Analysis of Medicare Pay-for-Performance Contracts

Hamsa Bastani
Stanford University, Electrical Engineering; hsridhar@stanford.edu

Mohsen Bayati
Stanford University, Graduate School of Business; bayati@stanford.edu

Mark Braverman
Princeton University, Computer Science; mbraverm@cs.princeton.edu

Ramki Gummadi
Stanford University, Graduate School of Business; gsrk84@gmail.com

Ramesh Johari
Stanford University, Management Science & Engineering; rjohari@stanford.edu

Medicare has sought to improve patient care through pay-for-performance (P4P) programs that better align hospitals' financial incentives with quality of service. However, the design of these policies is subject to a variety of practical and institutional constraints, such as the use of "small" performance-based incentives. We develop a framework based on a stylized principal-agent model to characterize the optimal P4P mechanism within any set of feasible mechanisms in the regime of small incentives. Importantly, our feasible set can be flexibly modified to include institutional constraints. We apply our results to examine debated design choices in existing Medicare P4P programs, and offer several insights and policy recommendations. In particular, we find that these mechanisms may benefit by incorporating bonuses for top-performers, and using a single performance cutoff to uniformly assess performance-based payments. We also examine a number of comparative statics that shed light on when P4P mechanisms are effective.

Key words: pay-for-performance, Medicare, mechanism design, contracts, principal-agent problems, information asymmetry, healthcare

1. Introduction

> *"Despite the growing use of [pay-for-performance] initiatives, there is little evidence on how best to design incentive programs in the health sector. Perhaps as a result of the paucity of evidence, there is tremendous variety in the approaches used in existing incentive programs."* - Agency of Healthcare Quality & Research (Rosenthal and Dudley 2006)

Aggregate healthcare spending in the US is over \$2.5 trillion a year, which constitutes 17% of the GDP, since 2010. These numbers represent the highest per-capita spending in the world. Government healthcare spending alone currently stands at \$1.2 trillion, and is expected to exceed \$2.3 trillion by 2019 (Truffer et al. 2010). Yet, there are serious concerns about the quality of care, including overuse (patients receiving unnecessary or harmful procedures) and underuse (patients not receiving care recommended by scientific evidence) of healthcare resources (Davis et al. 2010, Green 2012). One primary cause cited for this dilemma is distorted incentives for healthcare providers created by misaligned payment systems (Porter and Teisberg 2006). In particular, traditional payment models do not account for treatment outcomes or the degree to which the treatment is appropriate for the patient's condition (Institute of Medicine 2006). This leads to concerns over moral hazard, where providers may over-treat patients with unnecessary, profitable procedures and fail to provide necessary but less profitable treatments.

Consequently, there has been wide interest among healthcare payers to incorporate pay-for-performance (P4P) programs as an incentive to improve health care quality. These programs define a real-valued performance metric (based on treatment costs and/or patient outcomes), and levy a financial bonus or penalty from providers as a function of this metric. More than half of commercial health maintenance organizations are using P4P (Rosenthal and Dudley 2007), and the US Department of Health and Human Services has committed to tying 90 percent of all traditional Medicare payments to quality or value by 2018 (HHS 2015).

In this paper, we focus on Medicare's current P4P initiatives, which include (1) the Hospital Readmissions Reduction Program (HRRP), (2) the Hospital-Acquired Condition Reduction Program (HACRP), and (3) the Value Based Purchasing (VBP). These programs modify traditional fee-for-service payments to providers by adding a very small performance-based incentive (depending on the program, the performance-based payment is capped at 1-3% of the base Medicare reimbursement amount). The early results of these programs appear promising: readmissions rates have declined significantly across all targeted conditions (Suter et al. 2014), patient safety indicators have improved (CMS 2015a), and surveys suggests that the P4P policies have had a major impact on providers' efforts to improve along targeted performance metrics (Joynt et al. 2016).

On the other hand, despite the growing use of P4P initiatives, there is little practical guidance on how best to design these programs, which has resulted in a tremendous variety of approaches in practice (Rosenthal and Dudley 2006). In particular, the structure of these programs arises from discussions with various stakeholders, and is subject to a variety of institutional constraints; thus, they often depart from optimal designs advocated by economic theory (Dudley et al. 2004). For example, one widely-criticized issue is the use of a relatively small performance-based payment. The medical community has put forth arguments that larger performance-based payments may be substantially more effective since they can subsidize the provider's costs of improved care (Rosenthal and Dudley 2007), and may be more likely to alter treatment choices (and thereby costs or patient outcomes) (Jha 2013). Researchers in operations management have used a combination of game theory and empirical estimation to conclude that optimal contracts would indeed constitute larger performance-based payments (Aswani et al. 2016) in order to incentivize more providers (Zhang et al. 2016). Yet, Medicare has steadfastly maintained small incentives for all its P4P programs, possibly due to concerns that large performance-based penalties may lead to hospital bankruptcies, which may in turn worsen patient care (e.g., see Maizel and Garner 2014). Thus, arguably, the use of small performance-based payments is an important institutional constraint on any P4P program that Medicare is willing to consider. Unfortunately, there is a mismatch between the significant literature on optimal contracts, and the practical constraints of Medicare P4P implementation.

In this paper, we develop a framework based on a stylized principal-agent model to characterize the optimal P4P mechanism within any set of feasible mechanisms in the regime of small incentives. Importantly, our feasible set can be modified to consider institutional constraints; in particular, this allows us to apply our results to examine Medicare's design choices in existing P4P programs, yielding several insights and policy recommendations.

1.1. Summary of Contributions

We first formalize Medicare's problem using a principal-agent model with an information asymmetry characterized by Medicare's imperfect observation of the patient's health status and the resulting treatment options. We use this setting to study the efficiency of potential mechanisms.

Small Incentives. Our main result characterizes the optimal mechanism within any feasible set under a stylized setting (with two treatments and uncorrelated principal-agent utilities) when the performance-based incentive is small. To illustrate the utility of this result, we analyze the following two issues that are debated in P4P design (Rosenthal and Dudley 2006).

Penalties vs. Bonuses. Currently, all three Medicare mechanisms assess penalties for poor-performers, but only VBP pays bonuses for top-performers. The addition of a bonus is a debated issue: while it can further motivate providers to improve performance metrics, it also increases the overall cost of the mechanism. Using the HRRP mechanism as a case study, our model suggests that incorporating a bonus may improve Medicare's expected utility. In particular, since the incentive size is small, the benefit realized from incentivizing a broader spectrum of providers (to improve their treatment choice) is likely to outweigh the cost of the bonus. Extensive numerical experiments suggest that our insight is robust to changes in our modeling assumptions.

Single vs. Continuous Performance Threshold. Currently, the HACRP mechanism uses a single performance threshold (the mechanism rewards/penalizes all providers above/below a fixed performance threshold equally), while the other two Medicare mechanisms utilize continuous performance thresholds (the size of the bonus/penalty scales in magnitude with the provider's performance). A single threshold is desirable from a behavioral perspective since it sets a simple performance benchmark for providers to achieve; on the other hand, intuition suggests that a continuous performance threshold may be more effective overall since it incentivizes a wider range of providers who may find the single threshold too hard or easy (Rosenthal and Dudley 2006). Surprisingly, using the HACRP mechanism as a case study, our model illustrates that a single threshold is always more effective for small incentives, since it does not further dilute the incentive. Again, extensive numerical experiments demonstrate that this insight is robust.

Comparative Statics: Correlations. Next, we study the addition of correlations between the principal's and agent's utilities, as well as changes in the number of available treatments. For the purpose of gaining qualitative insights, we consider the class of feasible mechanisms to be all linear mechanisms, which allows us to analytically express the optimal mechanism. This analysis yields several insights about the conditions under which a P4P mechanism can be effective.

First, we find that a P4P mechanism is *always* beneficial if the agent's utilities are negatively correlated with the principal's utilities, necessitating the incentive and justifying its cost. This motivates the widespread use of P4P since there is significant evidence for the misalignment of incentives in healthcare settings (Green 2012). Second, P4P mechanisms are valuable when the variation in the provider's private preferences between treatments is small compared to the variation in the performance metric (e.g., readmissions rates). This situation is germane in healthcare settings where the outcomes for patients could vary significantly, while the provider's private utilities corresponding to different treatment options may only vary little in comparison (Cutler 2004). Finally, we note that the number of available treatments does not affect the efficiency of a mechanism.

1.2. Related Literature

The principal-agent framework has a rich history (Bolton and Dewatripont 2004, Holmstrom and Milgrom 1987, Hölmstrom 1979, Laffont and Martimort 2001, Mirrlees 1999) outside the context of healthcare applications. As motivated in Arrow (1963), one of the distinguishing features in the healthcare setting is the nature of information asymmetry between the principal and the agent. Asymmetric information could refer to either the action of the agent not being observable (contractible) or alternately, the agent having access to more information than the principal in deciding the action. Our model incorporates both these forms of information asymmetry. The uncertainty about patient risk factors (and thus, the utility of potential treatments) that are only observable to the agent translates to the uncertainty between the agent's effort and production levels in classical principal-agent terminology. An alternate view can also map this setting to one of imperfect observation of the agent's actions.

Our work is related to the literature on contract design in the healthcare setting. For example, Yaesoubi and Roberts (2011) consider a reimbursement problem involving a binary decision on a preventive treatment that incorporates both moral hazard as well as asymmetric information. When patients' risks are rank-ordered, the authors study the design of optimal reimbursement rates under assumptions involving the observability of the total number of patient visits. An alternate incentivization problem in healthcare markets is studied by Yamin and Gavious (2013), who take the perspective of incentivizing the patient. Bundorf et al. (2008) consider the effect of insurance policy pricing on the welfare of the market through consumer (patient) responses. In contrast, we focus exclusively on the strategic behavior of the provider, which is exercised in the form of choosing between available treatments based on observed patient covariates.

We note that modifying payments is not the only proposed technique to align agents' incentives with improved outcomes. Ma and Mak (2012) propose a mechanism to publicize each provider's rating (as a combined measure of its cost and quality) to drive consumer demand for services. In this model, the patient's action is relevant through the demand function, and the social planner serves to transfer information to the patient via ratings. Other work has also studied the effects of organizational structure on influencing provider behavior by focusing on emergency physicians and their autonomy in assigning patients (Chan 2012). Fong (2009) consider mechanisms that involve scoring providers to solve the adverse selection problem (providers selectively avoid high-risk patients using private information on patients) simultaneously with the moral hazard problem (choosing sub-optimal treatments). We note that we do not directly model adverse selection, but it is possible to utilize a broader interpretation of provider actions in our setting that corresponds to a selection of both patient type as well as the action. Within the broader principal-agent framework, Frankel (2010) reviews models that use a similar form of uncertainty about the agent's menu of decisions, but focus on limiting the agent's action space rather than modifying payments.

Our work is most closely related to a recent stream of papers in operations management studying Medicare's new initiatives. For example, Adida et al. (2016), Gupta and Mehrotra (2015) and Guo et al. (2016) study the bundled payments initiative, a new model where providers receive a single payment for an episode of care. Adida et al. (2016) and Guo et al. (2016) compare the bundled payment model against traditional fee-for-service, while Gupta and Mehrotra (2015) study how such bundled payments should be priced. Similarly, Andritsos and Tang (2015) analyze the readmissions program HRRP and compare the effect of different payment mechanisms on readmissions rates, but they do not study the information asymmetry between Medicare and the provider. Zhang et al. (2016) and Aswani et al. (2016) provide empirical and theoretical evidence that the current Medicare incentives are too small (for HRRP and the Shared Savings Program respectively) and suggested increasing the size of the incentive. In contrast, we take the constraint of "small incentives" as a given, and study the design of efficient mechanisms in this regime. Notably, Fuloria and Zenios (2001) and Jiang et al. (2012) study performance-based contracts in a principal-agent model for outpatient settings; however, their optimal contract designs do not address the issue of institutional constraints, which is the focus of our study.

1.3. Structure of the Paper

The remaining paper is organized as follows. Section 2 describes our principal-agent formulation and the P4P mechanism optimization problem. We characterize optimal mechanisms in the small incentive regime in Section 3. We apply our result to current Medicare programs to derive qualitative insights about HRRP (Section 4) and HACRP (Section 5). We examine comparative statics in a more general setting (with correlated principal-agent utilities) to illustrate when P4P policies are valuable in Section 6. Finally, we discuss potential future research and conclude in Section 7.

2. Model

We study a principal-agent model where the principal represents Medicare (or alternatively, any healthcare insurer), and agents represent providers such as hospitals, clinics, or physicians.

2.1. Setup

Consider a single patient who arrives for treatment at a randomly-chosen provider (we will show that this model generalizes straightfowardly to multiple patient-provider pairs). There are n fixed treatment choices. The patient is associated with patient-specific risk factors that determine both the baseline patient outcome and the utility of each treatment choice. In particular, a treatment may be appropriate for some patients but not others.

Treatments and their patient-specific effects are known to the agent, but may or may not be known to Medicare. In particular, the effects of medical treatments like surgeries or drugs may be known to Medicare through clinical trials and past studies. The effects of operational interventions (e.g., extending a patient's length of stay, scheduling a post-discharge phone call, creating checklists or guidelines for staff, etc.) may not be known to Medicare and are not reported in claims data; these are often the kinds of interventions considered for improving provider metrics under Medicare's P4P policies (e.g., see Berenholtz et al. 2004). We consider both settings in our results. However, in either case, Medicare cannot determine the optimal treatment for a patient since they do not observe all the relevant patient-specific risk factors. While some risk factors are observable to both Medicare and the provider (e.g., demographic information such as age or gender, and claims data from past visits such as previously reported diagnoses), other risk factors are observable only to the provider (e.g., electronic medical record data such as lab tests, and information from the current visit such as symptoms and diagnosis). We formalize this information asymmetry as follows:

The patient is associated with a vector of patient-level covariates $\bar{\boldsymbol{x}} = [\boldsymbol{x}_0, \boldsymbol{x}]$, where $\boldsymbol{x}_0 \in \mathbb{R}^{d_0}$ and $\boldsymbol{x} \in \mathbb{R}^d$. The principal only observes partial information $\boldsymbol{x}_0$ due to the information asymmetry[1], while the agent observes both $\boldsymbol{x}_0$ and $\boldsymbol{x}$. The principal can estimate a patient's baseline outcome

[1] In reality, Medicare may have access to some patient information from claims histories that the provider may lack for a new patient. However, we assume that the provider gathers all treatment-relevant information for the patient.

based on the observed patient covariates $\boldsymbol{x}_0$ by applying a known "risk-adjustment" vector $\beta_0 \in \mathbb{R}^{d_0}$. This models Medicare's risk-adjustment procedure, which adjusts the outcome of interest for differences in observed patient risk factors before computing provider performance metrics. β_0 is typically estimated from aggregated data on past patient outcomes.

