Single-Call Mechanisms

CHRISTOPHER A. WILKENS, University of California at Berkeley
BALASUBRAMANIAN SIVAN, University of Wisconsin-Madison

Truthfulness is fragile and demanding. It is oftentimes computationally harder than solving the original problem. Even worse, truthfulness can be utterly destroyed by small uncertainties in a mechanism's outcome. One obstacle is that *truthful payments depend on outcomes other than the one realized*, such as the lengths of non-shortest-paths in a shortest-path auction. Single-call mechanisms are a powerful tool that circumvents this obstacle — they implicitly charge truthful payments, guaranteeing truthfulness in expectation using only the outcome realized by the mechanism. The cost of such truthfulness is a trade-off between the expected quality of the outcome and the risk of large payments.

We largely settle when and to what extent single-call mechanisms are possible. The first single-call construction was discovered by Babaioff, Kleinberg, and Slivkins [2010] in single-parameter domains. They give a transformation that turns any monotone, single-parameter allocation rule into a truthful-in-expectation single-call mechanism. Our first result is a natural complement to [Babaioff et al. 2010]: we give a new transformation that produces a single-call VCG mechanism from any allocation rule for which VCG payments are truthful. Second, in both the single-parameter and VCG settings, we precisely characterize the possible transformations, showing that that a wide variety of transformations are possible but that all take a very simple form. Finally, we study the inherent trade-off between the expected quality of the outcome and the risk of large payments. We show that our construction and that of [Babaioff et al. 2010] simultaneously optimize a variety of metrics in their respective domains.

Our study is motivated by settings where uncertainty in a mechanism renders other known techniques untruthful. As an example, we analyze pay-per-click advertising auctions, where the truthfulness of the standard VCG-based auction is easily broken when the auctioneer's estimated click-through-rates are imprecise.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics; F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Incentive compatibility, single-call, efficiency, black-box reductions

1. INTRODUCTION

In their seminal work that sparked the field of Algorithmic Mechanism Design, Nisan and Ronen [2001] made a striking observation: naïvely computing VCG payments for shortest-path auctions requires computing "n versions of the original problem." In their case, it requires solving $n + 1$ different shortest path problems in a network. Over the next decade, as researchers studied computation in mechanisms, they repeatedly noticed that computing payments is harder than solving the original problem. Babaioff et al. [2008] exhibited a problem for which deterministic truthfulness is precisely $(n+1)$-times harder than the original problem. In the case of Nisan and Ronen's own path

Part of this work was done while the authors were interns at Microsoft Research, Redmond, WA.
Author's addresses: C. Wilkens, Computer Science Division, University of California at Berkeley; B. Sivan, Computer Sciences Department, University of Wisconsin-Madison.

auction, Hershberger et al. [2007] showed that computing VCG prices for a directed graph requires time equivalent to $\sqrt{n}$ shortest path computations.[1]

Surprisingly, Babaioff, Kleinberg, and Slivkins [2010] recently showed that randomization eliminated this difficulty for a large class of problems. They showed that, if in a single-parameter domain payments need only be *truthful in expectation*, then they may be computed by solving the original problem only once. They apply their result to Nisan and Ronen's path auctions to get a truthful-in-expectation mechanism that uses precisely one shortest-path computation and chooses the shortest path with probability arbitrarily close to 1. We call this a *single-call mechanism*.

The usefulness of Babaioff, Kleinberg, and Slivkins' result goes far beyond speeding up computation: *Their construction enables truthfulness in cases in which computing "n versions of the original problem" is informationally impossible.* To use again the Nisan-Ronen path auction, suppose that the graph represents a packet network with existing traffic. In this case, the actual transit times (i.e. costs to edges) may be increased by congestion. While it is possible to estimate congestion ex ante, it is generally impossible to precisely know its effect without transmitting a packet and explicitly measuring its transit time. Unfortunately, since VCG prices depend on the transit times for many different paths, naïvely computing them will inherit any estimation errors. Even worse, when bidders have conflicting beliefs about such errors, *naïvely computing"VCG" prices with bad estimates may not guarantee truthfulness* even if the errors are small enough that they not affect the path chosen by the mechanism. In such a case, truthfulness may be regained using a mechanism that only requires measurements along a single path, that is, a mechanism that only requires measurements returned by a single call to the shortest-path algorithm. We will concretely demonstrate this phenomenon later using an example based on pay-per-click advertising auctions.

An important question arises then: *In which mechanism design problems, and to what extent, are single-call mechanisms possible?* In this paper we study, and largely settle, this question. First, we show that this it is possible to transform any mechanism that charges VCG prices in expectation into a roughly equivalent single-call mechanism. While similar in spirit to [Babaioff et al. 2010], our reduction charges prices that are fundamentally different from the mechanism in that paper — they do not coincide even when applied to the same allocation rule. Second, we give characterization theorems, delineating precisely the single-call mechanisms that are possible, for both the VCG and single-parameter settings. Finally, single-call constructions offer a tradeoff between expectation and risk. Our characterization theorems allow us to derive lower bounds on this tradeoff, establishing that our VCG construction and the construction of [Babaioff et al. 2010] are optimal in a general sense.

Mechanisms, Allocations, and Payments. One cornerstone of mechanism design is the decomposition of a mechanism into two distinct parts: an allocation function and a payment function. This approach has borne much fruit — it first revealed fundamental relationships between allocation functions and their nearly unique truthful prices, and it subsequently allowed researchers to study the the two problems in isolation. Like [Babaioff et al. 2010], we leverage this decomposition to study payment techniques that apply to large classes of allocation functions — naturally, our primary requirement is that the allocation function may only be evaluated once.

[1]Interestingly, the undirected case is easier. Hershberger and Suri [2001; 2002] show that it only requires time equivalent to a single shortest-path computation. Their work is orthogonal to our own — single-call mechanisms achieve truthfulness in a limited-information setting using only one shortest-path computation, while [Hershberger et al. 2007; Hershberger and Suri 2001; 2002] assume complete information and study an algorithmic problem.

We will focus on single-call mechanisms for two classes of allocation functions that, together, comprise most allocation functions for which truthful payments are known: *monotone single-parameter functions* and *maximal in distributional range (MIDR) functions*.

An allocation function is said to be *single-parameter* if an agent's bid can be expressed as a single number. This setting was first studied by Myerson [1981] in the context of single-item auctions. Subsequent generalizations showed that truthful prices existed if and only if a single-parameter allocation is monotone and provided an explicit characterization of truthful payments. We will use one such characterization developed by Archer and Tardos [2001].

An allocation function is said to be *maximal in distributional range* (MIDR) if, for some fixed set of distributions over outcomes, the allocation always chooses one that maximizes the social welfare of the bidders. MIDR allocation functions are important because they are precisely the ones for which VCG payments are truthful [Dobzinski and Dughmi 2009].

Truthfulness Under Uncertainty. Our motivation for developing and optimizing single-call mechanisms comes from scenarios where nature prohibits computing an allocation more than once, most often due to parameter uncertainty. We give a few examples here; more generally, we conjecture that most mechanism design problems have similar variants.

In the uncertain shortest-path auction described earlier, truthful prices will depend on the incremental effect of transit times adjusted for congestion. If the auctioneer generates the network traffic, he may be able to predict the congestion in an edge better than the edge itself and use this prediction when computing the shortest path. However, each edge may individually disagree with the auctioneer's estimate, and these beliefs are generally unknown to the auctioneer. If the auctioneer were to simply compute VCG payments by combining his estimates with players' bids, the prices would likely not be truthful. On the other hand, we can require that payments are computed using measured transit times instead of estimates; however, it is informationally impossible to know the precise delay along edges that were not actually traversed. A single-call mechanism sidesteps this hurdle by using only the delays along traversed edges for which the delay had been precisely known.

