\quad - \frac{x}{1+(i-1)x} \cdot (2-4\delta).
\end{aligned}$$

Summing over j , we find that

$$(n-i) \cdot d(A_{i+1}) = \sum_{j=1}^{n-i} (d_L(l_j^{i+1}, r_j^{i+1}) + d_R(l_j^{i+1}, r_j^{i+1}))$$

is lower bounded by

$$\begin{aligned}
&\sum_{j=1}^{n-i} (d_L(l_{j+1}^i, r_{j+1}^i) + d_R(l_j^i, r_j^i)) \cdot \left(1 + \frac{x}{1+(i-1)x}\right) \\
&\quad - \sum_{j=1}^{n-i} \frac{x}{1+(i-1)x} \cdot (2-4\delta) \\
&= \sum_{j=1}^{n-i} (d_L(l_{j+1}^i, r_{j+1}^i) + d_R(l_j^i, r_j^i)) \cdot \left(1 + \frac{x}{1+(i-1)x}\right) \\
&\quad - \frac{x(2-4\delta)}{1+(i-1)x} \cdot (n-i).
\end{aligned}$$

First, we observe that the distances of the leftmost and rightmost points in A_i , namely $d_L(l_1^i, r_1^i)$ and $d_R(l_{n-i+1}^i, r_{n-i+1}^i)$, are not counted in the above expression. Recalling that for any input (l, r) mechanism f must satisfy $d_L(l, r), d_R(l, r) \leq 2$, we find that the above expression is lower bounded by

$$\begin{aligned}
&\sum_{j=1}^{n-i+1} (d_L(l_j^i, r_j^i) + d_R(l_j^i, r_j^i)) \cdot \left(1 + \frac{x}{1+(i-1)x}\right) \\
&\quad - \frac{x(2-4\delta)}{1+(i-1)x} \cdot (n-i) - 4 \cdot \left(1 + \frac{x}{1+(i-1)x}\right).
\end{aligned}$$

Next, recalling that input (l, r) mechanism f must satisfy $d_L(l, r) + d_R(l, r) \geq 2$, we find that the above expression is in turn lower bounded by

$$\begin{aligned}
&\sum_{j=1}^{n-i+1} (d_L(l_j^i, r_j^i) + d_R(l_j^i, r_j^i)) \\
&\quad + \frac{4x\delta}{1+(i-1)x} \cdot (n-i) - 4 \cdot \left(1 + \frac{x}{1+(i-1)x}\right).
\end{aligned} \tag{8}$$

But, as we have $x \leq t(1+x)$ and $t < 1/2$, this implies that $\frac{x}{1+(i-1)x} < \frac{1}{2}$ for all $i \geq 2$. Therefore (8) is lower bounded by

$$\sum_{j=1}^{n-i+1} (d_L(l_j^i, r_j^i) + d_R(l_j^i, r_j^i)) + \frac{4x\delta(n-i)}{1+(i-1)x} - 6.$$

Finally, dividing through by $n-i$, we find that indeed

$$\begin{aligned} d(A_{i+1}) &\geq \frac{1}{n-i} \cdot \left(\sum_{j=1}^{n-i+1} d_L(l_j^i, r_j^i) + d_R(l_j^i, r_j^i) \right) \\ &\quad + \frac{4x\delta}{1+(i-1)x} - \frac{6}{n-i} \\ &\geq \frac{1}{n-i+1} \cdot \left(\sum_{j=1}^{n-i+1} d_L(l_j^i, r_j^i) + d_R(l_j^i, r_j^i) \right) \\ &\quad + \frac{4x\delta}{1+(i-1)x} - \frac{6}{n-i} \\ &= d(A_i) + \frac{4x\delta}{1+(i-1)x} - \frac{6}{n-i}. \quad \square \end{aligned}$$

Given Lemma B.6, we are now ready to prove our core lemma, Lemma 2.7.