The agent can modify the patient's outcome using one of n potential treatment choices. Let the notation $[k]$ denote the set $\{1,...,k\}$. Each treatment is associated with (a possibly known or unknown) parameter vector $\beta_i \in \mathbb{R}^d$ for each $i \in [n]$, where we normalize $\|\beta_i\|_2 = 1$. (If the treatment parameters $\{\beta_i\}$ are unknown to Medicare, we consider the standard uniform prior over the d-dimensional unit sphere S^{d-1}.) The principal's utility for the patient under treatment i (with respect to any performance metric, e.g., readmissions or mortality rates) is

$$u_i^P = \boldsymbol{x}_0 \cdot \beta_0 + \boldsymbol{x} \cdot \beta_i \,,$$

for each $i \in [n]$. This model allows us to capture patient-specific treatment outcomes. For analytical tractability, we assume that the unobserved patient covariates $\boldsymbol{x}$ are i.i.d. random variables from a multivariate normal distribution $\mathcal{N}(0^d, \sigma^2 I)$, but we relax this assumption in our numerical results to show that our insights are robust to the choice of distribution. (We assume the observed patient covariates $\boldsymbol{x}_0$ are also i.i.d. random variables, but we do not make any assumptions on its distribution, except that it has finite mean.)

The agent (who is randomly chosen by the patient) may be partially interested in the principal's utility (captured by a *correlation parameter* $\rho \in [-1,1]$ below). The resulting correlation in principal-agent utilities may be positive due to altruistic concerns for the patient's welfare or downstream reputation effects (for instance, Medicare now publishes selected provider-specific performance metrics); it may also be negative, since increased healthcare spending by the agent may result in improved patient outcomes. The agent also has private treatment-specific preferences that we capture through a vector of random shocks $a \sim \mathcal{N}(0^n, \tau^2 I)$; these preferences may include treatment profits as well as the time and effort costs spent on administering a treatment at a particular healthcare facility. Thus, the utility of performing treatment i for agent t is

$$u_i^A = \rho \cdot (\boldsymbol{x}_0 \cdot \beta_0 + \boldsymbol{x} \cdot \beta_i) + a_i \,.$$

Naturally, the agent chooses the treatment $i^* = \arg\max_{i \in [n]} u_i^A$ that maximizes her overall utility. After treatment, the principal observes the resulting performance $u_{i^*}^P$; however, she never observes the unknown risk factors $\boldsymbol{x}$ or the agent's utility $u_{i^*}^A$. The principal achieves expected utility

$$u^P = \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[u_{i^*}^P\right] \,,$$

Multiple Patients. Our model straightforwardly generalizes to a setting with T patients, each of whom independently selects a random provider. Let the n treatments be fixed across all provider-patient pairs, and index patient covariates as well as provider preferences by $t \in [T]$. At each time t, the agent chooses $i_t^* = \max_{i \in [n]} u_{t,i}^A$. Then, the principal achieves overall expected utility

$$u^P = \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\sum_{t=1}^{T} u_{t,i_t^*}^P\right] = T \cdot \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[u_{i^*}^P\right] ,$$

where we note that the principal's expected utility for each patient t is independent of that of other patients (since the patient covariates $\bar{\boldsymbol{x}}_t$ and agent preferences a_t are drawn i.i.d. at each time t). Thus, it suffices to maximize the principal's expected utility for a single patient-provider pair.

Pay-for-performance. The principal's expected utility may suffer under the baseline mechanism described above if agents choose treatments that are substantially worse than first-best for the chosen performance metric, i.e., $\mathbb{E}\left[u_{i^*}^P\right] \ll \mathbb{E}\left[\max_{j \in [n]} u_j^P\right]$. For example, this may occur if the correlation parameter ρ is negative; in this case, adding a performance-based payment would alleviate the negative correlation in principal-agent utilities and thus, the agent may choose a treatment with better expected performance.

We consider a performance-based payment that is a function of the principal's risk-adjusted utility from a chosen treatment. Note that the risk-adjusted principal's utility is

$$u_i^P - \boldsymbol{x}_0 \cdot \beta_0 \;=\; \boldsymbol{x} \cdot \beta_i \, ,$$

which adjusts for the baseline expected outcome computed from observed patient covariates.

2.2. Mechanism Optimization

Consider a general mechanism where the principal offers the agent a performance-based incentive $\phi(z\,;\xi)$, which is an increasing function of the performance metric z and is parametrized in magnitude by ξ, i.e., the incentive size $|\phi(z\,;\xi)|$ is increasing in ξ for all z. We define $\phi(z\,;0) = 0$ to be the baseline setting with no performance-based incentive and consider values of $\xi \geq 0$.

We impose two mild regularity conditions on ϕ. First, we assume ϕ has a continuous and bounded second derivative with respect to the magnitude parameter ξ (but not necessarily the performance metric z). Second, we assume $|\phi|$ grows at most polynomially in z, i.e., for any fixed ξ, there exist constants C, k such that $|\phi(z\,;\xi)| \leq C \cdot z^k$ for all z; this ensures that the expectation of the principal's utility over the performance z is finite. We do not make any other functional assumptions on ϕ; this flexible parametrization helps us analyze Medicare's current policies, which are typically not smooth in the performance metric z (see §4-5).

We then have the following modified utility functions for the principal and the agent respectively:

$$\tilde{u}_i^P(\phi) = \boldsymbol{x}_0 \cdot \beta_0 + \boldsymbol{x} \cdot \beta_i - \phi(\boldsymbol{x} \cdot \beta_i \ ; \xi) \ ,$$
$$\tilde{u}_i^A(\phi) = \rho \cdot (\boldsymbol{x}_0 \cdot \beta_0 + \boldsymbol{x} \cdot \beta_i) + a_i + \phi(\boldsymbol{x} \cdot \beta_i \ ; \xi) \ .$$

Once again, the agent chooses action $i^* = \arg\max_{i \in [n]} \tilde{u}_i^A$. Given a feasible set of potential mechanisms Φ, the goal is to choose a mechanism $\phi \in \Phi$ that maximizes the principal's overall expected utility. In particular, we seek the optimal mechanism

$$\phi_{OPT} = \arg\max_{\phi \in \Phi} \ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}_{i^*}^P(\phi)\right] \ .$$

We present two sets of results. First, we consider arbitrary nonparametric mechanisms Φ for small values of the incentive size ξ. We characterize the optimal mechanism in a restricted setting with only two treatments ($n = 2$) and uncorrelated principal-agent utilities ($\rho = 0$) in §3. We apply this result to current Medicare P4P mechanisms to derive qualitative insights about their design in §4-5; numerical experiments suggest our insights are robust to modeling assumptions. Second, in §6, we study the impact of a nonzero correlation parameter ρ and different numbers of treatments n by considering a parametric class Φ, which allows us to analytically solve for the optimal mechanism. We examine several comparative statics to demonstrate when P4P mechanisms are valuable.

3. Small Incentive Schemes

As discussed earlier, Medicare has so far chosen to make a very small portion of its total payments performance-based; they are currently capped at 1-3% of traditional fee-for-service Medicare payments to the provider. Despite significant criticism that small incentives may be insufficient (Jha 2013, Rosenthal and Dudley 2007), Medicare has maintained its stance, possibly due to institutional constraints and concerns that large performance-based penalties may lead to hospital bankruptcies (e.g., Maizel and Garner 2014). Thus, in order to make feasible policy recommendations to Medicare, we focus on the performance of P4P mechanisms under the "small incentive" regime. For analytical tractability, we consider the simplified setting where the agent only has two actions and the correlation parameter is zero (we study a more general setting in §6).

Given a feasible class of mechanisms Φ, we wish to find a mechanism that maximizes the principal's expected utility $\mathbb{E}_{\bar{\boldsymbol{x}},a}[\tilde{u}_{i^*}^P]$ up to first-order in ξ. Since we are considering small values of ξ, we will neglect terms of order ξ^2 or higher. This motivates the following definition:

Definition 1 (Optimal Small-Incentive Mechanism). Given a class of mechanisms Φ, we define an optimal small-incentive mechanism to satisfy

$$\phi^* \in \arg\max_{\phi \in \Phi} \left\{ \frac{d}{d\xi} \ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}_{i^*}^P(\phi)\right] \big|_{\xi=0} \right\} \ .$$

In other words, ϕ^* is the mechanism that improves the principal's expected utility the most with an infinitesimal incentive size. (Recall that we defined $\phi(z\,;0) = 0$ for all ϕ.) The following lemma establishes that ϕ^* is indeed first-order optimal (see proof in Appendix A):

LEMMA 1. *Let* $\phi_{OPT} \in \arg\max_{\phi\in\Phi}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi)\right]$. *Then, for* $n=2$ *and* $\rho = 0$,

$$\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi_{OPT})\right] \le \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi^*)\right] + \mathcal{O}(\xi^2)\,.$$

The next corollary follows immediately (under the same setting) for finite mechanism classes:

COROLLARY 1. *If the class of feasible mechanisms* Φ *is finite, then there exists some threshold* $\xi_0 > 0$ *such that* $\phi^* = \phi_{OPT}$ *for all small incentives* $\xi \in [0, \xi_0]$.

Our main result in this section (Proposition 1) characterizes ϕ^* (see proof in Appendix A):

PROPOSITION 1. *Consider the case where* $n = 2$ *and* $\rho = 0$. *Then, a mechanism* ϕ *is the optimal small-incentive mechanism if and only if*

$$\phi \in \arg\max_{\phi\in\Phi}\left\{\mathbb{E}_z\left[\left(\frac{(1-\cos\theta)\cdot z}{\tau\sqrt{\pi}} - 1\right)\cdot\frac{d\phi(z\,;\xi)}{d\xi}\right]\Bigg|_{\xi=0}\right\},$$

where $z \sim \mathcal{N}(0,\sigma^2)$ *and* θ *is the angle between treatment parameters* β_1 *and* β_2.

The intuition behind Proposition 1 is that any performance-based mechanism results in two effects:

1. The agent switches her action to a better one due to the performance-based incentive. This expected benefit is captured in the first term:

$$\mathbb{E}_z\left[\frac{(1-\cos\theta)\cdot z}{\tau\sqrt{\pi}}\cdot\frac{d\phi(z\,;\xi)}{d\xi}\right].$$

   Note that this benefit is decreasing in $\tau = \mathbb{E}\left[|a_i|\right]$ (the expected magnitude of the agent's private preferences); in other words, the mechanism can easily induce the agent to improve her treatment choice when the magnitude of the agent's preferences is small. The benefit is also decreasing in $\cos\theta$ (the similarity between the two treatment parameters β_1 and β_2), i.e., outcomes do not differ significantly if and when the agent switches her action. When $\theta = 0$, the treatment parameters are the same and thus, no benefit is realized from switching actions.

2. The principal incurs a cost from paying the incentive amount (this may also be a payment if the incentive is a penalty). This expected cost is captured in the second term:

$$\mathbb{E}_z\left[\frac{d\phi(z\,;\xi)}{d\xi}\right].$$

Note that this cost is always negative when ϕ is always negative (the incentive consists of only penalties for poor-performers), but may be positive when ϕ is positive for at least some values of z (the incentive also includes bonuses for top-performers)[2].

The sum of the expected benefit from the agent's change in action and the expected cost of the incentive gives us the net improvement in the principal's utility by adding mechanism ϕ with an infinitesimally small size ξ over the baseline setting of $\phi = 0$.

We observe two immediate corollaries (under the same setting). They follow by taking $\Phi = \{\phi, 0\}$, and taking the expectation of the treatment parameters β_i (for each $i \in [n]$) over the uniform prior on S^{d-1} (the d-dimensional unit sphere) respectively.

Corollary 2. *A mechanism ϕ is better than the baseline mechanism if and only if*

$$\mathbb{E}_z \left[\left(\frac{(1-\cos\theta)\cdot z}{\tau\sqrt{\pi}} - 1 \right) \cdot \frac{d\phi(z\ ;\xi)}{d\xi} \right] \Bigg|_{\xi=0} > 0\,.$$

Corollary 3. *If the treatment parameters are unknown and the principal has a uniform prior over the d-dimensional unit sphere, then Proposition 1 holds if we replace* $\cos\theta$ *with* 0.

In the following sections, we apply Proposition 1 to extract insights about existing Medicare pay-for-performance incentive schemes and potential alternatives.

4. Hospital Readmissions Reduction Program

While most P4P programs penalize poor performers, a debated issue within P4P design is whether mechanisms should reward good performers with bonuses (Rosenthal and Dudley 2006). For instance, even within Medicare's three existing P4P programs, HRRP and HACRP are *penalty-only* mechanisms, while VBP is a *bonus-penalty* mechanism since it also gives bonuses to providers with good performance metrics. This issue has been extensively explored in the literature from a behavioral economics perspective. Some analysts argue that penalties may be more effective motivational tools than bonuses since people view potential losses differently from potential gains (Kahneman and Tversky 1979). Others have argued that providers dislike penalty-based approaches and, when faced with such negative incentives, they "game" the system (Morreim 1991, Werner et al. 2002). Our model complements this line of work by providing an alternate view: incorporating both bonuses and penalties can be beneficial when the incentive size is small because it incentivizes a broader spectrum of providers to improve their choice of treatment.

[2] In some cases, the principal may not internalize monetary gains from penalizing the agent as a gain in utility, e.g., the government may institute a penalty with the goal of improving overall efficiency, but does not value monetary benefits from those penalties. In this case, we can replace this term with $\mathbb{E}_z\left[\left(\frac{d\phi(z\ ;\xi)}{d\xi}\right)^+\right]$. However, Medicare reports all collected penalties as cost savings achieved by the program (for e.g., see CMS 2015a), suggesting that this is not the view adopted by Medicare.

We use HRRP as a case study to address the design choice of using penalties vs. bonuses. The current HRRP mechanism is penalty-only; we use Proposition 1 to compare the principal's expected utility under the current mechanism against an alternative bonus-penalty mechanism that pays a symmetric bonus when providers perform well. Our results suggest that adding a symmetric bonus to the HRRP mechanism would indeed improve Medicare's expected utility.