Machine scheduling offers another application for single-call mechanisms. In some applications (e.g. cloud services), it is common for machines to bid in terms of cost per unit time (or other resource). It is then the responsibility of the scheduler to estimate the time required for the job on that machine. If the scheduler's estimates differ from a machine's belief about a job's runtime, then we find ourselves in the same situation as the path auction — the standard truthful prices for this single-parameter setting will depend on machines' beliefs about the runtimes of jobs under alternate schedules. A single-call mechanism sidesteps this problem because it requires only the runtimes of jobs under the schedule chosen by the mechanism, which may be measured.

Another interesting example arises in the application of learning procedures such as multi-arm-bandits (MABs). In recurring mechanisms, it is natural for the auctioneer to run a learning algorithm across multiple auctions. For example, when an online advertising auction is repeated, the auctioneer tries to learn the likelihood that a particular ad will get clicked. Computing truthful prices requires knowing what would have happened if the learner had been initialized with a different set of bids. This setting was the original motivation of [Babaioff et al. 2010], where they showed that their single-call construction allowed a MAB to be implemented truthfully with $O(\sqrt{T})$ regret. This contrasts with results of Babaioff, Sharma, and Slivkins [2009] and Devanur and Kakade [2009] who showed that any universally truthful mechanism must have regret at least $\Omega(T^{\frac{2}{3}})$ for different measurements of regret.

Finally, in Section 5 we analyze *single-shot pay-per-click (PPC) advertising auctions*. A PPC advertising auction ranks bidders using their pay-per-click bid (i.e. they only pay when they receive a click) and an estimate of the probability of a click (the click-through rate, or CTR). If the bidders' estimates of their own CTRs are different from the auctioneer's, truthful prices necessarily depend on bidders' beliefs about the CTRs, which are unknown.

Single-Call Mechanisms and Reductions. Our tool for creating single-call mechanisms is *the single-call reduction*, the main object of study in this paper. A single-call reduction is a transformation that takes an allocation function as a black box and produces a truthful-in-expectation mechanism that calls the allocation function once.Since the expected payment is equal to the truthful payment for the resulting mechanism, the payments are dubbed *implicit*.

Babaioff, Kleinberg, and Slivkins [2010] discovered such a reduction for single-parameter domains. Using only the guarantee that the black-box allocation rule is monotone, their reduction produces a truthful-in-expectation mechanism that implements the same outcome as the original allocation rule with probability arbitrarily close to 1.[2]

VCG is a mechanism design framework much broader than single-parameter. *Can we construct similar single-call mechanisms that charge VCG prices?* We answer this in the affirmative by giving a reduction producing, for any MIDR allocation function, a single-call mechanism that charges VCG prices in expectation. Analogous to [Babaioff et al. 2010], our reduction transforms any MIDR allocation rule into a truthful-in-expectation mechanism that implements the same outcome as the original allocation rule with probability arbitrarily close to 1. However, our construction is fundamentally different in that the distribution of payments does not coincide with [Babaioff et al. 2010] when an allocation is both MIDR and single-parameter. This reduction can guarantee truthfulness in multi-parameter mechanisms with uncertainty, as described above, and can also be used to speed up payment computation in MIDR settings like Dughmi and Roughgarden's [2010] truthful FPTAS for welfare-maximization packing problems.

We next ask *what single-call reductions are possible*? Babaioff et al. generalize to a class of self-resampling procedures. Subsequent research [Hartline 2011] generalized further (and simplified substantially), but concisely characterizing single-call reductions remained an open question. We give tight characterization theorems, showing that a wide variety of reductions are possible and that payments have a very simple characterization in both scenarios. The key technical idea is a simple proof equating a reduction's expected payments with those required for truthfulness, giving a sharp characterization of the parameters in the reduction. Our technique is a very simple alternative to the contraction mapping argument in [Babaioff et al. 2010].

Finally, we ask *what are the best single-call reductions*? As noted above, known single-call reductions choose an outcome different from the original allocation rule with some small probability δ. The penalty for making δ small is that the payments may occasionally be very large — we study this tradeoff. Our study is not unprecedented: [Babaioff et al. 2010] asked, as an open question, if their reduction optimized payments with respect to the welfare loss, and Lahaie [2010] show a similar trade-

[2]The authors of [Babaioff et al. 2010] have observed that their construction may be extended to any domain where the bid space is convex.

off between the size and complexity of kernel-based payments achieving ϵ-incentive compatibility in single-call combinatorial auctions.

We study the tradeoff inherent to single-call mechanisms with respect to three measures of expectation — welfare, revenue, and a technical (but natural) precision metric — and two measures of risk — variance and worst-case payments. We show that our VCG reduction and the single-parameter reduction of [Babaioff et al. 2010] *simultaneously* optimize the tradeoff between expectation and risk for all these criteria.

2. PRELIMINARIES

A mechanism is a protocol among n rational agents that implements a social choice function over a set of outcomes $\mathcal{O}$. Agent i has preferences over outcomes $o \in \mathcal{O}$ given by a *valuation function* $v_i : \mathcal{O} \to \mathbf{R}$. The function v_i is private but is drawn from a publicly known set $V_i \subseteq \mathbf{R}^{\mathcal{O}}$.

A *deterministic direct revelation mechanism* $\mathcal{M}$ is a social choice function $A : V_1 \times \ldots V_n \to \mathcal{O}$, also known as an *allocation rule*, and a vector of payment functions $P_1, \ldots, P_n$ where $P_i : V_1 \times \ldots V_n \to \mathbf{R}$ is the amount that agent i pays to the mechanism designer. When a direct revelation[3] mechanism is instantiated, each agent reports a *bid* $b_i \in V_i$. The mechanism uses bids $b = (b_1, \ldots, b_n)$ to choose an outcome $A(b) \in \mathcal{O}$ and to compute payments $P_i(b)$. The utility $u_i(v_i, o)$ that agent i receives is $u_i(v_i, o) = v_i(o) - P_i$.

A mechanism is *truthful* (or incentive compatible) if bidding truthfully (i.e. $b_i = v_i$) is a dominant strategy. Formally, for each i, each $v_{-i} \in V_{-i}$, and every $v_i, v_i{}' \in V_i$, we have $u_i(v_i, A(v)) \geq u_i(v_i, A(v_i{}', v_{-i}))$, where v_{-i} denotes the vector of valuations for all agents except agent i.

A mechanism is *ex-post individually rational* (IR) if agents always get non-negative utility, and mechanism has *no positive transfers* (NPT) if for each agent i and each $v \in V$, $P_i(v) \geq 0$, i.e., the mechanism never pays a player money.

A randomized mechanism is a distribution over deterministic mechanisms. Thus, $A(b)$ and $P_i(b)$ are random variables. For randomized mechanisms, properties like truthfulness may be said to hold universally or in expectation. A randomized mechanism is *universally truthful* if it is truthful for every deterministic mechanism in its support. It is *truthful in expectation* if, in expectation over the randomization of the mechanism, truthful bidding is a dominant strategy. Henceforth, we use truthful, IR, and NPT to mean truthful in expectation unless otherwise noted.