Proof of Lemma 2.7. We note that for any $x > 0$ and $\delta > 0$, the series $\sum_{i=2}^{\infty} \frac{4x\delta}{1+(i-1)x}$ diverges. We may therefore fix some $x > 0$ such that $x \leq t(1+x)$, an m such that

$$\sum_{i=2}^{m-1} \frac{4x\delta}{1+(i-1)x} > 3, \quad (9)$$

and n such that

$$\sum_{i=2}^{m-1} \frac{6}{n-i} < 1, \quad (10)$$

and consider the hierarchy $H(x, n, m)$ with these parameters. By Lemma B.6, which held under the assumption that Lemma 2.7 does not hold for the pair (δ, t) , we have

$$\begin{aligned} d(A_m) - d(A_2) &= \sum_{i=2}^{m-1} (d(A_{i+1}) - d(A_i)) \\ &\geq \sum_{i=2}^{m-1} \frac{4x\delta}{1+(i-1)x} - \sum_{i=2}^{m-1} \frac{4}{n-i} \\ &> 2, \end{aligned}$$

where the last inequality follows from Equations (9) and (10). That is, $d(A_m) > d(A_2) + 2$. But, as observed before, the average distance of any alignment A must satisfy $4 \geq d(A) \geq 2$, and so we find that $4 \geq d(A_m) > d(A_2) + 2 \geq 4$, a contradiction. \square

B.2 Proof of Lemma 2.12

By definition of the formal variance $v(p, x, \varepsilon)$ and the constraints on Z_c , we have

$$\text{Var}(Z_c) \geq \inf\{v(p, x, \varepsilon) \mid (p, x) \in \Omega(\delta, t), \varepsilon \leq \alpha\}.$$

Note that for fixed p and x , the formal variance $v(p, x, \varepsilon)$ is

$$\begin{aligned} v(p, x, \varepsilon) &= px^2 + \frac{\left(\frac{1}{2} - px\right)^2}{1-p} - \frac{1}{4} \\ &\quad + \frac{p}{1-p} (\varepsilon^2 - (2x-1)\varepsilon). \end{aligned}$$

which is quadratic in ε , with an axis of symmetry at $\varepsilon = (x - \frac{1}{2})$. As $t \leq \frac{1}{2} - \alpha$, for all $x \geq 1-t$ and $\varepsilon \leq \alpha$ the following holds: $x - \frac{1}{2} \geq \frac{1}{2} - t \geq \alpha \geq \varepsilon$. By the above we conclude that for any fixed p and $x \geq 1-t$, the function $v(p, x, \varepsilon)$ is monotone decreasing in ε for all $\varepsilon \leq \alpha$, implying that

$$\begin{aligned} \text{Var}(Z_c) &\geq \inf\{v(p, x, \varepsilon) \mid (p, x) \in \Omega(\delta, t), \varepsilon \leq \alpha\} \\ &= \inf\{v(p, x, \alpha) \mid (p, x) \in \Omega(\delta, t)\} \\ &= \inf\{v(p, x) \mid (p, x) \in \Omega(\delta, t)\} \\ &= V(\delta, t). \quad \square \end{aligned}$$

B.3 Proof of Lemma 2.13

Recall the definition of $V(\delta, t)$,

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

In order to lower bound the above, we expand $v(p, x)$ and consider it as a function of x .

$$\begin{aligned} v(p, x) &= px^2 + \frac{\left(\frac{1}{2} + \alpha - px\right)^2}{1-p} - \left(\frac{1}{2} + \alpha\right)^2 \\ &= \left(\frac{p}{1-p}\right)x^2 - \frac{2p\left(\frac{1}{2} + \alpha\right)}{1-p}x \\ &\quad + \frac{\left(\frac{1}{2} + \alpha\right)^2}{1-p} - \left(\frac{1}{2} + \alpha\right)^2. \end{aligned}$$

For fixed p and α this expression is quadratic in x , with an axis of symmetry at $x = (\frac{1}{2} + \alpha)$. As $t \leq \frac{1}{2} - \alpha$, for all $x \geq 1-t$ we have that $x \geq 1-t \geq \frac{1}{2} + \alpha$ and so for any fixed p and $x \geq 1-t$, the function $v(p, x)$ is monotone increasing in x and therefore attains its minimum over the set $S_p \triangleq \{x \mid x \geq 1-t, \frac{1}{2} - \delta \leq px\}$ at the minimum $x \in S_p$; that is, at $x = \max\{1-t, (\frac{1}{2} - \delta)/p\}$. We consider the two cases corresponding to $p(1-t) \geq \frac{1}{2} - \delta$ and $p(1-t) \leq \frac{1}{2} - \delta$, for which the minimum is attained at $x = 1-t$ and $x = (\frac{1}{2} - \delta)/p$, respectively.