4.1. Background

Hospital readmissions are costly for Medicare and are associated with worse outcomes for patients. While some readmissions are not preventable, studies have shown that several measures can significantly lower readmissions rates, including (i) clarifying patient discharge instructions, (ii) coordinating with post-acute care providers and patients' primary care physicians, and (iii) reducing medical complications during the initial hospital stay (Boccuti and Casillas 2015). We can view these measures as potential provider treatment choices. However, these operational interventions vary widely in implementation (and thus, effectiveness) across hospitals and are not recorded or monitored by Medicare. HRRP seeks to reduce 30-day readmission rates by penalizing hospitals with relatively high risk-adjusted readmissions rates for a fixed set of conditions. In particular, if the provider's risk-adjusted readmissions rate is higher than that of the national average, then the provider's base reimbursement is penalized by a multiplier that is equal to the normalized difference in the two rates. Consistent with Medicare's focus on small incentives, the multiplier is currently capped at 3% (CMS 2015b), despite evidence that this incentive size may be too small to significantly improve provider performance (Zhang et al. 2016).

4.2. Description of HRRP Mechanism

We present a stylized model to map the HRRP mechanism to our framework. As in our model, Medicare estimates a risk-adjustment parameter β_0 based on past national readmissions data. For each patient, they use the observed covariates $\boldsymbol{x}_0$ to compute the average probability $r = \boldsymbol{x}_0 \cdot \beta_0$ that the patient will not be readmitted[3] (under the unknown distribution of treatments allotted by providers in the US). The provider has access to several interventions that affect the rate of patient readmissions, and may result in a higher or lower readmissions rate than the national average. Intervention i is associated with parameters (β_i, a_i) such that a patient's probability of not being readmitted under this intervention is $r + \boldsymbol{x} \cdot \beta_i$, and the provider's private preference for the intervention is a_i. Let the average cost of a readmission be normalized to 1. Then, the utility or cost savings to Medicare under treatment choice i is simply the (relative) probability of

[3] In practice, a complex hierarchical logit model is used; we assume a simple linear model in our stylized setting.

no-readmission: $z_i = \boldsymbol{x} \cdot \beta_i$. Thus, we can write the following utility functions for the principal and the agent respectively:

$$\tilde{u}_i^P = r + z_i - \phi(z_i \ ;\xi)\,,$$
$$\tilde{u}_i^A = \rho \cdot (r + z_i) + a_i + \phi(z_i \ ;\xi)\,.$$

We now focus on the HRRP P4P mechanism. Let $\alpha \in (0,1)$ be the provider's (base) reimbursement for an initial hospital stay. (Note that $\alpha < 1$ because readmissions are more expensive than initial admissions (Rizzo 2013).) Then, Medicare penalizes the provider's base reimbursement α by a multiplier given by the probability of readmission above the expected rate of $1 - r$. This multiplier is capped by $\xi > 0$. Thus, we can write the HRRP mechanism in our framework as

$$\begin{aligned}\phi_0(z \ ;\xi) &= \begin{cases} 0 & \text{if } z_i > 0 \\ \alpha z_i & \text{if } z_i \in [-\xi, 0] \\ -\alpha\xi & \text{if } z_i < -\xi \end{cases} \\ &= \alpha \cdot \max\left[\min\left(z_i, 0\right), -\xi\right]\,.\end{aligned}$$

The main design parameter is the choice of cap ξ (since α is a fixed constant based on reimbursement rates). The following result (see Appendix B for proof) shows that as the cap is increased, Medicare's payoffs are strictly increased and the provider's payoffs are strictly decreased:

Proposition 2. *For any n and ρ, the principal's and agent's utilities are monotonically increasing and decreasing respectively as a function of the cap ξ under the HRRP mechanism ϕ_0.*

Thus, providers may not find contract ϕ_0 individually-rational (i.e., profitable) if the cap is large. This may be one reason why Medicare severely limits the size of the penalty through a small choice of ξ (currently set at 0.03). Accordingly, we consider the regime of small ξ.

Note that ϕ_0 is *penalty-only*; we now compare it against an alternative bonus-penalty mechanism (see Figure 2).

4.3. Alternative Bonus-Penalty Mechanism

Consider the mechanism ϕ_1, which adds a symmetric (capped) bonus to providers who perform above average:

$$\begin{aligned}\phi_1(z \ ;\xi) &= \begin{cases} \alpha\xi & \text{if } z_i > \xi \\ \alpha z_i & \text{if } z_i \in [-\xi, \xi] \\ -\alpha\xi & \text{if } z_i < -\xi \end{cases} \\ &= \alpha \cdot \operatorname{sgn}\left[z_i\right] \cdot \min\left[|z_i|, \xi\right]\,.\end{aligned}$$

Given two possible HRRP mechanisms (penalty-only ϕ_0 and bonus-penalty ϕ_1), we apply Proposition 1 to evaluate which mechanism is small-incentive optimal (proof in Appendix B).

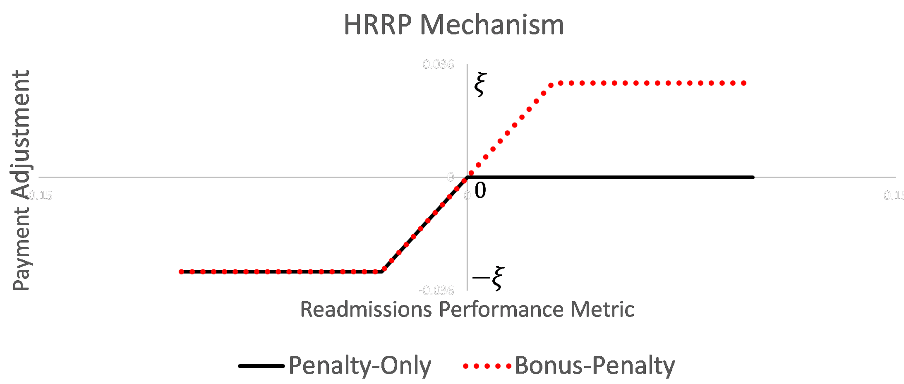

Figure 1 **Graphical comparison of the penalty-only ϕ_0 (black solid) and bonus-penalty ϕ_1 (red dashed) HRRP mechanisms as a function of the provider's performance metric.**

PROPOSITION 3. *Consider the case where $n = 2$ and $\rho = 0$. The bonus-penalty mechanism ϕ_1 is small-incentive optimal compared to the penalty-only mechanism ϕ_0 if and only if*

$$M < (1 - \cos\theta) \cdot \frac{\sqrt{2}}{\pi},$$

where we define $M = \tau/\sigma$, and θ is the angle between treatment parameters β_1 and β_2.

Recall that $\tau = \mathbb{E}_a [|a_i|]$ and $\sigma = \mathbb{E}_{\boldsymbol{x}} [|z_i|]$ for treatments $i \in [2]$. Thus, the bonus-penalty mechanism is more effective than the penalty-only mechanism for small values of the cap ξ if the expected magnitude of the agent's preferences is significantly smaller than the expected magnitude of the principal's utilities. In other words, when $M = \tau/\sigma$ is small, the mechanism can easily induce the agent to change her treatment choice, and the resulting benefit is increased by adding a bonus; on the other hand, when M is large, the agent is unlikely to change her treatment choice, and thus the cost of paying the bonus outweighs its benefit. Note that the result also depends on $1 - \cos\theta$; as noted in §3, the expected benefit from the agent changing her treatment choice is large only if the treatment parameters β_1 and β_2 are well-separated, i.e., patient outcomes differ significantly between treatments. The constant $\sqrt{2}/\pi$ is a result of our choice to model patient covariates $\boldsymbol{x}$ from a Gaussian distribution; however, our numerical experiments (§4.5) illustrate that this value varies only slightly under different distributions (e.g., uniform or exponential).

4.4. Simple Estimate

Generally, the ratio M is small in healthcare settings (see e.g., Cutler 2004), which suggests that the bonus-penalty mechanism may be more appropriate in practice. We now perform a very basic calculation to estimate if this is indeed the case for readmissions.

A provider typically chooses between action 1 (take no extra measure to reduce readmission risk) and action 2 (implement a specific intervention to reduce patient risk). One study on heart failure readmissions found that a post-discharge intervention that cost \$158 per patient (in 1997 dollars) reduced the probability of the patient being readmitted within 30 days from 22% (in the no-intervention case) to 6%, i.e., a 16% reduction in absolute probability of readmission (Anderson et al. 2005). On the other hand, empirical estimates place the cost of a 30-day heart failure readmission to be \$13,679 with a standard error of \$1,214 (in 2010 dollars), where the number of samples was 793 (Bayati et al. 2014). Controlling for inflation, \$158 dollars in 1997 is approximately \$215 in 2010. Based on these numbers, we assume the provider's utilities are $a_1 = 0$ and $a_2 = 215$, so $\tau = \sqrt{Var\,(a_i)} \approx 152$. Next, we assume the difference in Medicare's utilities is $z_2 - z_1 = 0.16 \times c$, where c is the cost of a patient readmission. Note that

$$\sigma = \sqrt{\frac{1}{2} Var\,(z_2 - z_1)} \approx \frac{0.16}{\sqrt{2}} \times \left(1214 \times \sqrt{793}\right) \approx 3868\,.$$

Also, we are in the setting with unknown treatment parameters β_1 and β_2, so we take the expectation over β_i under the uniform prior (as in Corollary 3) to get $\mathbb{E}_\beta\,[1 - \cos\theta] = 1$. Thus, we estimate

$$M \approx \frac{152}{3868} \ll \frac{\sqrt{2}}{\pi}\,.$$

While our estimates are rather coarse, the ratio between our estimated M and the threshold value of $\sqrt{2}/\pi$ is over an order of magnitude; thus, it seems likely that adding a symmetric bonus to the HRRP mechanism would indeed improve Medicare's expected utility. We leave a careful empirical investigation of this trade-off to future work.

4.5. Numerical Experiments

In this section, we relax our technical assumptions and numerically simulate various realistic scenarios. We find that our qualitative insights persist, suggesting that they are robust to our assumptions.

Data Generation. For simplicity, we consider the setting of unknown treatment parameters. (The setting with known treatment parameters simply introduces a constant factor shift, which is a function of the angles between possible treatments, but does not result in qualitatively different results.) Thus, the patient's treament-specific outcomes $\boldsymbol{x} \cdot \beta_i$ are i.i.d. random variables from $\mathcal{N}(0, \sigma^2)$. The observed baseline patient outcome $\boldsymbol{x}_0 \cdot \beta_0$ is a known constant that simply adds a constant shift to the principal's expected utility, so we normalize it to 0 without loss of generality. (Note that when $\rho \neq 0$, the baseline patient outcome also adds a constant shift to the agent's utilities, but this shift is equal across all treatments and thus, does not affect the resulting treatment choice.) We normalize $\sigma = 1$ and plot the principal's expected utility as a function of M (recall that $\tau = M\sigma$). Finally, we choose $\xi = 0.03$ to reflect the current cap imposed by Medicare's HRRP.

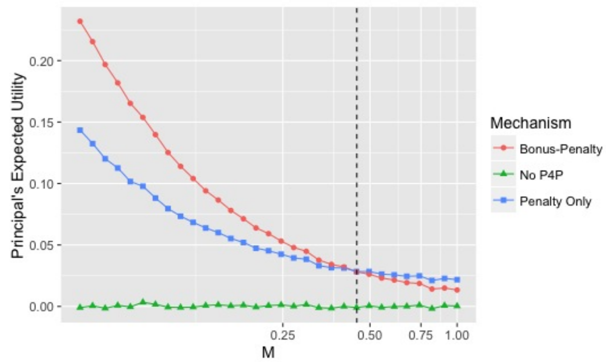

(a) $\rho = 0,\ n = 2$

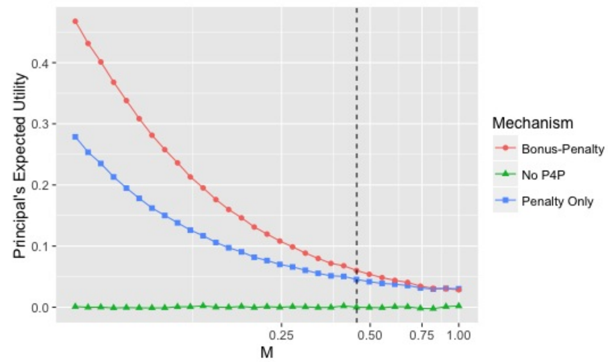

(b) $\rho = 0,\ n = 5$

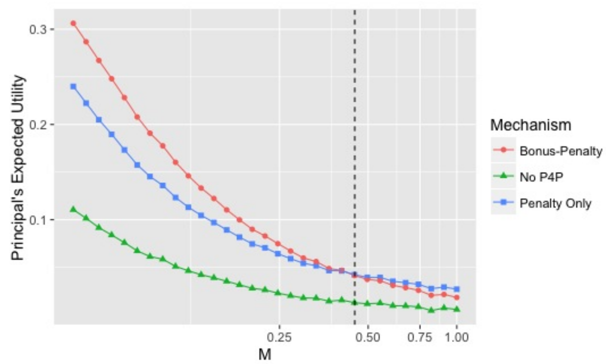

(c) $\rho = 0.01,\ n = 2$

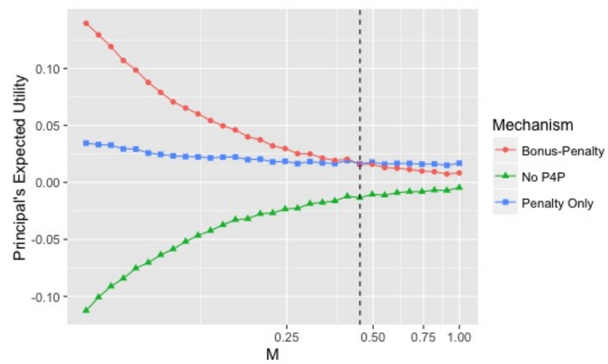

(d) $\rho = -0.01,\ n = 5$

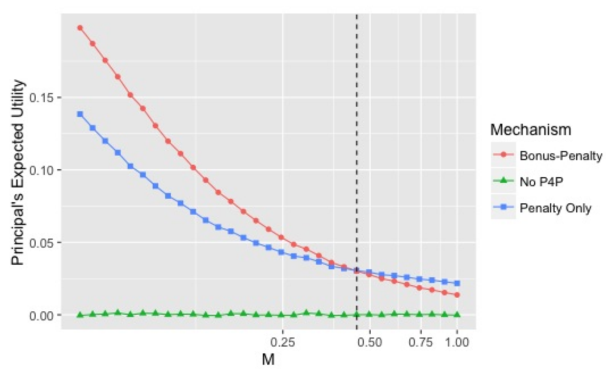

(e) $\rho = 0,\ n = 2$ with Uniform distribution

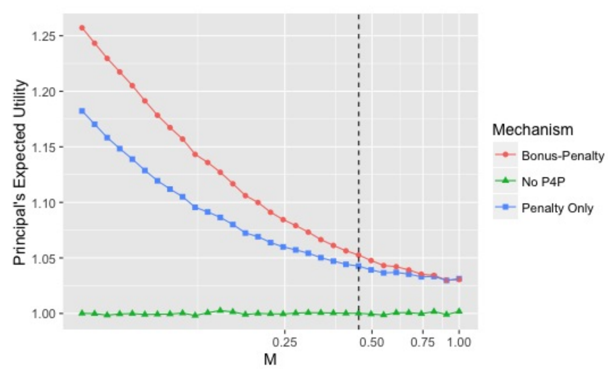

(f) $\rho = 0,\ n = 2$ with Exponential distribution

Figure 2 **Numerical simulations of HRRP mechanism as a function of the ratio $M = \tau/\sigma$. Figures (a)-(d) explore different values of ρ and n, and figures (e)-(f) consider utilities draw from the uniform and exponential distributions (rather than Gaussian) respectively. The vertical dashed line marks the threshold value of $M = \sqrt{2}/\pi$ given by Proposition 3.**

Each plot is averaged over 1 million randomly generated provider-agent pairs (we omit error bars since they are negligible due to the large number of trials).