MIDR Allocation Rules. MIDR mechanisms are variants of *VCG mechanisms*, mechanisms that maximize social welfare and charge "VCG payments". Formally, a VCG mechanism's social choice rule satisfies $A(v) \in \underset{o \in \mathcal{O}}{\operatorname{argmax}} \sum_j v_j(o)$, and its payments are $P_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j(A(v))$ for some function $h_i : V_{-i} \to \mathbb{R}$. VCG payments are the only universal technique known to induce truthful bidding. The most common implementation of VCG payments uses the Clarke-Pivot payment rule: set $h_i(v_{-i}) = \max_{o \in \mathcal{O}}(\sum_{j \neq i} v_j(o))$, which gives the only payments that simultaneously satisfy truthfulness, IR, and NPT.

More generally, any allocation rule that maximizes an affine function of agents' valuations can be truthfully implemented with VCG payments. Moreover, Roberts' theorem [Roberts 1979] implies that in a general setting (when $V_i = \mathbf{R}^{\mathcal{O}}$), if A is onto (every outcome can be realized), then A has truthful payments if and only if it is an

[3]"Direct revelation" means that an agent's bid b_i is an element of V_i. In general this need not be the case; however, by the revelation principle, any social choice rule that may be truthfully implemented may be implemented as a direct revelation mechanism that charges the same payments in equilibrium.

affine maximizer. If the "onto" restriction is relaxed, a social choice function is truthfully implementable with VCG payments if and only if it is (weighted) maximal-in-range (MIR) [Nisan and Ronen 2007] or, for randomized mechanisms, maximal-in-distributional-range (MIDR) [Dobzinski and Dughmi 2009]:

Definition 2.1. An allocation rule A is ***MIDR*** if there is a set $\mathcal{D}$ of probability distributions over outcomes such that A outputs a random sample from the distribution $D \in \mathcal{D}$ that maximizes expected welfare. Formally, for each $v \in V$, $A(v) = o \sim D^*$ where $D^* \in \underset{D \in \mathcal{D}}{\operatorname{argmax}} \mathbf{E}_{o \sim D}[\sum_i v_i(o)]$.

A weighted MIDR allocation rule maximizes the weighted social welfare $\sum_i w_i v_i(o)$ for $w_i \geq 0$.

Single-Parameter Domains. A larger class of social choice rules can be implemented when V_i is single dimensional. We say that a social choice rule has a single-parameter domain if $v_i(o) = t_i f_i(o)$ for some publicly known function $f_i : \mathcal{O} \to \mathbb{R}_+$. The value $t_i \in T_i$ is an agent's type (T_i is her type-space, and $T = T_1 \times \cdots \times T_n$), and submitting i's bid precisely requires stating $b_i = t_i$. When $T = \mathbb{R}^n_+$, we say that bidders have *positive types*. We also use $A_i(b) = f_i(A(b))$ as shorthand, and we say A is bounded if the functions A_i are bounded functions.

A single-parameter social choice rule may be implemented if and only if it is *monotone*, where $A : T \to \mathcal{O}$ is said to be monotone if for each agent i, for all $b_{-i} \in T_{-i}$ and for every two bids $b_i \geq b'_i$, we have $A_i(b_i, b_{-i}) \geq A_i(b'_i, b_{-i})$. This was first shown for a single item auction by Myerson [1981]; Archer and Tardos [2001] gave the current generalization:

THEOREM 2.2. *[Myerson + Archer-Tardos] For a single parameter domain, an allocation rule A has truthful payments $(P_1, \ldots, P_n)$ if and only if A is monotone. These payments take the form*

$$P_i(b) = h_i(b_{-i}) + b_i A_i(b_i, b_{-i}) - \int_0^{b_i} A_i(u, b_{-i})\, du,$$

where $h_i(b_{-i})$ is independent of b_i.

These payments simultaneously satisfy IR and NPT if and only if $P_i{}^0(b_{-i}) = 0$. Such a mechanism is said to be normalized.

3. SINGLE-CALL MECHANISMS

We call a mechanism a single-call mechanism if it only evaluates the allocation function once:

Definition 3.1. A *single-call mechanism* $\mathcal{M}$ for an allocation rule A is a truthful mechanism that has only oracle access to A and computes both the allocation and payments with a single call to A.

To construct a single-call mechanism, we must first specify the possible allocation functions A and then construct one procedure that yields a single-call mechanism for any A in this set. Thus, the tool for creating a single-call mechanism is a single-call reduction:

Definition 3.2. A *single-call reduction* is a procedure that takes any allocation function A from a fixed set (as a black box) and returns a single-call mechanism.

For example, the procedure of [Babaioff et al. 2010] is a single-call reduction that takes any A drawn from the set of all monotone, bounded, single-parameter allocation rules

ALGORITHM 1: Generic Single-Call Reduction $(\mu, \{\lambda_i\})$

input : Black box access to an allocation function A, which is drawn from a known set.
output: Truthful-in-expectation mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$.

Solicit bid vector b from agents;
Sample $\hat{b} \sim \mu_b$;
Realize the outcome $A(\hat{b})$; `//` $\mathcal{A}(b)$ `is the random function` $A(\hat{b})$ `where` $\hat{b} \sim \mu_b$
Charge payments $\lambda(A(\hat{b}), \hat{b}, b)$; `//` $\mathcal{P}_i(b)$ `is the random function` $\lambda_i(A(\hat{b}), \hat{b}, b)$ `where` $\hat{b} \sim \mu_b$

and returns a single-call mechanism. Similarly, our construction for VCG prices is a single-call reduction that takes any A that is MIDR and returns a single-call mechanism.

To formalize single-call reductions, we first note the following requirements:

— A reduction must take a bid vector b and a black-box allocation function A as input.
— A reduction must evaluate A on at most one bid vector $\hat{b}$, causing the outcome $A(\hat{b})$ to be realized.[4]
— A reduction must charge payments λ_i that are a function of b, $\hat{b}$, and $A(\hat{b})$ (and possibly its own randomness).

These requirements suggest the following generic definition of a single-call reduction to turn an allocation function A into a truthful-in-expectation single-call mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$:

(1) Solicit the bid vector b from agents.
(2) Use b to compute the modified bid vector $\hat{b}$. This implicitly defines a probability measure $\mu_b(B)$ denoting the probability of choosing $\hat{b} \in B \subseteq V_1 \times \cdots \times V_n$ as the modified (resampled) bid vector when b is the actual bid vector. When $\hat{b}_i \neq b_i$, we say that i's bid was resampled.
(3) Declare the outcome to be $A(\hat{b})$, i.e. evaluate A at the modified bid vector $\hat{b}$. This implicitly defines the allocation function $\mathcal{A}(b)$ which samples $\hat{b} \sim \mu_b$ and chooses the outcome $A(\hat{b})$. The resampling procedure must ensure that truthful payments $\mathcal{P}(b)$ exist for $\mathcal{A}(b)$; Note that $\mathcal{A}(b)$ and $\mathcal{P}(b)$ are random variables that depend on the randomly resampled bid vector $\hat{b}$. Also, $\mathcal{A}(b)$ and $\mathcal{P}(b)$ are randomized even if $A(b)$ and $P(b)$ are deterministic;
(4) Use b, $\hat{b}$, and $A(\hat{b})$ to compute payments $\lambda_i(A(\hat{b}), \hat{b}, b)$ that satisfy truthfulness in expectation, that is, charge player i a payment $\lambda_i(A(\hat{b}), \hat{b}, b)$ such that $\mathop{\mathbf{E}}_{\hat{b}}[\lambda_i(A(\hat{b}), \hat{b}, b)] = \mathop{\mathbf{E}}_{\hat{b}}[\mathcal{P}_i(b)]$.