Case 1: For fixed $p \geq \frac{\frac{1}{2}-\delta}{1-t}$, the minimum $x \in S_p$ is $x = 1-t$ and so the minimum value of $v(p, x)$ over all $x \in S_p$ is $v(p, 1-t)$, which is precisely equal to

$$p(1-t)^2 + \frac{\left(\frac{1}{2} + \alpha - p(1-t)\right)^2}{1-p} - \left(\frac{1}{2} + \alpha\right)^2.$$

Taking the derivative of $v(p, 1-t)$ with respect to p , we find that this function is monotone increasing in p ,

$$\frac{\partial}{\partial p} \left[p(1-t)^2 + \frac{\left(\frac{1}{2} + \alpha - p(1-t)\right)^2}{1-p} - \left(\frac{1}{2} + \alpha\right)^2 \right]$$

$$= \frac{(t + \alpha - \frac{1}{2})^2}{(1-p)^2} \geq 0.$$

So, the minimal value of $v(p, x)$ with $p \geq \frac{\frac{1}{2}-\delta}{1-t}$ and $x \in S_p$ is precisely $v(\frac{\frac{1}{2}-\delta}{1-t}, 1-t)$.

Case 2: For fixed $p \leq \frac{\frac{1}{2}-\delta}{1-t}$, the minimum $x \in S_p$ is $x = (1/2 - \delta)/p$ and so the minimum value of $v(p, x)$ over all $x \in S_p$ is $v(p, (\frac{1}{2} - \delta)/p)$, which we rewrite as a function of $x = (1/2 - \delta)/p$ as $v((1/2 - \delta)/x, x)$ and expand below.

$$v\left(\frac{\frac{1}{2}-\delta}{x}, x\right) = \left(\frac{1}{2} - \delta\right)x + \frac{(\alpha + \delta)^2}{1 - \frac{\frac{1}{2}-\delta}{x}} - \left(\frac{1}{2} + \alpha\right)^2.$$

Again, taking the derivative, this time with respect to x , we find that

$$\begin{aligned} \frac{\partial}{\partial x} \left[\left(\frac{1}{2} - \delta\right)x + \frac{(\alpha + \delta)^2}{1 - \frac{\frac{1}{2}-\delta}{x}} - \left(\frac{1}{2} + \alpha\right)^2 \right] \\ = \frac{\left(\frac{1}{2} - \delta\right)(x + 2\delta + \alpha - \frac{1}{2})(x - \frac{1}{2} - \alpha)}{(x - \frac{1}{2} + \delta)^2} \geq 0. \end{aligned}$$

That is, this bound is monotone increasing in x ($= \frac{\frac{1}{2}-\delta}{p}$), or monotone decreasing in p , and therefore the minimal value of $v(p, x)$ with $p \leq \frac{\frac{1}{2}-\delta}{1-t}$ and $x \in S_p$ is precisely $v(\frac{\frac{1}{2}-\delta}{1-t}, 1-t)$.

In summary, we find that indeed,

$$\begin{aligned} & \inf\{v(p, x) \mid (p, x) \in \Omega(\delta, t)\} \\ &= \inf\{\inf\{v(p, x) \mid x \in S_p\} \mid p \in [0, 1]\} \\ &\geq v\left(\frac{\frac{1}{2}-\delta}{1-t}, 1-t\right). \quad \square \end{aligned}$$

C Proof of Theorem 3.1: Omitted Lemmas

In this section we prove Theorem 3.1 by proving two contradictory lemmas, which are stated in §3.

Because we are proving an impossibility result, we can focus without loss of generality on 3-location inputs with n players. We denote such inputs by $\mathbf{x} = \{(x_1, n_1), (x_2, n_2), (x_3, n_3)\}$, indicating that n_i players are at location x_i , with $x_1 \leq x_2 \leq x_3$. We denote the set of inputs of this form by \mathcal{I}_3 . For an instance $\mathbf{x} = \{(x_1, n_1), (x_2, n_2), (x_3, n_3)\} \in \mathcal{I}_3$, we denote by $S(\mathbf{x})$ the set of possible values of social cost when facilities are placed on player locations. For example, when $x_2 - x_1 \leq x_3 - x_2$, $S(\mathbf{x}) = \{(x_2 - x_1)n_1, (x_2 - x_1)n_2, (x_3 - x_2)n_3\}$, where the three elements correspond to the social costs obtained by putting facilities at $\{x_2, x_3\}$, $\{x_1, x_3\}$ and $\{x_1, x_2\}$ respectively. Finally, we denote by $\{(s_i, p_i) \mid s_i \in S(\mathbf{x}), i \in I \subseteq [3]\}$ a distribution of social costs, indicating that cost s_i is incurred with probability p_i .