Figure 2(a) corresponds to the setting with $n = 2$ and $\rho = 0$, as considered in Corollary 3. As predicted by Proposition 3, we see that the bonus-penalty mechanisms yields higher expected utility for the principal when M is below the crossover point of $\sqrt{2}/\pi$ (given by the dashed line); above that threshold, the penalty-only mechanism is more desirable.

Figure 2(b) considers a larger number of treatments ($n=5$). Increasing the number of treatments provides more opportunities for the agent to switch her treatment choice to a better one; thus, we expect that introducing a P4P mechanism would be more valuable. Accordingly, we see that the principal's expected utility is nearly doubled using the same P4P mechanism when $n=5$ compared to when $n=2$. Interestingly, we also note that the crossover point for the bonus-penalty and penalty-only mechanisms is shifted higher, i.e., the bonus-penalty mechanism is more likely to yield higher expected utility for the principal when there are more treatment choices.

Figures 2(c)-(d) consider a positive and a negative value of ρ respectively[4]. As expected, when ρ is positive (negative), the principals' expected utility increases (decreases) across all mechanisms. Furthermore, note that the purpose of a P4P mechanism is to better align the incentives of the agent with those of the principal. Accordingly, the value of both the bonus-penalty and penalty-only mechanisms are smaller (with respect to no P4P mechanism) when $\rho>0$ (agent's utilities are positively aligned with the principal's utilities) and larger when $\rho<0$ (agent's utilities are negatively aligned with the principal's utilities). Finally, the gap between the penalty-only mechanism and the bonus-penalty mechanism decreases (increases) when ρ is positive (negative). This may be because a larger incentive is needed (i.e., adding a bonus) to encourage the agent to improve her treatment choice when $\rho<0$ (since the agent is more likely to prefer the worse treatment).

Figures 2(e)-(f) consider altering the assumption that the patient's treatment-specific outcomes $\boldsymbol{x}\cdot\beta_i$ and the agent's private preferences a_i are normal random variables. In (e), we draw them i.i.d. from a uniform distribution, i.e., $\boldsymbol{x}\cdot\beta_i \sim U[-\sigma,\sigma]$ and $a_i \sim U[-\tau,\tau]$. In (f), we draw them i.i.d from an exponential distribution, i.e., $\boldsymbol{x}\cdot\beta_i \sim Exp[1/\sigma]$ and $a_i \sim Exp[1/\tau]$. Our results remain qualitatively similar, suggesting that they are robust to distributional assumptions.

5. Hospital-Acquired Condition Reduction Mechanism

Next, we study the issue of using a single performance threshold (providers above/below this threshold are penalized/rewarded equally) or a continuous performance threshold (the size of the reward/penalty scales in magnitude with the provider's performance). A single-threshold mechanism is desirable from a behavioral perspective since it sets a simple performance benchmark for providers to achieve. However, such a mechanism fails to incentivize improvement above the fixed threshold, and furthermore, some providers may find the threshold too difficult to meet and opt not to engage. On the other hand, the continuous-threshold approach increases the likelihood that

[4] Note that ρ is not the Pearson correlation coefficient between the principal's utilities and the agent's utilities; the correlation coefficient is given by $\frac{\rho\sigma}{\sqrt{\rho^2\sigma^2+\tau^2}}$. This quantity approaches 1 when $M=\tau\to 0$.

all providers will have an incentive to engage and improve, and thus intuition suggests that it may be a more effective mechanism overall (Rosenthal and Dudley 2006). Surprisingly, our model illustrates that the single-threshold mechanism is always more effective at maximizing the principal's expected utility when the incentive size is very small, since it dilutes the size of the incentive less, making it more likely to alter a provider's treatment choice.

We use the HACRP mechanism as a case study to address the design choice of using a single vs. a continuous performance threshold. The current HACRP mechanism is single-threshold; we use Proposition 1 to compare the principal's expected utility under the current mechanism against an alternative continuous-threshold mechanism. Our results suggest that there is no cost to Medicare's expected utility by maintaining the current single-threshold structure.

5.1. Background

Hospital-acquired conditions (HACs) are complications that patients develop during a hospital stay, and are responsible for an estimated 100,000 patient deaths (Klevens et al. 2007) and billions of dollars in added costs (Scott et al. 2009) per year. Yet, studies suggest that HAC rates can be significantly reduced through improved clinical practices, including (i) implementing staff checklists to ensure adherence to guidelines, (ii) empowering nurses to stop procedures if a violation of the guidelines was observed, and (iii) educating the staff (Berenholtz et al. 2004). Once again, we can view these measures as potential provider treatment choices. HACRP seeks to reduce HAC incidence rates by penalizing hospitals with relatively high risk-adjusted HAC rates for a fixed set of conditions. In particular, if the provider's risk-adjusted HAC rate is among the highest quartile in the nation, then the providers' base reimbursement is penalized by a fixed (small) multiplier of 1% (CMS 2014).

5.2. Description of HACRP Mechanism

We present a stylized model to map the HACRP mechanism to our framework. Again, Medicare estimates a risk-adjustment parameter β_0 based on past national HAC data. For each patient, they use the observed covariates $\boldsymbol{x}_0$ to compute the posterior distribution across providers that the patient will not acquire a HAC; the threshold $r = \boldsymbol{x}_0 \cdot \beta_0$ is set at the lower 25^{th} percentile value of this distribution. The provider has access to several interventions that affect the rate of HAC incidence, and may result in a higher or lower HAC rate compared to the nation's highest quartile. Intervention i is associated with parameters (β_i, a_i) such that a patient's probability of not acquiring a HAC under this intervention is $r + \boldsymbol{x} \cdot \beta_i$, and the provider's private preference for the intervention is a_i. Let the average cost of a HAC be normalized to 1. Then, the utility or cost savings to Medicare under treatment choice i is simply the (relative) probability of no-HAC: $z_i = \boldsymbol{x} \cdot \beta_i$. Thus, we can write the following utility functions for the principal and agent respectively

$$\begin{aligned}\tilde{u}_i^P &= r + z_i - \phi(z_i \; ; \xi) \,, \\ \tilde{u}_i^A &= \rho \cdot (r + z_i) + a_i + \phi(z_i \; ; \xi) \,.\end{aligned}$$

We now focus on the HACRP P4P mechanism. Let α be the provider's (base) reimbursement for a HAC-free hospital stay. Then, Medicare penalizes the provider's base reimbursement α by a multiplier $\xi > 0$ given by the probability of a HAC above the threshold rate of $1 - r$. Thus, we can write the HACRP mechanism is our framework as

$$\begin{aligned}\phi_0(z \; ; \xi) &= \begin{cases} 0 & \text{if } z_i \geq 0 \\ -\alpha\xi & \text{if } z_i < 0 \end{cases} \\ &= -\alpha\xi \cdot \mathbb{I}\left[z_i < 0\right] \,.\end{aligned}$$

The main design parameter is the (small) choice of multiplier ξ (since α is a fixed constant based on reimbursement rates). Note that ϕ_0 has a *single-threshold*; we will now compare its performance against an alternative continuous-threshold mechanism (see Figure 4).

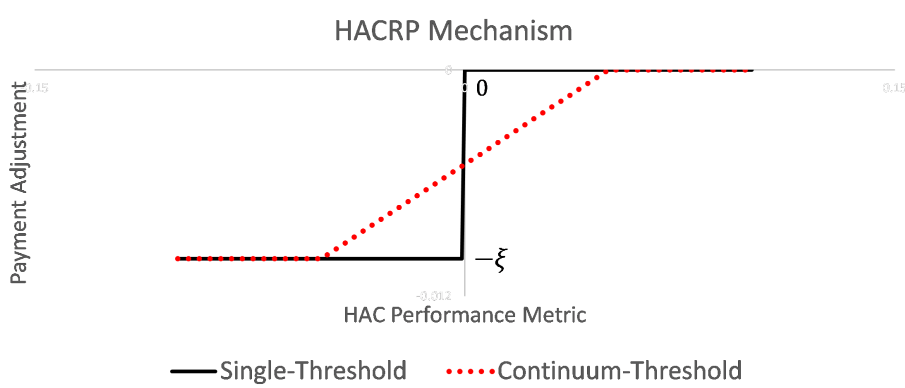

Figure 3 **Graphical comparison of the single-threshold ϕ_0 (black solid) and continuous-threshold ϕ_1 (red dashed) HACRP mechanisms as a function of the provider's performance metric.**

5.3. Alternative Continuous-Threshold Mechanism

Consider the mechanism ϕ_1, which introduces a gradual slope to the penalty function so that providers are penalized proportional to their performance (up to a cap). This slope is parametrized by the standard deviation σ of the performance z:

$$\phi_1(z \; ; \xi) = \begin{cases} 0 & \text{if } z_i > \sigma/2 \\ \frac{\alpha\xi}{\sigma}(z_i - \sigma/2) & \text{if } z_i \in [-\sigma/2, \sigma/2] \\ -\alpha\xi & \text{if } z_i < -\sigma/2 \end{cases}$$

$$= \frac{\alpha\xi}{\sigma} \cdot \max\left[\min\left(z_i - \sigma/2, 0\right), -\sigma\right] .$$

Note that both mechanisms ϕ_0 and ϕ_1 have the same area under the curve (see Fig 4). This ensures that the expected penalty is the same under both mechanisms; thus, any difference in the principal's expected utility is a result of the shape of the mechanism rather than the size of the incentive.

Given two possible HACRP mechanisms (single-threshold ϕ_0 and continuous-threshold ϕ_1), we use Proposition 1 to evaluate which mechanism is small-incentive optimal (proof in Appendix B).

PROPOSITION 4. *Consider the case where $n = 2$ and $\rho = 0$. The single-threshold mechanism ϕ_0 is always small-incentive optimal compared to the continuous-threshold mechanism ϕ_1.*

Thus, the single-threshold mechanism is more effective than the continuous-threshold mechanism for small values of the maximum penalty multiplier ξ. In other words, although the continuous-threshold mechanism penalizes a greater number of (poor) treatment choices, the dilution of the size of the incentive makes it less effective in the "small incentive" regime since it is less likely to alter a provider's treatment choice.

5.4. Numerical Experiments

We now relax our technical assumptions and numerically simulate analogous scenarios as in §4. We find that the threshold policy always outperforms the continuous-threshold policy across a wide range of parameters and distributions. However, the empirical improvement appears to be negligible compared to the value of adding a P4P mechanism.

Data Generation. We use the same problem parameters as outlined in the numerical experiments in §4. However, we choose $\xi = 0.01$ to reflect the current cap imposed by Medicare's HACRP. Once again, each plot is averaged over 1 million trials (randomly generated provider-agent pairs).

Once again, Figure 4(a) corresponds to $n = 2$ and $\rho = 0$ (as considered in Proposition 4). Figures 4(b)-(d) vary the number of treatments n and the parameter ρ. Figures 4(e)-(f) consider alternate distributions for the patient's treatment-specific outcomes and the agent's private preferences: (e) $\boldsymbol{x} \cdot \beta_i \sim U[-\sigma, \sigma]$ and $a_i \sim U[-\tau, \tau]$; while (f) $\boldsymbol{x} \cdot \beta_i \sim Exp[1/\sigma]$ and $a_i \sim Exp[1/\tau]$.

As suggested by the theory, we see that the single-threshold mechanism always yields (slightly) higher expected utility for the principal compared to the continuous-threshold mechanism across all parameters and distributions. However, this difference is often negligible, suggesting that either mechanism is viable. This remains surprising since intuition suggests that a continuous-threshold policy may incentivize a wider range of providers (e.g., see discussion in Rosenthal and Dudley 2006). Yet, we find that a single-threshold policy can achieve equivalent or better expected utility for the principal. Furthermore, a single-threshold policy may be desirable from a behavioral perspective since it sets a simple performance benchmark for providers to achieve (Rosenthal and Dudley 2006).