This general procedure is illustrated in Algorithm 1.

We describe a single-call reduction in the above framework by the tuple $(\mu, \{\lambda_i\})$, where μ implies specifying the resampling measure μ_b for all $b \in V_1 \times \cdots \times V_n$. Since payments should be finite, we require that λ_i be finite everywhere, and we also require that it be integrable. For the rest of this paper, we assume that λ_i's are deterministic.

[4]Strictly speaking, there may be settings where a single-call reduction could realize an outcome other than $A(\hat{b})$. However, our restriction follows naturally in scenarios where "computing $A(b)$" means realizing $A(b)$ and making measurements. It is also required for complete generality because there is no reason to believe that the designer knows how to realize any outcome other than $A(\hat{b})$.

For randomized λ_i's, the characterization theorems still hold with λ_i's replaced by their expectations over the randomness used.

We say that a reduction is *normalized* if $b_i(A(b)) = 0$ for all i implies $\lambda_i(A(\hat{b}), \hat{b}, b) = 0$, i.e. when every agent receives zero value, all payments are zero.

3.1. Optimal Reductions — Expectation vs. Risk

There are two downsides to the mechanisms produced by single-call reductions. First, there is a penalty in *expectation*, i.e., the expected outcome $\mathbf{E}_{\hat{b}}[A(\hat{b})]$ produced by the reduction is not identical to the desired outcome, $A(b)$. This modified outcome may reduce the expected welfare or revenue of the mechanism, or it may simply cause it to do the "wrong" thing.

Second, there is a penalty in *risk* because the payments λ may vary significantly, i.e. for a fixed b the payments at different resampled bids $\hat{b}$ could be very different. In particular, the magnitude of the payment charged by the single-call mechanism may be much larger than the payments in the original mechanism, i.e. it may be that $|\lambda_i| \gg |P_i|$ for certain outcomes.

Our characterization theorems reveal that there is a fundamental trade-off between expectation and risk. Thus, we call a reduction optimal if it minimizes risk with respect to a lower bound on the expectation.

3.1.1. Expectation. We study three criteria for measuring the expectation of a reduction: $\Pr(\hat{b} = b|b)$, social welfare, and revenue.

The first criterion, $\Pr(\hat{b} = b|b)$ (the precision), measures the likelihood that the reduction modifies players' bids. This criterion is natural when modifying bids is inherently undesirable:

Definition 3.3. The *precision* of a reduction α_P is the probability that the reduction does not alter any player's bid:

$$\alpha_P \equiv \min_b \Pr(\hat{b} = b|b) \ .$$

The other criteria measure standard quantities in mechanism design:

Definition 3.4. The *welfare approximation* α_W of a single-call reduction is given by the worst-case ratio between the welfare of the single-call mechanism and the welfare of the original allocation function:

$$\alpha_W = \min_{A,b} \frac{\mathbf{E}_{\hat{b}}\left[\sum_i b_i(\mathcal{A}_i(b))\right]}{\sum_i b_i(A_i(b))} \ .$$

When the welfare of A is zero, $\alpha_W = 1$ if the welfare of $\mathcal{A}$ is also zero and unbounded otherwise.

Definition 3.5. The *revenue approximation* α_R of a single-call reduction is given by the worst-case ratio between the revenue of the single-call mechanism and the revenue of the original allocation function:

$$\alpha_R = \min_{A,b} \frac{\mathbf{E}_{\hat{b}}\left[\sum_i \mathcal{P}_i(b)\right]}{\sum_i P_i(b)} \ .$$

When the revenue of A is zero, then $\alpha_R = 1$ when the revenue of $\mathcal{A}$ is also zero and unbounded otherwise.

In the case of continuous spaces we replace $\min/\max$ with $\inf/\sup$ as appropriate for infinite domains.

3.1.2. Risk. We measure risk through both the variance of payments and their worst-case magnitude.[5] In order to make a meaningful comparison across different allocation functions and bids, we normalize by players' bids:[6]

Definition 3.6. Decompose λ_i into terms which depend only on the payoff to a single bidder j (i.e. on $b_j(A(\hat{b}))$ instead of $A(\hat{b})$):

$$\lambda_i(A(\hat{b}), \hat{b}, b) = \sum_j \lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)$$

(our characterizations in Sections 4 and 6 show that this is possible for our settings). Then the *bid-normalized payments* of the reduction are given by

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \quad .$$

We can thus write the variance of bid-normalized payments as

$$\max_{A,i} \mathrm{Var}_{\hat{b}\sim\mu_b} \left(\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \right)$$

and the worst-case magnitude as

$$\max_{A,i,\hat{b}} \left| \sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \right|$$

where we replace $\min/\max$ with $\inf/\sup$ as appropriate for infinite domains.

3.1.3. Optimality. We define an optimal reduction as one that *simultaneously* optimizes the six-way trade-off between expectation and risk:

Definition 3.7. A single-call reduction optimizes the *variance of*/*worst-case* payments with respect to *precision*/*welfare*/*revenue* for a set of allocation functions if for every bid b, it minimizes the variance of/worst-case normalized payments over all possible reductions that achieve a precision of α_P / welfare approximation of α_W / revenue approximation of α_R.

4. MAXIMAL-IN-DISTRIBUTIONAL-RANGE REDUCTIONS

In this section, we show how to construct a single-call reduction for MIDR allocation rules, i.e. we show how to construct a randomized, truthful mechanism from an arbitrary MIDR allocation rule A using only a single black-box call to A. The main results are Theorem 4.1, a characterization of all reductions that use VCG payments for an arbitrary MIDR allocation rule, and an explicit construction that optimizes the expectation-risk tradeoff.

[5]Intuition suggests optimizing with respect to a high-probability bound. Unfortunately, this is problematic because ignoring low-probability events can dramatically change the expected payment. Thus, in general it is not reasonable to conclude a priori that low-probability events can be ignored.
[6]Intuition also suggests normalizing by the truthful prices for A (i.e. by P_i), but constant allocation functions such as $A_i(b) = 1$ have $P_i = 0$, making this impossible. Bid-normalized payments are a next logical choice.

Truthful payments for MIDR allocation rules are given by VCG payments with the Clarke-Pivot rule:[7]

$$\mathbf{E}[p_i] \;=\; \mathbf{E}[\text{total welfare of bidders without } i] - \mathbf{E}[\text{total welfare of bidders } j \neq i \text{ with } i] \quad (1)$$

(where the expectation is over the randomization in the given MIDR allocation rule). The reduction comes from this formula for $\mathbf{E}[p_i]$: we need to measure the welfare without agent i (the first term in the RHS), so, with some probability, we ignore agent i and maximize the welfare of the remaining agents. Intuitively, this is equivalent to evaluating the allocation function where i's bid is changed to a "zero" bid while other bids remain the same.

Unfortunately, having removed agent i, even with a small probability, means that computing truthful payments for agent $j \neq i$ requires knowing the allocation where both i and j are ignored. By induction, a single-call mechanism must generate all sets of agents $M \subseteq [n]$ with some probability. Thus, we get an intuitive picture of the reduction's behavior: it will randomly pick a set of bidders $M \subseteq [n]$ and zero the bids of agents not in M.