C.1 Proof of Lemma 3.2

In this section we establish that for any family of mechanisms $\{f_\theta\}_{\theta \in [0,1]}$ satisfying the conditions of Theorem 3.1, the mechanism f_1 must have an unbounded approximation ratio for the social cost objective. We start by proving that f_1 must in fact be *deterministic*. To do so, we rely on the notion of *partial group strategyproofness*, or partial GSP for short, introduced by (Lu et al. 2010).

Definition C.1. A partial group strategyproof (*partial GSP*) mechanism for facility location problems is a mechanism for which a group of players at the same location cannot benefit from misreporting their locations simultaneously.

As (Lu et al. 2010) observed, SP implies partial GSP.

Lemma C.2 ((Lu et al. 2010)). Any SP mechanism for 2-facility location is also partial GSP.

Armed with Lemma C.2, we now move on to stating and proving our characterization of 0-variance SP mechanisms for 2-facility location social cost minimization. That is, we characterize SP mechanisms which always produce the same social cost on a given instance.

Lemma C.3. Restricted to 3-location instances \mathcal{I}_3 , all 0-variance SP mechanisms that place facilities on player locations are deterministic.

Proof. Fix a 0-variance SP mechanism f that always places facilities on player locations.

For a 3-location instance

$$\mathbf{x} = \{(x_1, n_1), (x_2, n_2), (x_3, n_3)\} \in \mathcal{I}_3$$

where $x_1 \leq x_2 \leq x_3$, let the *balance ratio* $r(\mathbf{x})$ of \mathbf{x} be such that

$$r(\mathbf{x}) = \begin{cases} (x_2 - x_1)/(x_3 - x_2), & \text{if } x_2 - x_1 \leq x_3 - x_2 \\ (x_3 - x_2)/(x_2 - x_1), & \text{otherwise} \end{cases}.$$

If $x_2 - x_1 \leq x_3 - x_2$, we call x_1 the *near end* of \mathbf{x} and x_3 the *far end*. Otherwise x_3 is the near end and x_1 is the far end. Particularly, when $x_2 - x_1 = x_3 - x_2$, both ends can be the near end or the far end. When talking about a particular instance, we scale the instance and the mechanism itself at the same time retaining all relevant properties, thereby drastically simplifying the discussion. We will show that both far and near end of an instance are both output deterministically. That is, each of these points is output with probability exactly 0 or 1. As f always chooses exactly two locations and places facilities on player locations, this implies that f is deterministic.

We first prove that on any instance $\mathbf{x} \in \mathcal{I}_3$, mechanism f outputs the far end with probability either 0 or 1. That is (up to rescaling), for any input $\mathbf{x} = \{(-t, a), (0, b), (1, c)\}$ where $t \leq 1$, if we let $A = -t, B = 0, C = 1$ denote respectively the leftmost, middle and rightmost group of players in the instance \mathbf{x} , then f outputs C with probability exactly 0 or 1. Clearly, $S(\mathbf{x}) = \{at, bt, c\}$. Suppose f places a facility at C with probability $p \in (0, 1)$; then the cost to players in C is $(1-p)$. Pick a small $\delta > 0$ such that $\delta < 1-p$, $1 + \delta \neq at$ and $1 + \delta \neq bt$. As a 0-variance mechanism, on

instance $\mathbf{x}' = \{(-t, a), (0, b), (1 + \delta, c)\}$, f cannot randomize nontrivially between putting a facility at $1 + \delta$ or not. If f puts a facility at $1 + \delta$ on \mathbf{x}' , the group C in \mathbf{x} will deviate to $1 + \delta$ such that their cost will decrease to $\delta < 1 - p$. If f does not put a facility at $1 + \delta$, players at $1 + \delta$ in \mathbf{x}' will deviate to 1, decreasing their cost from $1 + \delta$ to $p\delta + (1 - p)(1 + \delta)$. Partial GSP is ruined in both cases. We conclude that f acts deterministically on the far end of any instance. As a corollary, on any instance \mathbf{x} whose balance ratio is $r(\mathbf{x}) = 1$, mechanism f acts completely deterministically.