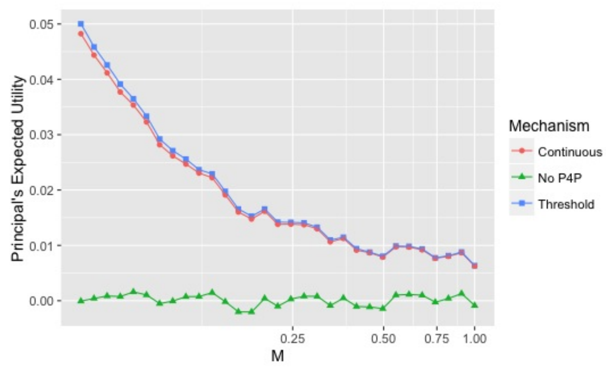

(a) $\rho = 0,\ n = 2$

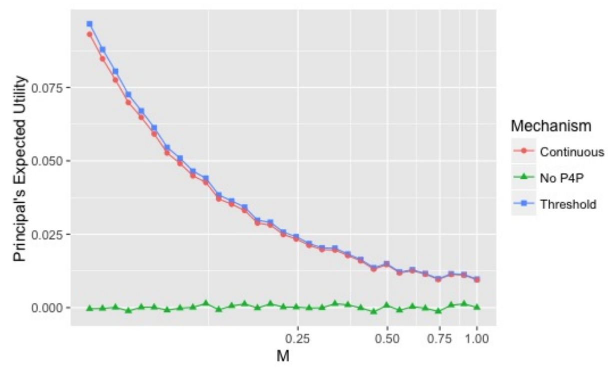

(b) $\rho = 0,\ n = 5$

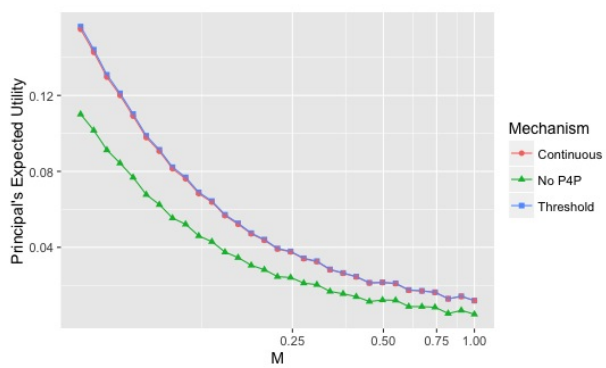

(c) $\rho = 0.01,\ n = 2$

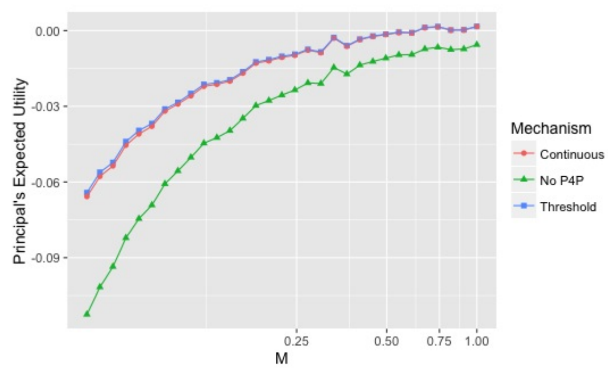

(d) $\rho = -0.01,\ n = 5$

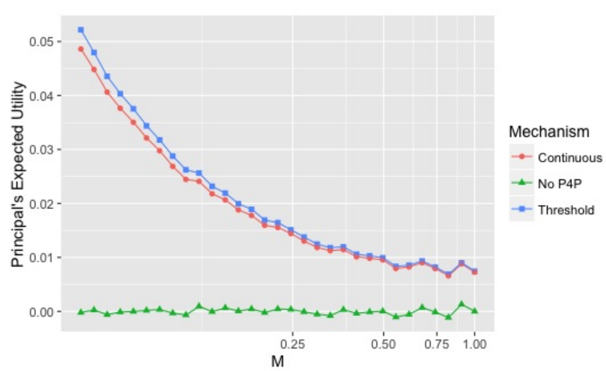

(e) $\rho = 0,\ n = 2$ with Uniform distribution

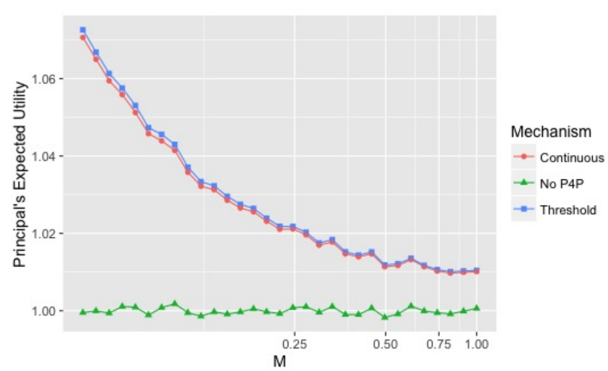

(f) $\rho = 0,\ n = 2$ with Exponential distribution

Figure 4 Numerical simulations of HACRP mechanisms as a function of the ratio $M = \tau/\sigma$. Figures (a)-(d) explore different values of ρ and n, and figures (e)-(f) consider utilities draw from the uniform and exponential distributions (rather than Gaussian) respectively.

These simulations also reiterate two points that were noted earlier in §4. In particular, a P4P mechanism is more valuable when (i) there are more treatment choices, and (ii) $\rho < 0$ so the agent's utilities are negatively correlated with the principal's utilities.

6. Comparative Statics

We now study the addition of a nonzero correlation parameter ρ in the principal-agent utilities, as well as changes in the number of available treatments n. For the purpose of gaining qualitative insights, we remove the "small incentive" constraint and consider the class of feasible mechanisms to be all linear mechanisms, which allows us to analytically express the optimal mechanism. We examine a number of comparative statics that illustrate when a P4P mechanism can be effective at improving the principal's expected utility. Importantly, we find that a P4P mechanism is *always* beneficial if the correlation parameter is negative, i.e., the agent's utilities are negatively correlated with the principal's utilities. In practice, this may often be the case, since providers can undertake costly interventions (e.g., extending a patient's length of stay for observation) to improve patient outcomes, but these interventions are not reimbursed under traditional fee-for-service, thereby creating misaligned incentives (Green 2012). Such settings necessitate a performance-based incentive.

6.1. Linear Mechanisms

In order to solve for the optimal mechanism, we limit the class of feasible mechanisms Φ to the set of linear mechanisms (see Definition 2). Linear contracts are used abundantly in practice due to their simplicity and flexibility. Recent work has also shown that they are optimally robust in moral hazard problems (Carroll 2015). We formally define this class below:

Definition 2 (Linear Mechanisms). The class of linear mechanisms is

$$\Phi \;=\; \left\{ \phi_\xi \;\middle|\; \phi_\xi(z) = \xi z \;\; \text{for} \;\; \xi \in [0,1) \right\} .$$

Furthermore, for analytical tractability, we consider the setting where Medicare does not know the treatment parameters β_i. As discussed earlier, this is a realistic assumption for Medicare's P4P policies, since providers often improve performance metrics through operational interventions (e.g., Berenholtz et al. 2004); however, such interventions vary in implementation (and therefore, efficacy) by provider, and are not reported in claims data, making it difficult for Medicare to estimate their associated treatment parameters. Thus, we write the principal-agent utilities as follows:

$$\tilde{u}_i^P\left(\phi_\xi\right) = z_0 + z_i - \phi_\xi\left(z_i\right) \,,$$
$$\tilde{u}_i^A\left(\phi_\xi\right) = \rho\left(z_0 + z_i\right) + a_i + \phi_\xi\left(z_i\right) \,,$$

where the known baseline outcome is $z_0 = \boldsymbol{x}_0 \cdot \beta_0$, the treatment-specific outcomes are $z_i \sim \mathcal{N}(0,\sigma^2)$ and the agent's private treatment-specific preferences are $a_i \sim \mathcal{N}(0,\tau^2)$. In particular, we no longer model $z_i = \boldsymbol{x} \cdot \beta_i$; this is because, from Medicare's perspective, the unobserved patient covariates $\boldsymbol{x}$ and the treatment parameters β_i are now both unknown quantities (that are known to the agent). Thus, we can simplify the problem by considering their product z_i (the source of information asymmetry) to directly be a normal random variable. Then, we define $\bar{\boldsymbol{z}} = [z_0, \boldsymbol{z}]$, where $\boldsymbol{z} \in \mathbb{R}^n$ is the vector whose components are z_i.

In order to evaluate the optimal mechanism, we define an efficiency function:

DEFINITION 3 (THE EFFICIENCY FUNCTION). The efficiency $\eta(\xi)$ is the ratio of the principal's (risk-adjusted) expected utility from mechanism ϕ_ξ, and the first-best (risk-adjusted) outcome. Thus,

$$\eta(\xi) \equiv \frac{\mathbb{E}_{\bar{\boldsymbol{z}},a}\left[\tilde{u}^P_{i^*}(\phi_\xi)\right] - \mathbb{E}_{\bar{\boldsymbol{z}}}\left[z_0\right]}{\mathbb{E}_{\bar{\boldsymbol{z}}}\left[\max_j u^P_j\right] - \mathbb{E}_{\bar{\boldsymbol{z}}}\left[z_0\right]},$$

where we have adjusted for the baseline patient outcome z_0.

The principal clearly wishes to choose a mechanism ϕ_ξ that maximizes the efficiency function. Proposition 5 derives an exact analytical expression for $\eta(\xi)$.

REMARK 1. The efficiency function normalizes the principal's expected utility (which was directly studied in §3) by the first-best outcome (which is a constant in expectation). We introduce this normalization to simplify our analytical expressions and comparative statics.

PROPOSITION 5. *The efficiency function is given by*

$$\eta(\xi) = (1-\xi)\cdot\left(1 + \frac{M^2}{(\rho+\xi)^2}\right)^{-1/2},$$

where we have defined the ratio $M = \tau/\sigma$.

Thus, we find that the efficiency depends only on two parameters:

1. $\boldsymbol{M = \tau/\sigma}$: The efficiency of any mechanism ϕ_ξ is decreasing in the ratio M. This insight was noted in Proposition 3 as well: if $\tau = \mathbb{E}_a\left[|a_i|\right]$ (the expected magnitude of the agent's preferences) is significantly smaller than $\sigma = \mathbb{E}_{\boldsymbol{z}}\left[|z_i|\right]$ (the expected magnitude of the principal's utilities), then a relatively cheap performance-based incentive can induce the agent to improve her treatment choice by a significant margin.
2. $\boldsymbol{\rho}$: The efficiency of any mechanism ϕ_ξ is increasing in the parameter ρ; in other words, as the agent's utilities are more aligned with the principal's utilities, the expected outcome approaches the first-best outcome (where the agent directly maximizes the principal's utility).

Surprisingly, the efficiency function does not depend on the number of treatments n.

We seek the *optimal mechanism* ϕ_{ξ^*}, where

$$\xi^* = \arg\max_{\xi\in[0,1]}\left\{\eta(\xi)\right\}.$$

PROPOSITION 6. *There exists a unique optimal $\xi^* = \arg\max_{[0,1]}\eta(\xi)$, where either $\xi^* \in \{0,1\}$ or ξ^* is the unique real zero of the cubic polynomial*

$$p(x) = (x+\rho)^3 + 2M^2 x + (\rho - 1)\,M^2.$$

The ξ-incentive mechanism improves efficiency over the baseline, i.e. $\xi^* > 0$, when

$$\rho^3 + M^2 (\rho - 1) < 0 .$$

We observe the following immediate corollary:

COROLLARY 4. *There is always a non-trivial linear P4P mechanism that improves the principal's expected utility when the agent's utilities are negatively correlated with the principal's utilities, i.e., the optimal incentive ξ^* is positive when ρ is negative.*

We now perform comparative statics with respect to the two parameters ρ and M to gain insight about the conditions under which a linear P4P mechanism can improve the principal's expected utility.

PROPOSITION 7. *Under a linear mechanism, the size of the optimal incentive ξ^**

(a) decreases as the agent's utility is increasingly aligned with the principal's utility, i.e.,

$$\frac{d\xi^*}{d\rho} < 0 ,$$

(b) increases as the expected magnitude of the agent's private preferences increases with respect to the expected magnitude of the principal's utilities, i.e.,

$$\frac{d\xi^*}{dM} > 0 .$$

The purpose of a P4P mechanism is to better align the utility funtions of the principal and agent. As ρ increases, the principal-agent utilities become increasingly aligned in the baseline setting, thereby removing the need for a large performance-based incentive. In contrast, we noted earlier that the efficiency of *any* linear mechanism is increasing in ρ. On the other hand, as $M = \tau/\sigma$ increases, the agent's private preferences begin to outweigh the agent's interest in the principal's utility, and thus a larger performance-based incentive is needed to align the principal-agent utilities. In contrast, we noted earlier that the efficiency of *any* linear mechanism is decreasing in M.

7. Discussion & Concluding Remarks

We analyze Medicare's pay-for-performance programs, paying particular attention to the use of "small incentives." We develop a framework for comparing the efficiency of any feasible set of mechanisms (subject to Medicare's practical and institutional constraints) under the regime of small incentives using a stylized principal-agent model. We then apply this result to current Medicare policies to examine two debated design choices: the use of rewards vs. bonuses, and the use of single or continuous performance thresholds. Our analysis yields several insights and policy recommendations for Medicare. We also examine several comparative statics to provide intuition on the conditions under which P4P mechanisms are effective.

In practice, many medical treatments are delivered by a chain of providers that are not part of the same organization (and are reimbursed separately)[5]. Our analysis and insights can be easily extended to this network setting, where agents choose both a treatment and a possible referral to a specialist, by employing a broader interpretation of the agent's action space.

There are several avenues for future work. A careful empirical investigation of provider and agent utility functions can provide finer insight into design choices for feasible mechanisms. For example, Fuloria and Zenios (2001) and Aswani et al. (2016) outline empirical methods for estimating these parameters. However, these papers often assume that providers are risk-neutral and fully rational agents, which may not be the case in practice. In particular, Rosenthal and Dudley (2006) discuss a few behavioral effects that may affect how providers respond to such contracts; for instance, losses (penalties) are viewed differently from gains (rewards) (Kahneman and Tversky 1979), and negative nudges like penalties may be more likely to encourage providers to "game" the system (Morreim 1991). Thus, the study of such behavioral effects may greatly augment our understanding of providers' true response functions and therefore, help optimize P4P mechanism design.

Appendix

A. Small Incentive Schemes

Let $\boldsymbol{z} \in \mathbb{R}^n$, where each $z_i = \boldsymbol{x} \cdot \beta_i$ for each $i \in [n]$. Also, let $z_0 = \boldsymbol{x}_0 \cdot \beta_0$.

A.1. Basic Lemmas

We first state some simple lemmas that will be useful for the proofs.

Lemma 2. *If we have a random variable $\boldsymbol{x} \sim \mathcal{N}(0^d, \sigma^2 I)$ and a unit-normed constant $\beta \in \mathbb{R}^d$ that satisfies $\|\beta\|_2 = 1$, then $\boldsymbol{x} \cdot \beta$ is a scalar random variable with distribution $\mathcal{N}(0, \sigma^2)$.*

Proof of Lemma 2 There exists a unitary transformation $U \in \mathbb{R}^{d \times d}$ such that $U\beta = e_1$, where $e_1 = [1, 0, ..., 0] \in \mathbb{R}^d$. We can then write

$$\begin{aligned} \boldsymbol{x}^T \beta &= (U\boldsymbol{x})^T (U\beta) \\ &= (U\boldsymbol{x})_1 \ \sim \ \mathcal{N}(0, \sigma^2), \end{aligned}$$

where we have used the fact that both inner products and $\boldsymbol{x}$ (a spherically symmetric normal variable) are invariant under unitary transformations.