4.1. Characterizing Truthfulness

We consider reductions in which i's resampled bid $\hat{b}_i$ is always b_i or zero,[8] where "zero" means that the agent has a valuation of zero for all outcomes. That is, the resampling measure $\mu_b(B)$ represents a discrete distribution over the bids $\{\hat{b}^M\}$ where $M \subseteq [n]$ is a set of agents and

$$\hat{b}_i^M = \begin{cases} b_i & i \in M \\ 0 & i \notin M \end{cases}$$

Resampling to $\hat{b}^M$ is equivalent to ignoring the welfare of agents outside M and evaluating A at b.

In the most general setting, our restriction to zeroing reductions is without loss of generality because b and zero are the only bids that are guaranteed to be valid inputs to A for all MIDR allocation functions A. That said, even if a multi-parameter bid structure were known, VCG payments do not depend on the outcome at any other bid. Thus, intuition suggests that resampling to other bids will not be helpful even if it is possible. This intuition can be formalized, but we do not do it here.

Let $\pi(M)$ be a distribution over sets $M \subseteq [n]$. We define the *associated coefficients* $c_i^\pi(M)$ as:

$$c_i^\pi(M) = \begin{cases} -1, & i \in M \\ \frac{\pi(M \cup \{i\})}{\pi(M)}, & i \notin M \end{cases}$$

Intuitively, c_i^π is the weighting that ensures $-\pi(M \cup \{i\}) c_i^\pi(M \cup \{i\}) = \pi(M) c_i^\pi(M)$ (where $i \notin M$) to match the terms in (1).

We prove the following characterization of all truthful MIDR reductions $(\pi, \{\lambda_i\})$ that work for all MIDR A:

[7]If we relax the no positive transfers requirement, a trivial way to construct a single-call mechanism is to ignore the first term in (1). However, the resulting mechanism would make a huge loss because no agent would ever pay the mechanism.

[8]Even if explicit "zero" bids are not known to the reduction, we assume that the reduction can induce A to optimize the utility of an arbitrary subset of agents. Note that a black-box allocation function can only be turned into a truthful mechanism (even if multiple calls to A are allowed) if it can ignore at least one bidder at a time, so our assumption is not unreasonable.

THEOREM 4.1. *A normalized single-call reduction, with VCG payments, for the set of all MIDR allocation rules satisfies truthfulness, individual rationality, and no positive transfers in an ex-post sense if and only if it takes the form $(\pi, \{\lambda_i\})$ where $\pi(M)$ is a distribution over sets $M \subseteq [n]$, the coefficients $c_i^{\pi(M)}$ are finite, and payments take the form*

$$\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = c_i^{\pi}(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) \ .$$

PROOF. Recall that in general, a multi-parameter allocation function that can be rendered truthful by VCG payments must be MIDR. Thus, our reduction must ensure that $\mathcal{A}$ is MIDR — a necessary and sufficient condition for this is that $\mu_b(B)$ can be defined by a distribution $\pi(M)$ where $\pi(M)$ is the probability of selecting $\hat{b} = \hat{b}^M$ (see the full version for details).

Next, we write VCG payments for $\mathcal{A}$ that satisfy individual rationality and no positive transfers using the Clarke-Pivot payment rule:

$$\begin{aligned}\mathbf{E}[\mathcal{P}_i] &= \sum_{j \neq i} \sum_{M \subseteq [n]} \pi(M) b_j(A(\hat{b}^{M \setminus \{i\}})) - \sum_{j \neq i} \sum_{M \subseteq [n]} \pi(M) b_j(A(\hat{b}^M)) \\ &= \sum_{M | i \notin M} \pi(M \cup \{i\}) \sum_{j \neq i} b_j(A(\hat{b}^M)) - \sum_{M | i \in M} \pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) \ . \end{aligned} \tag{2}$$

By definition of $\lambda_i(A(\hat{b}^M), \hat{b}^M, b)$, we know that the expected payment made by i will be

$$\mathbf{E}[\mathcal{P}_i] = \sum_{M \subseteq [n]} \pi(M) \lambda_i(A(\hat{b}^M), \hat{b}^M, b) \ . \tag{3}$$

The two formulas for payments in (2) and (3) must be equal:

$$\sum_{M \subseteq [n]} \pi(M) \lambda_i(A(\hat{b}^M), \hat{b}^M, b) = \sum_{M | i \notin M} \pi(M \cup \{i\}) \sum_{j \neq i} b_j(A(\hat{b}^M)) - \sum_{M | i \in M} \pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) \ .$$

Since A may be any MIDR allocation function, the only way this can hold is when terms corresponding to each M are equal, i.e., for all $i, \ M$

$$\pi(M) \lambda_i(A(\hat{b}^M), \hat{b}^M, b) = \begin{cases} \pi(M \cup \{i\}) \sum_{j \neq i} b_j(A(\hat{b}^M), & i \notin M \\ -\pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) & i \in M \ . \end{cases} \tag{4}$$

To see that this is necessary, construct two allocation functions A and A' such that $b_j(A(\hat{b}^M)) = b_j(A'(\hat{b}^M))$ for all $M \neq \bar{M}$ and $b_j(A(\hat{b}^{\bar{M}})) = 0$. It immediately follows that if the reduction works for both A and A', then (4) must hold for $\bar{M}$ under A. Since $\bar{M}$ is arbitrary, it follows that (4) must hold for all M.

The theorem immediately follows from the above equality. □

4.2. A Single-Call MIDR Reduction

We now give an explicit single-call reduction for MIDR allocation functions. Our reduction `MIDRtoMech`(A, γ) (illustrated in Algorithm 2) is defined by the following resampling distribution $\bar{\pi}$ parameterized by a constant $\gamma \in (0, 1)$:

$$\bar{\pi}(M) = \gamma^{n-|M|}(1-\gamma)^{|M|} \tag{5}$$

That is, each agent i is independently dropped from M with probability γ. Thus sampling from the distribution $\bar{\pi}$ is computationally easy. Following Theorem 4.1, we

ALGORITHM 2: `MIDRtoMech`(A, γ) — A single-call reduction for MIDR allocation functions

input : MIDR allocation function A.
output: Truthful-in-expectation mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$.

Solicit bids b from agents;
for $i \in [n]$ **do**
    with probability $1 - \gamma$
        Add agent i to set M;
    otherwise
        Drop agent i from M;
Realize the outcome $A(\hat{b}^M)$;
Charge payments
$\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = \left(\sum_{j \neq i} b_j(A(\hat{b}^M))\right) \times \begin{cases} -1, & i \in M \\ \frac{1-\gamma}{\gamma}, & i \notin M \end{cases}$;

charge payments $\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = c_i^{\bar{\pi}}(M) \sum_{j \neq i} b_i(A(\hat{b}^M))$ where

$$c_i^{\bar{\pi}}(M) = \begin{cases} -1, & i \in M \\ \frac{1-\gamma}{\gamma}, & i \notin M \end{cases}$$

COROLLARY 4.2 (OF THEOREM 4.1). *The mechanism*

$$\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\}) = \texttt{MIDRtoMech}(A, \gamma)$$

calls A once and it satisfies truthfulness, individual rationality, and no positive transfers in an ex-post sense for all MIDR A.

4.3. Optimal Single-Call MIDR Reductions

We now prove that the construction `MIDRtoMech`(A, γ) is optimal for the definitions of optimality given in Section 3. Theorem 4.1 implies that the bid-normalized payments will be

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} = (n-1)c_i^{\pi}(M)$$

Thus, it is sufficient to optimize the variance as $\max_i \mathrm{Var}_{M \sim \pi} c_i^{\pi}(M)$ and the worst-case as $\max_{i,M} |c_i^{\pi}(M)|$.