We now prove that on any instance, f outputs the near end with probability either 0 or 1. To this end, we first consider the instance $\mathbf{x} = \{(-1, a), (0, b), (1, c)\}$. By the previous paragraph, we have that, as $r(\mathbf{x}) = 1$, the probability that location -1 is output some $p \in \{0, 1\}$. We prove that for all $0 < t \leq 1$, on input $\mathbf{x}_t = \{(-t, a), (0, b), (1, c)\}$ mechanism f outputs location $-t$ with probability p_t precisely p , and in particular the probability of the near end being output is 0 or 1. There are two cases to consider, depending on the value of p .

Case 1: $p = 0$. If $p_t > 0$, players at -1 in \mathbf{x} will deviate to $-t$ in order to decrease their cost from 1 to $p_t \cdot (1 - t) + (1 - p_t)$, contradicting partial GSP. Therefore $p_t = 0 = p$.

Case 2: $p = 1$. This case is more intricate. We define a sequence $\{l_i\}_{i=0}^{\infty}$ where $l_0 = 1$ and $l_{i+1} = (l_i^2 + 2l_i)/(2.5 + 2l_i)$ and prove by induction that for all $k \geq 0$, on any input \mathbf{x}_t satisfying $r(\mathbf{x}_t) = t \geq l_k$, mechanism f outputs the near end of \mathbf{x}_t with probability $p_t = 1 (= p)$. The base case corresponds to $\mathbf{x}_t = \mathbf{x}$, and so trivially $p_t = p = 1$. For the inductive step, consider some instance \mathbf{x}_t with $r(\mathbf{x}_t) = t$ satisfying $l_i > t \geq l_{i+1}$, and suppose $p_t < 1$. By the inductive hypothesis, on input $\mathbf{x}' = \{(-l_i, a), (0, b), (1, c)\}$ the probability of f outputting $-l_i$ is 1. Therefore, by partial GSP, as group A in \mathbf{x} should not benefit from deviating to $-l_i$, we must have $(1 - p) \cdot t \leq l_i - t$, or put otherwise

$$p_t \cdot t \geq 2t - l_i. \quad (11)$$

On the other hand, consider the instance $\mathbf{x}'' = \{(-t, a), (\frac{l_i - t}{1 + l_i}, b), (1, c)\}$. Note that since

$$r(\mathbf{x}'') = \frac{(l_i - t)/(1 + l_i) + t}{1 - (l_i - t)/(1 + l_i)} = l_i,$$

by the induction hypothesis together with Case 1, f chooses the near and far end of \mathbf{x}'' with probability 0 or 1 each, and as f always outputs exactly two facilities, each on a distinct player location, this implies that f performs deterministically on \mathbf{x}'' . By partial GSP, location $\frac{l_i - t}{1 + l_i}$ in \mathbf{x}'' must get a facility, or else the players at this location will deviate to 0 in order to decrease their cost from $\frac{l_i - t}{1 + l_i} + t$ to at most $\frac{l_i - t}{1 + l_i} + p_t \cdot t$. Now, by Case 1, the far end of \mathbf{x}_t is chosen by f with probability 0 or 1. As f always outputs two facilities on input \mathbf{x}_t , the far end must therefore be chosen with probability precisely 1, else the expected number of output facilities would be strictly less than two. Likewise, group B

in \mathbf{x}_t must get a facility with probability precisely $1 - p_t$, and so the cost for players in group B on input \mathbf{x}_t is precisely $p_t \cdot t$. Consequently, again invoking partial GSP of f , we find that the players at group B in \mathbf{x}_t must not benefit from deviating to $\frac{l_i - t}{1 + l_i}$ and so we must have

$$p_t \cdot t \leq \frac{l_i - t}{1 + l_i}. \quad (12)$$

Combining Equations (11) and (12), we get

$$2t - l_i \leq \frac{l_i - t}{1 + l_i},$$

which implies that $t \leq \frac{l_i^2 + 2l_i}{3 + 2l_i} < \frac{l_i^2 + 2l_i}{2.5 + 2l_i} = l_{i+1}$, a contradiction, and so we conclude that $p_t = 1$.