□

[5] An example of such a chain may include a primary care physician, a specialist, a surgeon, a hospital, and a physical therapist. The overall outcome of the treatment depends on the treatment choices of all agents along this chain, but the payer only observes the overall outcome of the process, which makes direct pay-for-performance mechanisms to each provider infeasible.

LEMMA 3. *Let θ_{ij} be the angle between vectors β_i and β_j for any $i, j \in [n]$. Then,*

$$\boldsymbol{z} \sim \mathcal{N}(0, \sigma^2 \Sigma),$$

with $\Sigma_{i,j} = \cos\theta_{ij}$.

Proof of Lemma 3 Let $B \in \mathbb{R}^{n\times d}$ be the matrix whose i^{th} row is β_i. Then, $\boldsymbol{z} = B\boldsymbol{x}$ so $\boldsymbol{z}$ is a linear transformation of $\boldsymbol{x}$ and therefore jointly Gaussian as well. Using Lemma 2, note that

$$z_i - z_j \;=\; \boldsymbol{x}\cdot(\beta_i - \beta_j) \;\sim\; \mathcal{N}(0, \|\beta_i - \beta_j\|_2^2 \sigma^2).$$

The law of cosines gives us

$$\|\beta_i - \beta_j\|_2^2 \;=\; \|\beta_i\|_2^2 + \|\beta_j\|_2^2 - 2\cdot\|\beta_i\|_2^2 \cdot \|\beta_j\|_2^2 \cdot \cos\theta_{ij} \;=\; 2 - 2\cos\theta_{ij},$$

so $z_i - z_j \sim \mathcal{N}(0, 2\sigma^2(1-\cos\theta_{ij}))$. Then, we evaluate

$$\mathbb{E}[z_i z_j] \;=\; \frac{1}{2}\left(\mathbb{E}[z_i^2] + \mathbb{E}[z_j^2] - \mathbb{E}[(z_i - z_j)^2]\right) \;=\; \sigma^2\cos\theta_{ij}.$$

Thus, $\Sigma_{i,j} = \cos\theta_{ij}$.

□

A.2. Proof of Lemma 1

LEMMA 4. *Let mechanism $\phi(z\ ;\xi)$ have a uniformly bounded and continuous second derivative $\left|\frac{d^2\phi}{d\xi}\right|$ with respect to the incentive size ξ. Then, for $n = 2$ and $\rho = 0$, the principal's expected utility $\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}_{i^*}^P(\phi(z\ ;\xi))\right]$ also has a uniformly bounded and continuous second derivative with respect to ξ.*

Proof of Lemma 4 First, note that

$$\frac{d}{d\xi}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}_{i^*}^P(\phi(z\ ;\xi)) - u_{i^*}^P\right] = \frac{d}{d\xi}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}_{i^*}^P(\phi(z\ ;\xi))\right].$$

Thus, it suffices to show that $\frac{d^2}{d\xi^2}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}_{i^*}^P(\phi(z\ ;\xi)) - u_{i^*}^P\right]$ is uniformly bounded and continuous. We now evaluate $\tilde{u}_{i^*}^P(\phi(z\ ;\xi)) - u_{i^*}^P$, assuming $n = 2$ and $\rho = 0$. Using our new notation:

$$\tilde{u}_i^P = z_0 + z_i - \phi(z_i\ ;\xi),$$
$$\tilde{u}_i^A = a_i + \phi(z_i\ ;\xi).$$

Recall that $\boldsymbol{z}$ is a jointly Gaussian random variable (Lemma 3) and $a \sim \mathcal{N}(0, \tau^2 I)$ by definition.

Assume without loss of generality that treatment 1 yields higher performance for the principal, i.e., $z_1 > z_2$. We now have 3 possible cases:

1. $a_1 > a_2$: the agent would have chosen the optimal action with and without the added incentive. (Note that $a_1 > a_2$ implies $a_1 + \phi(z_1\ ;\xi) > a_2 + \phi(z_2\ ;\xi)$ since we $z_1 > z_2$ and ϕ is non-decreasing in z.) Thus, the incentive induces no change in the agent's action and costs the principal $\phi(z_1\ ;\xi)$.

2. $a_1 < a_2$ but $a_1 + \phi(z_1\ ;\xi) > a_2 + \phi(z_2\ ;\xi)$: the agent would have chosen the sub-optimal action without the incentive but chooses the optimal action with the incentive. Thus, the incentive results in an added utility of $z_1 - z_2$ (through the agent's change in action) and costs $\phi(z_1\ ;\xi)$.

3. $a_1 + \phi(z_1\ ;\xi) < a_2 + \phi(z_2\ ;\xi)$: the agent chooses the sub-optimal action even with the incentive. Thus, the incentive induces no change in the agent's action and costs the principal $\phi(z_2\ ;\xi)$.

Summing the terms above, we define

$$\begin{aligned}D(\phi) \equiv &- \phi(z_1\ ;\xi)\cdot \mathbb{I}[a_1 > a_2] + (z_1 - z_2 - \phi(z_1\ ;\xi))\cdot \mathbb{I}[a_2 > a_1 \text{ and } a_1 + \phi(z_1\ ;\xi) > a_2 + \phi(z_2\ ;\xi)] \\ &- \phi(z_2\ ;\xi)\cdot \mathbb{I}[a_1 + \phi(z_1\ ;\xi) < a_2 + \phi(z_2\ ;\xi)]\,,\end{aligned}$$

where we note that

$$\mathbb{E}_{\boldsymbol{z},a}\left[D(\phi) \mid z_1 > z_2\right] \;=\; \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}\left(\phi\left(z\ ;\xi\right)\right) - u^P_{i^*}\right]\,.$$

We first take the expectation of $D(\phi)$ over the random vector $a \sim \mathcal{N}(0,\tau^2 I)$. Note that $\mathbb{E}_a\left[D(\phi) \mid z_1 > z_2\right] = \mathbb{E}_a\left[D(\phi)\right]$ since $\boldsymbol{z}, a$ are independent random variables. Also, we can simplify $\Pr[a_1 > a_2] = 1/2$ by symmetry and $a_2 - a_1 \sim \mathcal{N}(0, 2\tau^2)$. Thus, by taking $F(\cdot)$ to be the cdf of $\mathcal{N}(0,2\tau^2)$, we have the following identities:

$$\begin{aligned}\Pr[a_2 > a_1 \text{ and } a_1 + \phi(z_1\ ;\xi) > a_2 - \phi(z_2\ ;\xi)] &= \Pr[0 < a_2 - a_1 < \phi(z_1\ ;\xi) - \phi(z_2\ ;\xi)] \\ &= F\left(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi)\right) - 1/2\,, \\ \Pr[a_1 + \phi(z_1\ ;\xi) < a_2 + \phi(z_2\ ;\xi)] &= 1 - F(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi))\,.\end{aligned}$$

Substituting, we have

$$\begin{aligned}\mathbb{E}_a\left[D(\phi) \mid z_1 > z_2\right] = &- \frac{\phi(z_1\ ;\xi)}{2} + (z_1 - z_2 - \phi(z_1\ ;\xi))\cdot\left[F\left(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi)\right) - \frac{1}{2}\right] \\ &- \phi(z_2\ ;\xi)\cdot\left[1 - F(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi))\right] \\ = &\; F\left(\phi\left(z_1\ ;\xi\right) - \phi\left(z_2\ ;\xi\right)\right)\cdot\left[z_1 - z_2 - \phi(z_1\ ;\xi) + \phi(z_2\ ;\xi)\right] - \phi(z_2\ ;\xi) - \frac{z_1 - z_2}{2}\,.\end{aligned}$$

Next, note that we can commute

$$\frac{d^2\mathbb{E}_{\boldsymbol{z},a}\left[D(\phi) \mid z_1 > z_2\right]}{d\xi^2} = \mathbb{E}_{\boldsymbol{z}}\left[\frac{d^2\mathbb{E}_a\left[D(\phi)\right]}{d\xi^2}\middle| z_1 > z_2\right]\,.$$

Here, we first used Fubini's theorem to change the order of expectations (note that $|D(\phi)|$ is integrable since we assumed $|\phi(z\ ;\xi)|$ is bounded by a polynomial in z and z is a Gaussian random variable); next, we used Leibniz's integral rule to move in the derivative with respect to ξ (note that $D(\phi)$ and its first two derivatives with respect to ξ are clearly continuous since we assumed ϕ satisfies these properties).

From the definition of $D(\phi)$, it is clear that $\mathbb{E}_a\left[D(\phi)\right]$ has a continuous and uniformly bounded second derivative if $\frac{d^2}{d\xi^2}\phi$ is continuous and uniformly bounded. (Note that F is the cdf of a normal variable and thus, all its derivatives are continuous and bounded.) It can then be easily verified that taking the conditional expectation with respect to the normal variable $\boldsymbol{z}$ preserves this property.

□

Proof of Lemma 1 Fix a mechanism ϕ and define $f(\xi) = \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}\left(\phi\left(z\ ;\xi\right)\right)\right]$. From Lemma 4, we know that f has a uniformly bounded and continuous second derivative with respect to ξ. We can then apply a first-degree Taylor series expansion around zero:

$$f(\xi) = f(0) + \xi\cdot\left.\frac{df}{d\xi}\right|_{\xi=0} + R_1\,,$$

and we can bound the Lagrange remainder term R_1 (for some $\xi^* \in (0,\xi)$):

$$\begin{aligned}|R_1| &= \frac{1}{2}\left|\frac{d^2 f(\xi^*)}{d\xi^2}\right| \cdot \xi^2 \\ &\leq C\xi^2 ,\end{aligned}$$

for some constant C (half the uniform bound on the second derivative of f). Then, for any mechanism ϕ,

$$\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi)\right] = \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[u^P_{i^*}\right] + \xi \cdot \left.\frac{d\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi)\right]}{d\xi}\right|_{\xi=0} + \mathcal{O}\left(\xi^2\right) .$$

Then, by definition of ϕ^*,

$$\begin{aligned}\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi_{OPT})\right] - \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi^*)\right] &= \xi \cdot \left(\left.\frac{d\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi_{OPT})\right]}{d\xi}\right|_{\xi=0} - \left.\frac{d\mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi^*)\right]}{d\xi}\right|_{\xi=0}\right) + \mathcal{O}\left(\xi^2\right) \\ &\leq \mathcal{O}\left(\xi^2\right) .\end{aligned}$$

□

A.3. Proof of Proposition 1

Proof of Proposition 1 By Definition 1, we wish to find

$$\phi^* \in \arg\max_{\phi\in\Phi}\left\{\frac{d}{d\xi}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi)\right]\Big|_{\xi=0}\right\} .$$

Recall from the proof of Lemma 4 that for any $\phi \in \Phi$,

$$\frac{d}{d\xi}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^P_{i^*}(\phi(z\ ;\xi))\right]\Big|_{\xi=0} = \frac{d}{d\xi}\ \mathbb{E}_{\boldsymbol{z},a}\left[D(\phi)|z_1 > z_2\right]\Big|_{\xi=0} = \mathbb{E}_{\boldsymbol{z}}\left[\frac{d}{d\xi}\ \mathbb{E}_a\left[D(\phi)\right]\Big|_{\xi=0}\Big| z_1 > z_2\right] ,$$

where we derived

$$\mathbb{E}_a\left[D(\phi)\right] = F\left(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi)\right)\cdot\left(z_1 - z_2 - \phi(z_1\ ;\xi) + \phi(z_2\ ;\xi)\right) - \phi(z_2\ ;\xi) - \frac{z_1 - z_2}{2} .$$

Recall that F is the cdf of the random variable distributed as $\mathcal{N}(0, 2\tau^2)$. Next, we evaluate $\frac{d}{d\xi}\left\{\mathbb{E}_a D(\phi)\right\}\big|_{\xi=0}$. For ease of notation, we define

$$\psi(z) := \left.\frac{d\phi(z\ ;\xi)}{d\xi}\right|_{\xi=0} .$$

Note that we can simplify

$$\begin{aligned}\frac{d}{d\xi}F\left(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi)\right)\Big|_{\xi=0} &= f\left(\phi(z_1\ ;\xi) - \phi(z_2\ ;\xi)\right)\cdot\left(\frac{d\phi(z_1\ ;\xi)}{d\xi} - \frac{d\phi(z_2\ ;\xi)}{d\xi}\right)\Big|_{\xi=0} \\ &= \frac{1}{2\tau\sqrt{\pi}}\left[\psi(z_1) - \psi(z_2)\right] ,\end{aligned}$$

using that fact that $\phi(z;\xi = 0) = 0$ and $f(0) = 1/(2\tau\sqrt{\pi})$. Substituting and using $F(0) = 1/2$ yields

$$\begin{aligned}\frac{d}{d\xi}\mathbb{E}_a\left[D(\phi)\right]\Big|_{\xi=0} &= \frac{z_1 - z_2 - \phi(z_1;0) + \phi(z_2;0)}{2\tau\sqrt{\pi}}\left[\psi(z_1) - \psi(z_2)\right] + F\left(\phi(z_1;0) - \phi(z_2;0)\right)\cdot\left[-\psi(z_1) + \psi(z_2)\right] - \psi(z_2) \\ &= \frac{z_1 - z_2}{2\tau\sqrt{\pi}}\left[\psi(z_1) - \psi(z_2)\right] + \frac{1}{2}\left[-\psi(z_1) + \psi(z_2)\right] - \psi(z_2) \\ &= \frac{z_1 - z_2}{2\tau\sqrt{\pi}}\left[\psi(z_1) - \psi(z_2)\right] - \frac{1}{2}\left[\psi(z_1) + \psi(z_2)\right] .\end{aligned}$$

Finally, we can evaluate $\frac{d}{d\xi}\left\{\mathbb{E}_{\boldsymbol{z},a}\left[D(\phi)\mid z_1>z_2\right]\right\}\Big|_{\xi=0}$ by taking the expectation with respect to $\boldsymbol{z}$. Taking θ to be the angle between vectors β_1 and β_2, it follows from Lemma 3 that

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \sim \mathcal{N}\left(0,\ \sigma^2 \begin{bmatrix} 1 & \cos\theta \\ \cos\theta & 1 \end{bmatrix}\right).$$