4.3.1. Optimizing Risk vs. Precision

THEOREM 4.3. *The reduction* `MIDRtoMech`(A, γ) *uniquely minimizes both the payment variance and the worst-case payment among all reductions that achieve a precision of at least $\alpha_P = (1-\gamma)^n$.*

That is, for any other distribution π with precision $\pi([n]) \geq (1-\gamma)^n$, the payment variance is larger, i.e.

$$\max_i \mathrm{Var}_{M \sim \pi} c_i^{\pi}(M) > \max_i \mathrm{Var}_{M \sim \bar{\pi}} c_i^{\bar{\pi}}(M) \enspace ,$$

and the worst-case payment is larger, i.e.

$$\max_{i,M} |c_i^{\pi}(M)| > \max_{i,M} |c_i^{\bar{\pi}}(M)| \enspace .$$

PROOF. First we prove optimality for the worst-case payment $\max_{i,M} |c_i^{\pi}(M)|$ by contradiction. Assume that some distribution $\pi(M)$ does as well as $\bar{\pi}(M)$. Then it must be that $\max_{i,M} c_i^{\pi}(M) \leq \max_{i,M} c_i^{\bar{\pi}}(M)$ (the largest coefficient is not bigger), and $\pi([n]) \geq \bar{\pi}([n]) = \alpha_P$ (it respects the lower bound on precision). Since $\max c_i^{\bar{\pi}}(M) = \frac{1-\gamma}{\gamma}$, it must be that for all M and $i \notin M$,

$$\frac{\pi(M \cup \{i\})}{\pi(M)} \leq \max_{i,M} c_i^{\bar{\pi}}(M) = \frac{1-\gamma}{\gamma} = \frac{\bar{\pi}(M \cup \{i\})}{\bar{\pi}(M)} \quad .$$

Therefore, for any bidder i, it must be that

$$\frac{\pi([n])}{\pi([n] \setminus \{i\})} \leq \frac{\bar{\pi}([n])}{\bar{\pi}([n] \setminus \{i\})} \quad .$$

Since $\pi([n]) \geq \bar{\pi}([n])$, it follows that $\pi([n] \setminus \{i\}) \geq \bar{\pi}([n] \setminus \{i\})$. Repeating this argument, it follows by induction that $\pi(M) \geq \bar{\pi}(M)$ for any set M.

However, we also know that both $\pi(M)$ and $\bar{\pi}(M)$ are distributions so both have to sum to one over all M. Given that $\pi(M) \geq \bar{\pi}(M)$ for all M, this implies $\pi(M) = \bar{\pi}(M)$. Thus, $\bar{\pi}(M)$ is uniquely optimal.

We prove that $\bar{\pi}$ optimizes the payment variance in the full version. □

4.3.2. Optimizing Risk vs. Welfare. A natural optimization metric is the social welfare of $\mathcal{A}$ (indeed, this was an open question from [Babaioff et al. 2010] in the single-parameter setting).

Unfortunately, since MIDR allocation rules may generate negative utilities and remain MIDR under additive shifts of the valuation function, one can make the welfare approximation arbitrarily bad (indeed, even undefined) by subtracting a constant from each player's valuation. Thus, if valuation functions may be negative, we cannot meaningfully optimize the loss in social welfare.

However, when valuation functions are known to be nonnegative, then the following lemma shows that the worst-case welfare approximation is bounded:

LEMMA 4.4. *The reduction* `MIDRtoMech`(A, γ) *obtains an* $\alpha_W = \min_i \Pr_\pi(i \in M) = 1 - \gamma$ *approximation to the social welfare, and there is an allocation function* A *and bid* b *such that this bound is tight.*

Using this lemma, we can show that `MIDRtoMech`(A, γ) is optimal. (Proofs are given in the full version.)

THEOREM 4.5. *The reduction* `MIDRtoMech`(A, γ) *minimizes payment variance and worst-case payments among all reductions that achieve a welfare approximation of at least* $\alpha_W = 1 - \gamma$.

4.3.3. Optimizing Risk vs. Revenue. The following lemma implies that a lower bound on the factor of approximation to revenue is equivalent to a lower bound on precision.

LEMMA 4.6. *The reduction* `MIDRtoMech`(A, γ) *obtains an* $\alpha_\pi = \pi([n]) = (1-\gamma)^n$ *approximation to the revenue, and this is tight.*

Since Theorem 4.3 says that `MIDRtoMech`(A, γ) optimizes payments with respect to precision, it similarly follows that it optimizes payments with respect to revenue. (Proofs are given in the full version.)

THEOREM 4.7. *The reduction* `MIDRtoMech`(A, γ) *minimizes payment variance and the worst-case payment among all reductions that guarantee an* $\alpha_R = (1-\gamma)^n$ *approximation to revenue.*

5. A SINGLE-CALL APPLICATION — PPC ADAUCTIONS

Pay-per-click (PPC) AdAuctions are a prime example of mechanisms in which uncertainty can destroy truthfulness. There is a deep literature on truthful ad auctions, much of which makes a powerful assumption: the likelihood that a user clicks in any given setting is a commonly-held belief. In reality, this simply is not true. Auctioneers make their best effort to estimate the likelihood of a click; however, anecdotal evidence [Jabas 2010] suggests that advertisers manipulate their bids according to the perceived accuracy of the auctioneer's estimates. As we will illustrate in this section, even if the auctioneer's estimates are good enough to (say) maximize welfare given the current bids, they are not sufficient to compute truthful prices. We show that single-call mechanisms can recover truthfulness in PPC ad auctions in spite of these conflicting beliefs.

In a standard PPC ad auction, n advertisers compete for $m \ll n$ slots. The value to an advertiser depends on the likelihood of a click, called the click-through-rate (CTR) c, and the value to the advertiser once the user has clicked, the value-per-click v. The expected value to an advertiser is thus cv. The auctioneer's job is to assign advertisers to slots and compute per-click payments — bidders are only charged when a click occurs. Both tasks require knowing the CTRs for common objectives like welfare or revenue maximization, so the auctioneer must also maintain estimates of the CTRs, which we denote by c'.

Researchers generally acknowledge that, in reality, both c and v may depend arbitrarily on the outcome — they certainly depend on the quality and relevance of the particular ad being shown, but they also depend on where the ad is shown and on which other ads are shown nearby. However, for analytical tractability, the parameters c and v are often assumed to have a very restricted structure. We discuss two different structures to illustrate the pervasiveness of the problem caused by estimation error and to show how different single-call reductions may be applied.