We are still to show that l_k tends to 0 as k tends to infinity. Note that $l_k > 0$ for all k , and

$$\frac{l_{i+1}}{l_i} = \frac{l_i + 2}{2l_i + 2.5} \leq \max\left\{\frac{l_i}{2l_i}, \frac{2}{2.5}\right\} = \frac{2}{2.5}.$$

Therefore $0 < l_k \leq (\frac{2}{2.5})^k$. Clearly $\lim_{k \rightarrow \infty} l_k = 0$.

From the above we conclude that for a 3-location instance $\mathbf{x} \in \mathcal{I}_3$, if $r(\mathbf{x}) > 0$, f does not randomize nontrivially on both ends of \mathbf{x} . If $r(\mathbf{x}) = 0$, \mathbf{x} must be a 2-location, or even a 1-location instance, on which there is only one way to put 2 facilities. Altogether we conclude that, f acts deterministically on both ends of any 3-location instance, or equivalently, f is deterministic restricted to 3-location instances. \square

Given Lemma C.3, we may safely assume that f_1 is a deterministic mechanism whenever restricted to 3-location instances. We will rely on the following characterization of deterministic SP mechanisms for the 2-facility location problem, established by (Fotakis and Tzamos 2013a, Theorem 3.3).

Lemma C.4 ((Fotakis and Tzamos 2013a)). *Let f be any SP mechanism for 2-facility location with a bounded approximation ratio for the social cost. Then, restricted to 3-location instances with $n \geq 5$ players, either there exists a unique dictator $j \in [n]$ such that for all instances $\mathbf{x} \in \mathcal{I}_3$ a facility is allocated to player j , or for all instances $\mathbf{x} \in \mathcal{I}_3$ the two facilities are placed on $lt(\mathbf{x})$ and $rt(\mathbf{x})$.*

Using this characterization and Lemma C.3 we can now prove Lemma 3.2.

Proof of Lemma 3.2. We prove that f_1 neither chooses the two extremes nor has a dictator, and therefore by Corollary C.4 is not a bounded mechanism. Let α be the approximation ratio of f_0 . Consider instance $\mathbf{x} = \{(-1, n), (0, n), (1, 1)\}$ (i.e., n players at -1 , n at 0 and 1 at 1) where $n \geq \max\{3\alpha, 2\}$. Clearly $S(\mathbf{x}) = \{1, n\}$. Let $C_0 = C(f_0, \mathbf{x})$. Then, by virtue of f_0 being α -approximate and by Markov's Inequality, we have

$$\Pr[C_0 = n] \leq \Pr[C_0 \geq n] \leq \frac{\mathbb{E}[C_0]}{n} \leq \frac{\alpha}{3\alpha} = \frac{1}{3}. \quad (13)$$

If the deterministic mechanism f_1 puts a facility at 1, thereby producing social cost $C(f_1, \mathbf{x}) = n$, then by continuity of expected social cost, there is a $0 < \theta' < 1$ satisfying $\mathbb{E}[C(f_{\theta'}, \mathbf{x})] = \frac{1}{2}(1 + n)$, and therefore $\Pr[C(f_{\theta'}, \mathbf{x}) = n] = \frac{1}{2}$. Pick such a θ' and let $C_{\theta'} = C(f_{\theta'}, \mathbf{x})$. For a random variable C chosen from the distribution $\{(1, 1 - p), (n, p)\}$ we have $\text{Var}(C) = (n - 1)^2 \cdot (p - p^2)$, which is monotone increasing in p for all $p \leq \frac{1}{2}$. By Equation (13) we thus obtain

$$\begin{aligned} \text{Var}(C_0) &\leq \text{Var}(\{(1, 2/3), (n, 1/3)\}) \\ &< \text{Var}(\{(1, 1/2), (n, 1/2)\}) = \text{Var}(C_{\theta'}), \end{aligned}$$

and also clearly $\text{Var}(C_{\theta'}) > 0 = \text{Var}(C_1)$, a contradiction to monotonicity of $\text{Var}(f_\theta, \mathbf{x})$.