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{z}}\left[(z_1-z_2)\left[\psi(z_1)-\psi(z_2)\right]\Big| z_1>z_2\right] &= 2\cdot\mathbb{E}_{\boldsymbol{z}}\left[(z_1-z_2)\left[\psi(z_1)-\psi(z_2)\right]\cdot\mathbb{I}\left[z_1>z_2\right]\right] \\
&= 2\cdot\mathbb{E}_{\boldsymbol{z}}\left[\left(z_1\psi(z_1)+z_2\psi(z_2)-z_1\psi(z_2)-z_2\psi(z_1)\right)\cdot\mathbb{I}\left[z_1>z_2\right]\right] \\
&= 2\cdot\mathbb{E}_{\boldsymbol{z}}\left[z_1\psi(z_1)\cdot\mathbb{I}\left[z_1>z_2\right]\right]+2\cdot\mathbb{E}_{\boldsymbol{z}}\left[z_1\psi(z_1)\cdot\mathbb{I}\left[z_2>z_1\right]\right] \\
&\quad -2\cdot\mathbb{E}_{\boldsymbol{z}}\left[z_1\psi(z_2)\cdot\mathbb{I}\left[z_1>z_2\right]\right]-2\cdot\mathbb{E}_{\boldsymbol{z}}\left[z_1\psi(z_2)\cdot\mathbb{I}\left[z_2>z_1\right]\right] \\
&= 2\cdot\mathbb{E}_{\boldsymbol{z}}\left[z_1\psi(z_1)\right]-2\cdot\mathbb{E}_{\boldsymbol{z}}\left[z_1\psi(z_2)\right] \\
&= 2\cdot\mathbb{E}_{z}\left[z\psi(z)\right]-2\cdot\mathbb{E}_{\boldsymbol{z}}\left[(z_1-\cos\theta\cdot z_2)\,\psi(z_2)\right]-2\cdot\mathbb{E}_{\boldsymbol{z}}\left[\cos\theta\cdot z_2\psi(z_2)\right] \\
&= 2\left(1-\cos\theta\right)\cdot\mathbb{E}_{z}\left[z\psi(z)\right],
\end{aligned}$$

where $z\sim\mathcal{N}\left(0,\sigma_2^2\right)$. Note that we have used index exchangeability of z_1, z_2 combined with the fact that $\mathbb{I}\left[z_1>z_2\right]+\mathbb{I}\left[z_2>z_1\right]=1$, as well as the independence of $z_1-\cos\theta\cdot z_2$ and z_2. Similarly, we have

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{z}}\left[\psi(z_1)+\psi(z_2)\Big| z_1>z_2\right] &= 2\cdot\mathbb{E}_{\boldsymbol{z}}\left[\psi(z_1)+\psi(z_2)\cdot\mathbb{I}\left[z_1>z_2\right]\right] \\
&= 2\cdot\mathbb{E}_{\boldsymbol{z}}\left[\psi(z_1)\cdot\mathbb{I}\left[z_1>z_2\right]\right]+2\cdot\mathbb{E}_{\boldsymbol{z}}\left[\psi(z_1)\cdot\mathbb{I}\left[z_2>z_1\right]\right] \\
&= 2\cdot\mathbb{E}_{z}\left[\psi(z)\right].
\end{aligned}$$

Then, we have that

$$\begin{aligned}
\frac{d}{d\xi}\ \mathbb{E}_{\bar{\boldsymbol{x}},a}\left[\tilde{u}^{P}_{i^*}(\phi)\right]\Big|_{\xi=0} &= \frac{d}{d\xi}\mathbb{E}_{\boldsymbol{z},a}\left[D(\phi)|z_1>z_2\right]\Big|_{\xi=0} \\
&= \frac{1-\cos\theta}{\tau\sqrt{\pi}}\cdot\mathbb{E}_{z}\left[z\psi(z)\right]-\mathbb{E}_{z}\left[\psi(z)\right] \\
&= \mathbb{E}_{z}\left[\left(\frac{(1-\cos\theta)\cdot z}{\tau\sqrt{\pi}}-1\right)\cdot\frac{d\phi(z\ ;\xi)}{d\xi}\right]\Bigg|_{\xi=0}.
\end{aligned}$$

□

B. Analysis of Medicare Pay-for-Performance Policies

Proof of Proposition 2 The HRRP mechanism is $\phi_0=\alpha\cdot P_i$, where we define $P_i=\xi\cdot\max\left[\min\left(z_i,0\right),-\xi\right]$ as the penalty amount for treatment i. Let i^* be the agent's optimal action under cap ξ and i^o be the agent's optimal action under cap $\xi+\delta$ (where we take $\delta>0$).

We can verify that for each treatment choice i,

$$\frac{dP_i(\xi)}{d\xi}=-\xi\cdot\mathbb{I}[z_i<-\xi]\le 0.$$

Using this result, we have the following inequalities by definition of i^o:

$$\begin{aligned}
u^A_{i^o}+P_{i^o}(\xi+\delta) &\le u^A_{i^o}+P_{i^o}(\xi) \\
&\le u^A_{i^*}+P_{i^*}(\xi).
\end{aligned}$$

Taking $\delta \to 0$, we get $\frac{du^A_{i^*}}{d\xi} \leq 0$. Thus, the agent's utility is monotonically decreasing in the cap ξ.

Now, we consider the principal's utility. If the agent's action stays the same for the caps ξ and $\xi + \delta$, i.e., $i^* = i^o$, then the principal's utility clearly increases as P_{i^*} is decreasing in ξ. On the other hand, if $i^* \neq i^o$, we have the following inequalities by definition of i^o and i^*:

$$\begin{aligned} u^A_{i^*} + P_{i^*}(\xi+\delta) &\leq u^A_{i^o} + P_{i^o}(\xi+\delta) \\ &\leq u^A_{i^*} + P_{i^*}(\xi) - P_{i^o}(\xi) + P_{i^o}(\xi+\delta)\,, \\ P_{i^*}(\xi+\delta) - P_{i^*}(\xi) &\leq P_{i^o}(\xi+\delta) - P_{i^o}(\xi)\,. \end{aligned}$$

Simplifying the expression and taking $\delta \to 0$, we get

$$\frac{dP_{i^*}}{d\xi} \leq \frac{dP_{i^o}}{d\xi}\,.$$

This in turn implies that $z_{i^o} \geq z_{i^*}$, i.e. the performance metric is either the same or improved under the new treatment choice i^o. Since $\xi < 1$, we can verify $\frac{d}{dz_i}\tilde{u}^P_i \geq 1 - \xi > 0$, i.e. the principal's utility increases overall if the performance metric is improved. Thus, whether the treatment choice changes or not, the principal's utility is increasing in the cap ξ.

□

Proof of Proposition 3 We apply Proposition 1 to both mechanisms ϕ_0 and ϕ_1. First, we evaluate

$$\begin{aligned} \left.\frac{d\phi_0}{d\xi}\right|_{\xi=0} &= -\xi \cdot \mathbb{I}\left[z < -\xi\right]\Big|_{\xi=0} \\ &= -\xi \cdot \mathbb{I}\left[z < 0\right]\,. \end{aligned}$$

Then, letting f be the pdf of $z \sim \mathcal{N}(0, \sigma^2)$, we can compute

$$\begin{aligned} \mathbb{E}_z\left[\left(\frac{(1-\cos\theta)\cdot z}{\tau\sqrt{\pi}} - 1\right)\cdot \frac{d\phi_0(z\;;\xi)}{d\xi}\right]\Bigg|_{\xi=0} &= -\frac{\xi\cdot(1-\cos\theta)}{\tau\sqrt{\pi}}\int_{-\infty}^0 zf(z)\cdot dz + \xi\,\Pr\left[z<0\right] \\ &= \frac{\xi}{2}\cdot\left(\frac{\sigma\sqrt{2}}{\tau\pi} + 1\right)\,. \end{aligned}$$

Analogously, consider mechanism ϕ_1. First, we can re-write

$$\phi_1 = \xi \cdot \max\left[\min\left[z_i, 0\right], -\xi\right] + \xi \cdot \min\left[\max\left[z_i, 0\right], \xi\right]\,,$$

which yields

$$\left.\frac{d\phi_1}{d\xi}\right|_{\xi=0} = -\xi \cdot \mathbb{I}\left[z<0\right] + \xi \cdot \mathbb{I}\left[z>0\right]\,.$$

Then, we can compute

$$\begin{aligned} \mathbb{E}_z\left[\left(\frac{z}{\tau\sqrt{\pi}} - 1\right)\cdot\frac{d\phi_1(z\;;\xi)}{d\xi}\right]\Bigg|_{\xi=0} &= \frac{\xi}{\tau\sqrt{\pi}}\left(-\int_{-\infty}^0 zf(z)\cdot dz + \int_0^\infty zf(z)\cdot dz\right) + \xi\left(\Pr\left[z<0\right] - \Pr\left[z>0\right]\right) \\ &= \frac{2\xi}{\tau\sqrt{\pi}}\int_0^\infty zf(z)\cdot dz \\ &= \frac{\xi\sigma\sqrt{2}}{\tau\pi}\,. \end{aligned}$$

Simplifying, we find that ϕ_1 is the optimal small-incentive mechanism if and only if

$$\frac{\tau}{\sigma} < \frac{\sqrt{2}}{\pi}\,.$$

□

Proof of Proposition 4 Once again, we apply Proposition 1 to both mechanisms ϕ_0 and ϕ_1. First, consider mechanism ϕ_0:

$$\left.\frac{d\phi_0}{d\xi}\right|_{\xi=0} = -\xi \cdot \mathbb{I}\left[z < 0\right].$$

Thus, $\left.\frac{d\phi_0}{d\xi}\right|_{\xi=0}$ is the same for both the penalty-only HRRP mechanism and the single-threshold HACRP mechanism. Thus, we can re-use the expression from the proof of Proposition 3

$$\mathbb{E}_z\left[\left(\frac{(1-\cos\theta)\cdot z}{\tau\sqrt{\pi}} - 1\right)\cdot\frac{d\phi_0(z\ ;\xi)}{d\xi}\right]\Bigg|_{\xi=0} = \frac{\xi\,(1-\cos\theta)}{2}\cdot\left(\frac{\sigma\sqrt{2}}{\tau\pi}+1\right).$$

Next, we evaluate mechanism ϕ_1. Note that

$$\left.\frac{d\phi_1}{d\xi}\right|_{\xi=0} = \frac{\xi}{\sigma}\cdot\max\left[\min\left(z_i - \sigma/2, 0\right), -\sigma\right].$$

Then, we can compute

$$\begin{aligned}
\mathbb{E}_z\left[\left(\frac{z}{\tau\sqrt{\pi}} - 1\right)\cdot\frac{d\phi_1(z\ ;\xi)}{d\xi}\right]\Bigg|_{\xi=0} &= \frac{\xi}{\tau\sqrt{\pi}}\left(-\int_{-\infty}^{-\sigma/2} zf(z)\cdot dz + \int_{-\sigma/2}^{\sigma/2}\frac{z-\sigma/2}{\sigma}zf(z)\cdot dz\right)\\
&\quad + \xi\left(\Pr\left[z < -\sigma/2\right] - \int_{-\sigma/2}^{\sigma/2}\frac{z-\sigma/2}{\sigma}f(z)\cdot dz\right)\\
&= \frac{\xi}{\tau\sqrt{\pi}}\left(-\frac{2\sigma}{e^{1/8}\sqrt{2\pi}} + \sigma\cdot\mathrm{erf}\left(\frac{1}{2\sqrt{2}}\right)\right) + \frac{\xi}{2}\left(\mathrm{erfc}\left(\frac{1}{2\sqrt{2}}\right) + \mathrm{erf}\left(\frac{1}{2\sqrt{2}}\right)\right)\\
&= \frac{\xi\sigma}{\tau\sqrt{\pi}}\left(-\frac{2}{e^{1/8}\sqrt{2\pi}} + \mathrm{erf}\left(\frac{1}{2\sqrt{2}}\right)\right) + \frac{\xi}{2}.
\end{aligned}$$

Simplifying, we find that

$$\begin{aligned}
\mathbb{E}_z\left[\left(\frac{z}{\tau\sqrt{\pi}} - 1\right)\cdot\left(\frac{d\phi_0(z\ ;\xi)}{d\xi} - \frac{d\phi_1(z\ ;\xi)}{d\xi}\right)\right]\Bigg|_{\xi=0} &= \frac{\xi\sigma}{\tau\pi\sqrt{2}}\cdot\left(1 - \frac{1}{\sqrt{2\pi}}\left(-\frac{2}{e^{1/8}\sqrt{2\pi}} + \mathrm{erf}\left(\frac{1}{2\sqrt{2}}\right)\right)\right)\\
&> 0,
\end{aligned}$$

since we can verify numerically that

$$-\frac{2}{e^{1/8}\sqrt{2\pi}} + \mathrm{erf}\left(\frac{1}{2\sqrt{2}}\right) < 0.$$

Thus, ϕ_0 is always the optimal small-incentive mechanism.

□

C. Optimal Linear Mechanisms

In this section, we consider the case where the treatment parameters β_i are unknown to the principal so we directly model the treatment-specific outcomes as random variables, i.e., $\boldsymbol{z} \sim \mathcal{N}\left(0, \sigma^2 I\right)$.