Outcome-Independent Values and Separable CTRs. In the ad auction literature, it is common to assume that a bidder's value-per-click v_i is independent of the assignment and that the CTR is separable, that is, it takes the form $c = \alpha_j \beta_i$, where β_i depends only on the ad and α_j depends only on the slot $j \in [m]$ where the ad is shown. Unfortunately, even in this restricted setting, estimation errors may break the truthfulness of VCG prices (even in the special case where $\beta_i = 1$).[9]

In the language of allocations and payments, truthfulness is broken because the auctioneer only knows an estimate of A and thus does not have enough information to compute true VCG prices. However, once ads are shown, clicks may be measured, giving an unbiased estimate of bidders' values. Unfortunately, this can only be done once — since the auctioneer only has one opportunity to show ads to the user, these unbiased estimates can only be measured under a single advertiser-slot assignment. Fortunately, *these unbiased estimates are exactly the information required to compute truthful payments using a single-call mechanism.*

Since a player's bid b_i is merely its value-per-click v_i, this version of a PPC ad auction is a single-parameter domain and we can apply the result of [Babaioff et al. 2010]. Their result says that we can turn any monotone allocation rule into a truthful-in-expectation mechanism — maximizing welfare subject to estimates α'_j and β'_i is a

[9]For example, consider a two bidder setting where $v_1 = 1.1$ and $v_2 = 1$. Let $\alpha_1 = 0.1$, $\alpha_2 = 0.09$, and $\beta_1 = \beta_2 = 1$. Assume the auctioneer knows all parameters correctly except α_1, which he estimates $\hat{\alpha}_1 = 0.11$. It is easy to check that the first bidder would rather say $b_1 = 0$ than $b_1 = v_1 = 1.1$ in a standard VCG pay-per-click auction. Thus, a small estimation error in a single parameter can incentivise bidders to radically misstate their bids, even though the estimates are good enough to guarantee the welfare-maximizing allocation. This example is covered in detail in the full paper.

monotone allocation rule as long as the estimates α'_j have the same order as α_j (i.e. $\alpha'_{j_1} \geq \alpha'_{j_2}$ if $\alpha_{j_1} \geq \alpha_{j_2}$), so [Babaioff et al. 2010] gives a truthful mechanism for almost any estimates:

THEOREM 5.1. *Consider a single-parameter PPC auction with separable CTRs and let A^{PPC} be the allocation rule that maximizes welfare using estimated CTR parameters α'_j and β'_i, where the estimates α'_j are properly ordered. Then* `SPtoMechBKS`(A^{PPC}, γ)*, the single-call reduction of [Babaioff et al. 2010], gives a mechanism that is truthful in expectation and has expected welfare within a factor of $(1-\gamma)^n$ of A^{PPC}.*

Outcome-Dependent Values and CTRs. While most research uses single-parameter models for analytical tractability, an advertiser's value-per-click v really depends on the advertiser-slot assignment chosen by the auctioneer as noted earlier. As in the preceding single-parameter setting, estimated CTRs are insufficient to guarantee truthfulness; however, the reduction of [Babaioff et al. 2010] no-longer applies in such a multi-parameter domain — we show how our MIDR single-call reduction can be used to recover truthfulness.

To capture the dependence on the advertiser-slot assignment, we assume that a bidder's CTR $c_{i,j}$ and value-per-click $v_{i,j}$ depend arbitrarily on both the bidder i and the slot j. Since the only allocation rules that have truthful prices in general multi-parameter domains are MIDR, we assume that the auctioneer can generate a MIDR allocation, specifically we assume the auctioneer can query an oracle to determine the allocation that maximizes the welfare of any set of bidders under the actual bid b (but not necessarily for an arbitrary bid b) and apply our MIDR reduction:

THEOREM 5.2. *Consider a multi-parameter PPC auction where a bidder's value-per-click $v_{i,j}$ depends on the bidder and the slot. Let A^{PPC} be an allocation rule that chooses the advertiser-slot assignment returned by the welfare-maximizing oracle described above. Then the mechanism* `MIDRtoMech`(A^{PPC}, γ) *is truthful in expectation and approximates the welfare of A^{PPC} to within a factor of $(1-\gamma)$.*

6. SINGLE-PARAMETER REDUCTIONS

In this section, we characterize truthful reductions for single-parameter domains and show that the construction of [Babaioff et al. 2010] is optimal. Theorem 6.1 characterizes all reductions that are truthful for an arbitrary monotone, bounded, single-parameter allocation function A. Our characterization is more general than the self resampling procedures described by Babaioff et al. and shows that a wide variety of probability measures may be used to construct a truthful reduction. Theorem 6.3 shows that the construction given in Babaioff et al. is optimal among such reductions for a fixed bound on the precision, welfare approximation, or revenue approximation of the reduction.

6.1. Characterizing Single-Call Reductions

For the sake of intuition, we start with the special case that the resampling measure μ_b has a nicely behaved density representation $f_b(\hat{b})$ (the resampling density) that is continuous in $\hat{b}$ and b. The proof for arbitrary measures μ_b requires significant measure theory and is given in the full version.

Define the coefficients $c_i^f(\hat{b}, b)$ as $c_i^f(\hat{b}, b) = 1 - \frac{1}{b_i}\int_0^{b_i} \frac{f_{u,b_{-i}}(\hat{b})}{f_b(\hat{b})} du$ when $b_i \neq 0$, and to be 0 when $b_i = 0$. We characterize truthful reductions as follows:

THEOREM 6.1. *A normalized single-parameter reduction $(f, \{\lambda_i\})$ for the set of all monotone bounded single-parameter allocation functions satisfies truthfulness, individ-*

ual rationality and no positive transfers in an ex-post sense if and only if the following conditions are met:

(1) *The resampling density f_b is such that the single-call mechanism's randomized allocation procedure $\mathcal{A}_i(b)$ is monotone in expectation, i.e., for all agents i, for all b, and $b'_i \geq b_i$, $\mathbf{E}_{\hat{b}\sim f_b}[\mathcal{A}_i(b'_i, b_{-i})] \geq \mathbf{E}_{\hat{b}\sim f_b}[\mathcal{A}_i(b)]$. (See below.)*

(2) *The resampling density f_b is such that $f_b(\hat{b}) \neq 0$ if $\int_0^{b_i} f_{u,b_{-i}}(\hat{b})du \neq 0$.*[10]

(3) *The payment functions $\lambda_i(A(\hat{b}), \hat{b}, b)$ satisfy: $\lambda_i(A(\hat{b}), \hat{b}, b) = b_i c_i^f(\hat{b}, b) A_i(\hat{b})$ almost surely, i.e. for all $\hat{b}$ except possibly a set with probability zero under f_b.*

PROOF. (Sketch of necessity. See the full version for a complete proof of both necessity and sufficiency.) The proof is analogous to Theorem 4.1, using the Archer-Tardos payments in place of VCG ones. The Archer-Tardos characterization says that truthful payments exist for a single-parameter allocation function if and only if it is monotone, hence the requirement that $\mathcal{A}$ be monotone in expectation.

Truthful payments (if they exist) must be given by the Archer-Tardos formula $\mathbf{E}[\mathcal{P}_i] = b_i \mathbf{E}_{\hat{b}\sim f_b}[\mathcal{A}_i(b)] - \int_0^{b_i} \mathbf{E}_{\hat{b}\sim f_{u,b_{-i}}}[\mathcal{A}_i(u, b_{-i})]du$. Additionally, we can express the expected price as $\mathbf{E}[\mathcal{P}_i] = \int_{\hat{b}\in\mathbb{R}^n} f_b(\hat{b})\lambda_i(A(\hat{b}), \hat{b}, b)d\hat{b}$ by construction. Equating these two formulae for $\mathbf{E}[\mathcal{P}_i]$ as in Theorem 4.1 and rearranging the Archer-Tardos payment formula gives

$$\int_{\hat{b}\in\mathbb{R}^n} f_b(\hat{b})\lambda_i(A(\hat{b}), \hat{b}, b)d\hat{b} = \int_{\hat{b}\in\mathbb{R}^n} f_b(\hat{b}) b_i A_i(\hat{b}) \left(1 - \frac{1}{b_i}\int_0^{b_i} \frac{f_{u,b_{-i}}(\hat{b})}{f_b(\hat{b})}du\right) d\hat{b} \ . \qquad (6)$$

To prove necessity, we must show that the integrands of (6) are equal almost surely. If we had that (6) is required to hold for arbitrary functions A, it is intuitive that the only way to guarantee (6) is to match the integrands; however, we only require (6) to hold for monotone A. In the full version, we show that this is still enough to establish that the integrands are equal almost surely. □

Unfortunately, our assumption that μ_b has a density representation is unreasonable. Most significantly, one would expect $\hat{b} = b$ with some nonzero probability, implying that μ_b would have at least one atom for most interesting distributions. In particular, the distribution used in the BKS transformation has such an atom, so it cannot be analyzed in this fashion.