We conclude that, given the location vector \mathbf{x} , f_1 puts facilities at -1 and 0 . In particular, f_1 neither chooses the two extremes (which are -1 and 1) nor has a dictator (because any player can be the one located at 1), and hence has an unbounded approximation ratio. \square

C.2 Proof of Lemma 3.3

In this section we establish that for any family of SP mechanisms $\{f_\theta\}_{\theta \in [0,1]}$ satisfying the conditions of Theorem 3.1, the mechanism f_1 must have a bounded approximation ratio for the social cost objective.

Lemma C.5. *Let $\{f_\theta\}_{\theta \in [0,1]}$ be a family of SP mechanisms satisfying the conditions of Theorem 3.1. If f_0 restricted to n -player 3-location instances has a bounded approximation ratio $\alpha = \alpha(n)$, then for any n -player 3-location input $\mathbf{x} \in \mathcal{I}_3$, if $S(\mathbf{x}) = \{s_1, s_2, s_3\}$ and $s_3 > 40\alpha \cdot \text{opt}(\mathbf{x})$, then $C(f_1, \mathbf{x}) \neq s_3$.*

Proof. Without loss of generality let $s_1 = 1$. In addition, let $t > 20$. Assuming $s_3 = 2t\alpha > 20\alpha$, we proceed by 2 cases. Throughout the proof we rely on the previously-stated simple observation that for a random variable C chosen from the distribution $\{(1, 1 - p), (z, p)\}$ we have $\text{Var}(C) = (z - 1)^2 \cdot (p - p^2)$, which is monotone increasing in p for all $p \leq \frac{1}{2}$ and monotone decreasing in p for $p \geq \frac{1}{2}$.

Case 1: $s_2 > t \cdot \alpha$. We prove that $C(f_1, \mathbf{x}) = s_1$. Otherwise, $C(f_1, \mathbf{x}) \geq s_2 > \alpha \geq C(f_0, \mathbf{x})$ and by continuity of expected social cost there exists some $\theta \in (0, 1)$ such that $\mathbb{E}[C(f_\theta, \mathbf{x})] = \frac{1}{2}(s_1 + s_2)$. Let $C_0 = C(f_0, \mathbf{x})$ and $C_\theta = C(f_\theta, \mathbf{x})$. Since f_0 is α -approximate, by Markov's Inequality we have

$$\Pr[C_0 = s_1] = 1 - \Pr[C_0 \geq s_2] \geq 1 - \frac{\alpha}{t\alpha} = 1 - \frac{1}{t}.$$

Therefore, as shifting all the mass of C_0 's distribution from cost $s_2 > t \cdot \alpha > \alpha = \mathbb{E}[C_0]$ to cost s_3 can only serve to increase the variance, and by our observation that $\text{Var}(\{(1, 1 - p), (z - p)\})$ is monotone increasing in $p \leq \frac{1}{2}$ (e.g., $\frac{1}{t} \leq \frac{1}{2}$), we have

$$\begin{aligned} \text{Var}(C_0) &\leq \text{Var}(\{(s_1, 1 - 1/t), (s_3, 1/t)\}) \\ &= \text{Var}(\{(1, 1 - 1/t), (2t\alpha, 1/t)\}) \end{aligned}$$

$$\begin{aligned} &= (2t\alpha - 1)^2 \cdot (1/t - 1/t^2) \\ &\leq 4t^2\alpha^2 \cdot (1/t) \\ &= 4t\alpha^2. \end{aligned}$$

On the other hand, for C_θ we have $\mathbb{E}[C_\theta] = \frac{1}{2}(s_1 + s_2)$ and so shifting all the mass from s_3 to s_1 and part of the mass from s_2 to s_1 (in order to keep the expected cost unchanged) can only decrease the variance,⁶ we have

$$\begin{aligned} \text{Var}(C_\theta) &\geq \text{Var}(\{(s_1, 1/2), (s_2, 1/2)\}) \\ &\geq \text{Var}(\{(1, 1/2), (t\alpha, 1/2)\}) \\ &= (t\alpha - 1)^2 \cdot \left(\frac{1}{2} - \frac{1}{4}\right) \\ &= \frac{t^2\alpha^2 - 2t\alpha + 1}{4} \\ &> \frac{t^2\alpha^2 - 2t\alpha}{4}. \end{aligned}$$

But for $t > 20$ this implies that

$$\text{Var}(C_\theta) - \text{Var}(C_0) > 0,$$

contradicting monotonicity of variance. Therefore in this case, $C(f_1, \mathbf{x}) = s_1 \neq s_3$.