To handle general measures μ_b we apply the same ideas using tools from measure theory. See the full version for details.

6.2. The BKS Reduction for Positive Types

The central construction of Babaioff, Kleinberg, and Slivkins [2010] is a reduction for scenarios where bidders have positive types.[11]

Their resampling procedure (implicitly defining μ_b) is described Algorithm 3. In the language of our characterization, the coefficients c_i^{BKS} are

$$c_i^{BKS}(\hat{b}, b) = \begin{cases} 1, & \hat{b}_i = b_i \\ 1 - \frac{1}{\gamma} & \textit{otherwise.} \end{cases}$$

[10]This condition effectively requires $c_i^f(\hat{b}, b)$ to be finite.

[11]They also give a reduction that applies to more general type spaces, but we do not state it here.

ALGORITHM 3: `SPtoMechBKS`(A, γ) — The BKS reduction for single-parameter domains

input : Bounded, monotone allocation function A.
output: Truthful-in-expectation mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$.

Solicit bids b from agents;
for $i \in [n]$ **do**
  with probability $1 - \gamma$
    Set $\hat{b}_i = b_i$;
  otherwise
    Sample x_i uniformly at random from $[0, \hat{b}_i]$;
    Set $\hat{b}_i = b_i x_i^{\frac{1}{1-\gamma}}$;
Realize the outcome $A(\hat{b})$;
Charge payments
$\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = b_i A_i(\hat{b}) \times \begin{cases} 1, & \hat{b}_i = b_i \\ \frac{1-\gamma}{\gamma}, & \hat{b}_i < b_i \end{cases}$;

They proved that `SPtoMechBKS`(A, γ) is truthful. This fact can be easily derived from Theorem 6.1:

THEOREM 6.2 (BABAIOFF, KLEINBERG, AND SLIVKINS 2010.). *For all monotone, bounded, single-parameter allocation rules A, the single-call mechanism given by* `SPtoMechBKS`(A, γ) *satisfies truthfulness and no positive transfers in an ex-post sense and is ex-post universally individually rational.*

6.3. Optimal Single-Call Reductions

Analogous to our MIDR construction, we show that, the BKS construction for positive types is optimal with respect to precision, welfare, and revenue as defined in Section 3. Using our characterization from Theorem 6.1, the bid-normalized payments we wish to optimize will be

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} = \frac{c_i^{\mu}(\hat{b}, b) b_i A_i(\hat{b})}{b_i A_i(\hat{b})} = c_i^{\mu}(\hat{b}, b) \ .$$

Thus, optimizing variance of normalized payments is equivalent to optimizing $\max_i \mathrm{Var}_{\hat{b} \sim \mu_b} c_i^{\mu}(\hat{b}, b)$, and optimizing the worst-case normalized payment is equivalent to optimizing $\sup_{i, \hat{b}} |c_i^{\mu}(\hat{b}, b)|$.

For this section, we make a "nice distribution" assumption that for any $u \neq b_i$, $\Pr(\hat{b}_i = u | b) = 0$. That is, if we compute the marginal distribution of $\hat{b}_i$, the only bid $\hat{b}_i$ that has an atom is b_i (other bids only have positive density). We handle the general case in the full version.

Our main result is that the BKS transformation is optimal:

THEOREM 6.3. *The single-call reduction* `SPtoMechBKS`(A, γ) *optimizes the variance of bid-normalized payments and the worst-case bid-normalized payment for every b subject to a lower bound $\alpha = (1 - \gamma)^n \in (\frac{1}{e}, 1)$ on the precision, the welfare approximation, or the revenue approximation.*

We prove Theorem 6.3 by showing that the three risk metrics are equivalent for interesting reductions in this setting. Precise statements and proofs are given in the full version.

7. ACKNOWLEDGMENTS

First, we want to thank Kamal Jain for showing us the issues with pay-per-click advertising auctions and for suggesting [Babaioff et al. 2010] as a possible solution.

Second, we would like to thank many people who have provided invaluable feedback and suggestions. In particular, we would like to thank Alex Slivkins for his substantial help in the revising process, Robert Kleinberg for his technical suggestions, and Christos Papadimitriou for his helpful comments.

REFERENCES

ARCHER, A. AND TARDOS, E. 2001. Truthful mechanisms for one-parameter agents. In *FOCS '01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. IEEE Computer Society, Washington, DC, USA, 482.

BABAIOFF, M., BLUMROSEN, L., NAOR, M., AND SCHAPIRA, M. 2008. Informational overhead of incentive compatibility. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 88–97.

BABAIOFF, M., KLEINBERG, R. D., AND SLIVKINS, A. 2010. Truthful mechanisms with implicit payment computation. In *EC '10: Proceedings of the 11th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 43–52.

BABAIOFF, M., SHARMA, Y., AND SLIVKINS, A. 2009. Characterizing truthful multi-armed bandit mechanisms: extended abstract. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*. ACM, New York, NY, USA, 79–88.

DEVANUR, N. R. AND KAKADE, S. M. 2009. The price of truthfulness for pay-per-click auctions. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*. ACM, New York, NY, USA, 99–106.

DOBZINSKI, S. AND DUGHMI, S. 2009. On the power of randomization in algorithmic mechanism design. In *FOCS*. 505–514.

DUGHMI, S. AND ROUGHGARDEN, T. 2010. Black-box randomized reductions in algorithmic mechanism design. In *Proceedings of the 51st IEEE symposium on Foundations of Computer Science*. 775–784.

HARTLINE, J. D. 2011. Approximation in economic design. Draft.

HERSHBERGER, J. AND SURI, S. 2001. Vickrey prices and shortest paths: What is an edge worth? In *Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. FOCS '01. IEEE Computer Society, Washington, DC, USA.

HERSHBERGER, J. AND SURI, S. 2002. Erratum to "Vickrey pricing and shortest paths: What is an edge worth?". In *Proceedings of the 43rd Symposium on Foundations of Computer Science*. FOCS '02. IEEE Computer Society, Washington, DC, USA.

HERSHBERGER, J., SURI, S., AND BHOSLE, A. 2007. On the difficulty of some shortest path problems. *ACM Trans. Algorithms 3*, 5:1–5:15.

JABAS, D. 2010. Private communication.

LAHAIE, S. 2010. Stability and incentive compatibility in a kernel-based combinatorial auction. In *AAAI*.

MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research 6,* 1, 58–73.

NISAN, N. AND RONEN, A. 2001. Algorithmic mechanism design. *Games and Economic Behavior 35,* 1-2, 166–196.

NISAN, N. AND RONEN, A. 2007. Computationally feasible VCG mechanisms. *J. Artif. Intell. Res. (JAIR) 29*, 19–47.

ROBERTS, K. 1979. The characterization of implementable social choice rules. In *Aggretaion and Revelation of Preferences*, J-J.Laffont (ed.), North Holland Publishing Company.