Case 2: $s_2 \leq t \cdot \alpha$. If $C(f_1, \mathbf{x}) = s_3$, then by continuity of expected social cost there exists some $\theta \in (0, 1)$ such that $\mathbb{E}[C(f_\theta, \mathbf{x})] = \frac{1}{2}(s_2 + s_3)$. Let $C_0 = C(f_0, \mathbf{x})$, $C_\theta = C(f_\theta, \mathbf{x})$. Again, by Markov's Inequality and f_0 being α -approximate, we have

$$\Pr[C_0 = s_3] \leq \Pr[C_0 \geq s_3] \leq \frac{\alpha}{s_3} = \frac{1}{2t},$$

and so we have, by a similar argument to Case 1, that

$$\text{Var}(C_0) \leq \text{Var}(\{(s_1, 1 - 1/2t), (s_3, 1/2t)\}) \leq 2t\alpha^2.$$

On the other hand,

$$\begin{aligned} \text{Var}(C_\theta) &\geq \text{Var}(\{(s_2, 1/2), (s_3, 1/2)\}) \\ &\geq \text{Var}(\{(t\alpha, 1/2), (2t\alpha, 1/2)\}) \\ &= t^2\alpha^2 \cdot \text{Var}(\{(0, 1/2), (1, 1/2)\}) \\ &= \frac{1}{4}t^2\alpha^2. \end{aligned}$$

But as $t > 20 > 8$, we have

$$\text{Var}(C_\theta) - \text{Var}(C_0) = \frac{t\alpha^2}{4}(t - 8) > 0,$$

again contradicting monotonicity of variance. Therefore in this case, too, $C(f_1, \mathbf{x}) \neq s_3$.

We conclude that when $s_3 = 2t\alpha > 40\alpha$, we have $C(f_1, \mathbf{x}) \neq s_3$. \square

⁶This, as the difference between s_3 and $\mathbb{E}[C_\theta] = \frac{1}{2}(s_1 + s_2)$ is greater than the differences between both other costs and $\frac{1}{2}(s_1 + s_2)$, which are both equal to $\frac{1}{2}(s_2 - s_1)$.

Lemma C.6. For an n -player 3-location instance \mathbf{x} , if $S(\mathbf{x}) = \{s_1, s_2, s_3\}$ where $s_1 \leq s_2 \leq s_3$, then $s_2 \leq (n-2) \cdot s_1 = (n-2) \cdot \text{opt}(\mathbf{x})$.

Proof. Without loss of generality suppose $\mathbf{x} = \{(-1, a), (0, b), (t, c)\}$, where $a + b + c = n$, $a, b, c \geq 1$ and $t \geq 1$, in which case for all $d, e \in \{a, b, c\}$, we have $\frac{d}{e} \leq n-2$. Clearly $S(\mathbf{x}) = \{a, b, ct\}$. Now, regardless of the ordering of $S(\mathbf{x})$, we find that s_1 is at least some e in $\{a, b, c\}$, as $t \geq 1$. Moreover, s_2 is at most some d in $\{a, b, c\}$, as $s_2 \leq s_3$ and either $s_2 \neq ct$ or $s_3 \neq ct$. Consequently we find that

$$\frac{s_2}{s_1} \leq \frac{d}{e} \leq n-2. \quad \square$$

Proof of Lemma 3.3. Let $\alpha = \alpha(n)$ be the approximation ratio of f_0 restricted to n -player 3-location instances. For any 3-location instance \mathbf{x} , if $s_3 \leq 40\alpha \cdot \text{opt}(\mathbf{x})$, mechanism f_1 is 40α -approximate. Else, by Lemma C.5 and Lemma C.6, $C(f_1, \mathbf{x}) \leq s_2 \leq (n-2) \cdot s_1$, and so f_1 is $(n-2)$ -approximate. In both cases the approximation ratio of f_1 is bounded by $\max\{(n-2), 40\alpha\}$ for all \mathbf{x} . \square