Proof of Proposition 5 Recall that

$$\tilde{u}_i^A\left(\phi_\xi\right) = \rho\left(z_0 + z_i\right) + a_i + \xi\cdot z_i.$$

For ease of notation, we define

$$\begin{aligned}
y_i &= \tilde{u}_i^A\left(\phi_\xi\right) - \rho\cdot z_0\\
&= \left(\rho+\xi\right)\cdot z_i + a_i,
\end{aligned}$$

so we have $i^* = \arg\max_i \tilde{u}_i^A(\phi_\xi) = \arg\max_i y_i$. Since z_i, a_i are independent normal random variables, we know

$$y_i \sim \mathcal{N}\left(0, (\rho+\xi)^2\sigma^2 + \tau^2\right).$$

We want to compute the efficiency function

$$\begin{aligned}\eta(\xi) &= \frac{\mathbb{E}_{\bar{\boldsymbol{z}},a}\left[\tilde{u}_{i^*}^P(\phi_\xi)\right] - \mathbb{E}_{\bar{\boldsymbol{z}}}\left[z_0\right]}{\mathbb{E}_{\bar{\boldsymbol{z}}}\left[\max_i u_i^P\right] - \mathbb{E}_{\bar{\boldsymbol{z}}}\left[z_0\right]} \\ &= \frac{(1-\xi)\cdot\mathbb{E}_{\boldsymbol{z},a}\left[z_{i^*}\right]}{\mathbb{E}_{\boldsymbol{z}}\left[\max_i z_i\right]}.\end{aligned}$$

We first evaluate the numerator. Without loss of generality, we can take $i^* = 1$ and write

$$\begin{aligned}\mathbb{E}_{\boldsymbol{z},a}\left[z_{i^*}\right] &= \mathbb{E}_{\gamma\sim\max_{i\in[n]} y_i}\left[\mathbb{E}_{\boldsymbol{z},a}\left[z_1 \mid y_1 = \theta, y_2 \ldots y_n \le \gamma\right]\right] \\ &= \mathbb{E}_{\gamma\sim\max_{i\in[n]} y_i}\left[\mathbb{E}_{\boldsymbol{z},a}\left[z_1 \mid y_1 = \gamma\right]\right],\end{aligned}$$

where γ is a random variable distributed as the maximum of n i.i.d. draws from $\mathcal{N}(0, (\rho+\xi)^2\sigma^2+\tau^2)$. Note that the first equality follows from symmetry over i and the second follows from the independence of z_1 with respect to $\{y_j\}_{j\neq 1}$. The pair (z_i, y_i) is clearly a jointly Gaussian random variable, satisfying

$$\begin{bmatrix} z_i \\ y_i \end{bmatrix} \sim \mathcal{N}\left(0,\ \sigma^2\begin{bmatrix} 1 & \rho+\xi \\ \rho+\xi & (\rho+\xi)^2 + \tau^2/\sigma^2 \end{bmatrix}\right).$$

Then, we can evaluate

$$\begin{aligned}\mathbb{E}_{\boldsymbol{z},a}\left[z_1 \mid y_1 = \gamma\right] &= \mathbb{E}_{\boldsymbol{z}}\left[z_1\right] + \frac{\operatorname{Cov}(z_1, y_1)}{\operatorname{Var}(y_1)}\left(\theta - \mathbb{E}_{\boldsymbol{z},a}\left[y_1\right]\right) \\ &= \frac{(\rho+\xi)\cdot\sigma^2}{(\rho+\xi)^2\cdot\sigma^2+\tau^2}\cdot\gamma.\end{aligned}$$

Thus, we find the numerator of $\eta(\xi)$ is

$$\begin{aligned}(1-\xi)\cdot\mathbb{E}_{\boldsymbol{z},a}\left[z_{i^*}\right] &= (1-\xi)\cdot\frac{(\rho+\xi)\cdot\sigma^2}{(\rho+\xi)^2\sigma^2+\tau^2}\cdot\mathbb{E}_{\boldsymbol{z},a}\left[\max_{i\in[n]} y^i\right] \\ &= (1-\xi)\cdot\frac{(\rho+\xi)\cdot\sigma^2}{\sqrt{(\rho+\xi)^2\cdot\sigma^2+\tau^2}}\cdot\bar{\gamma}_n.\end{aligned}$$

where we have defined $\bar{\gamma}_n$ to be the expected value of the maximum of n i.i.d. standard normal variables (distributed as $\mathcal{N}(0,1)$). Note that we have the following closed form expression

$$\bar{\gamma}_n \equiv \int_{-\infty}^{\infty} \frac{y}{\sqrt{2\pi}}\exp\left[-y^2/2\right]\cdot\Phi(y)^{n-1}dy,$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution $\mathcal{N}(0,1)$.

Now, we evaluate the denominator of $\eta(\xi)$. Since the z_i are i.i.d. draws from $\mathcal{N}(0,\sigma^2)$, we have

$$\mathbb{E}_{\boldsymbol{z}}\left[\max_{i\in[n]} z_i\right] = \sigma\bar{\gamma}_n.$$

Thus, we can compute

$$\begin{aligned}\eta(\xi) &= \frac{(1-\xi)\cdot\frac{(\rho+\xi)\cdot\sigma^2}{\sqrt{(\rho+\xi)^2\sigma^2+\tau^2}}\cdot\bar{\gamma}_n}{\sigma\bar{\gamma}_n} \\ &= (1-\xi)\cdot\frac{(\rho+\xi)\cdot\sigma}{\sqrt{(\rho+\xi)^2\sigma^2+\tau^2}} \\ &= (1-\xi)\cdot\left(1+\frac{\tau^2}{(\rho+\xi)^2\cdot\sigma^2}\right)^{-1/2}.\end{aligned}$$

Substituting $M = \tau/\sigma$ yields the result.

□

Proof of Proposition 6 We will first show that $\eta(\xi)$ has a unique maximizer ξ^* in the restricted interval $[0,1]$. We begin by writing

$$\eta(\xi) = (1-\xi)\cdot g(\xi)\,.$$

where we have defined

$$g(\xi) = \left(1 + \frac{M^2}{(\rho+\xi)^2}\right)^{-1/2},$$

and $M = \tau/\sigma$. Note that we can write

$$\begin{aligned} g'(\xi) &= \frac{M^2}{(\rho+\xi)^3}\cdot\left(1+\frac{M^2}{(\rho+\xi)^2}\right)^{-3/2} \\ &= \frac{M^2}{(\rho+\xi)^3}\cdot g(\xi)^3\,, \end{aligned}$$

so

$$\eta'(\xi) = \frac{M^2(1-\xi)}{(\rho+\xi)^3}\cdot g(\xi)^3 - g(\xi)\,.$$

The optimal incentive (if it does not lie at the endpoints) satisfies $\eta'(\xi^*) = 0$. Note that $g(\xi) > 0$ always, so we can simplify:

$$\begin{aligned} \frac{M^2(1-\xi^*)}{(\rho+\xi^*)^3} - g(\xi^*)^{-2} &= 0\,, \\ \frac{M^2(1-\xi^*)}{(\rho+\xi^*)^3} - \left(1+\frac{M^2}{(\rho+\xi^*)^2}\right) &= 0\,, \\ (\rho+\xi^*)^3 + (\rho+2\xi^*-1)M^2 &= 0\,. \end{aligned}$$

Define the cubic polynomial

$$p(x) = (x+\rho)^3 + 2M^2x + (\rho-1)\,M^2\,.$$

By definition, ξ^* is a zero of $p(x)$. Additionally, note that

$$\frac{dp}{dx} = 3\,(x+\rho)^2 + 2M^2 > 0\,,$$

for every x, ρ, M. This shows that $p(x)$ is strictly increasing on the real line and therefore has only one zero. Thus, the optimal incentive ξ^* is either the unique maximizer of $p(x)$, or one of the endpoints $\{0,1\}$.

Next, we can surmise whether the P4P mechanism increases efficiency over the baseline setting if $\eta'(\xi) > 0$, i.e., the efficiency $\eta(\xi)$ is increasing when evaluated at $\xi = 0$ so it must be that $\xi^* > 0$.

$$\eta'(0) = \frac{1}{\sqrt{1+M^2/\rho^2}}\cdot\left(\frac{1}{\rho+\rho^3/M^2} - 1\right) > 0\,.$$

Re-arranging terms, we can write this condition as

$$\rho^3 + M^2\,(\rho-1) < 0\,.$$

□

Proof of Proposition 7 Recall that ξ^* satisfies

$$(\rho+\xi^*)^3+(\rho+2\xi^*-1)\cdot M^2 \ =0\,.$$

We can implicitly differentiate with respect to ρ and simplify:

$$\begin{aligned}
&3\,(\rho+\xi^*)^2\cdot\left(1+\frac{d\xi^*}{d\rho}\right)+\left(1+2\frac{d\xi^*}{d\rho}\right)M^2=0\,,\\
&\frac{d\xi^*}{d\rho}=-\frac{3\,(\rho+\xi^*)^2+M^2}{3\,(\rho+\xi^*)^2+2M^2}\,.
\end{aligned}$$

Thus, we have

$$-1<\frac{d\xi^*}{d\rho}<-\frac{1}{2}\,.$$

Next, we implicitly differentiate with respect to M:

$$3\,(\rho+\xi^*)^2\cdot\frac{d\xi^*}{dM}+2M\cdot(\rho+2\xi^*-1)+2M^2\cdot\frac{d\xi^*}{dM}=0\,.$$

Simplifying yields

$$\begin{aligned}
\frac{d\xi^*}{dM}&=-\frac{2M\cdot(\rho+2\xi^*-1)}{3\,(\rho+\xi^*)^2+2M^2}\\
&=\frac{2}{M}\cdot\frac{(\rho+\xi^*)^3}{3\,(\rho+\xi^*)^2+2M^2}\,,
\end{aligned}$$

where the last equality follows from the original polynomial equation satisfied by ξ^*. Thus, $\operatorname{sgn}\left[\frac{d\xi^*}{dM}\right]=\operatorname{sgn}\left[\rho+\xi^*\right]$. We can re-write the polynomial equation satisfied by ξ^* as follows:

$$x^3+2M^2x=(\rho+1)\,M^2\,,$$

where $x=\rho+\xi^*$. Note that $(\rho+1)\geq 0$ by definition so it must be the case that $x>0$. Thus, we have that

$$\frac{d\xi^*}{dM}>0\,.$$

□

Acknowledgments

Bastani acknowledges the support of NSF Graduate Research Fellowship DGE-114747; Bayati acknowledges support of NSF grants CCF:1216011 and CAREER award CMMI:1554140; Braverman acknowledges support of NSF grants CCF:1216011, CAREER grant CCF-1149888, CCF-1525342, and a Packard Fellowship in Science and Engineering.

This paper has benefitted from valuable feedback from Ken Moon, Andrzej Skrzypacz, Takuo Sugaya, Seungjin Whang, and Stefanos Zenios, as well as various seminar participants.

References

Adida, Elodie, Hamed Mamani, Shima Nassiri. 2016. Bundled payment vs. fee-for-service: Impact of payment scheme on performance. *Management Science* .

Anderson, Cheryl, BV Deepak, Yaw Amoateng-Adjepong, Stuart Zarich. 2005. Benefits of comprehensive inpatient education and discharge planning combined with outpatient support in elderly patients with congestive heart failure. *Congestive Heart Failure* **11**(6) 315–321.

Andritsos, Dimitrios, Christopher S Tang. 2015. Incentive programs for reducing readmissions when patient care is co-produced. *Working Paper* .

Arrow, K. 1963. Uncertainty and the welfare economics of medical care. *The American Economic Review* **53**(5).

Aswani, Anil, Zuo-Jun Max Shen, Auyon Siddiq. 2016. Data-driven incentive design in the medicare shared savings program. *Working Paper* .

Bayati, Mohsen, Mark Braverman, Michael Gillam, Karen Mack, George Ruiz, Mark Smith, Eric Horvitz. 2014. Data-driven decisions for reducing readmissions for heart failure: General methodology and case study .

Berenholtz, Sean M, Peter J Pronovost, Pamela A Lipsett, Deborah Hobson, Karen Earsing, Jason E Farley, Shelley Milanovich, Elizabeth Garrett-Mayer, Bradford D Winters, Haya R Rubin, et al. 2004. Eliminating catheter-related bloodstream infections in the intensive care unit*. *Critical care medicine* **32**(10) 2014–2020.

Boccuti, Cristina, Giselle Casillas. 2015. Aiming for fewer hospital u-turns: The medicare hospital readmission reduction program. *Kaiser Family Foundation Issue Brief* .

Bolton, Patrick, Mathias Dewatripont. 2004. *Contract Theory*. The MIT Press.

Bundorf, M. Kate, Jonathan D. Levin, Neale Mahoney. 2008. Pricing and welfare in health plan choice. *NBER Working Paper* (14153).

Carroll, Gabriel. 2015. Robustness and linear contracts. *The American Economic Review* **105**(2) 536–563.

Chan, David C. 2012. Organizational structure and moral hazard among emergency department physicians. *Working paper* .

CMS. 2014. Hospital-acquired condition (hac) reduction program URL `"https://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/AcuteInpatientPPS/HAC-Reduction-Program.html"`.

CMS. 2015a. Fiscal year (fy) 2016 results for the cms hospital-acquired conditions (hac) reduction program URL `"https://www.cms.gov/Newsroom/MediaReleaseDatabase/Fact-sheets/2015-Fact-sheets-items/2015-12-10-2.html"`.

CMS. 2015b. Readmissions reduction program URL `"https://www.cms.gov/medicare/medicare-fee-for-service-payment/acuteinpatientpps/readmissions-reduction-program.html"`.

Cutler, David. 2004. *Your Money Or Your Life: Strong Medicine for America's Health Care System*. Oxford University Press.

Davis, K., C. Schoen, K. Stremikis, C. Fund. 2010. *Mirror, mirror on the wall: How the performance of the US health care system compares internationally: 2010 update*. Commonwealth Fund.

Dudley, R Adams, Anne Frolich, David L Robinowitz, Jason A Talavera, Peter Broadhead, Harold S Luft. 2004. Strategies to support quality-based purchasing. *Agency for Healthcare Research and Quality* .

Fong, K. 2009. Evaluating skilled experts: Optimal scoring rules for surgeons. *Working Paper* .

Frankel, Alex. 2010. Aligned delegation To appear in American Economic Review.

Fuloria, Prashant C., Stefanos A. Zenios. 2001. Outcomes-adjusted reimbursement in a health-care delivery system. *Management Science* **47**(6) 735–751.

Green, Linda V. 2012. Operations management forum-the vital role of operations analysis in improving healthcare delivery. *Manufacturing & Service Operations Management* **14**(4) 488–494.

Guo, Pengfei, Christopher S Tang, Yulan Wang, Ming Zhao. 2016. The impact of reimbursement policy on patient welfare, readmission rate and waiting time in a public healthcare system: Fee-for-service vs. bundled payment. *Working Paper* .

Gupta, Diwakar, Mili Mehrotra. 2015. Bundled payments for healthcare services: Proposer selection and information sharing. *Operations Research* **63**(4) 772–788.

HHS. 2015. Better, smarter, healthier: In historic announcement, hhs sets clear goals and timeline for shifting medicare reimbursements from volume to value.

Holmstrom, B., P. Milgrom. 1987. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica* **55**(2).

Hölmstrom, Bengt. 1979. Moral hazard and observability. *The Bell Journal of Economics* 74–91.

Institute of Medicine, Consensus Report. 2006. Rewarding provider performance : aligning incentives in medicare. Committee on Redesigning Health Insurance Performance Measures, Payment, and Performance Improvement Programs, Board on Health Care Services.

Jha, Ashish K. 2013. Time to get serious about pay for performance. *JAMA* **309**(4) 347–348.

Jiang, Houyuan, Zhan Pang, Sergei Savin. 2012. Performance-based contracts for outpatient medical services. *Manufacturing & Service Operations Management* **14**(4) 654–669.

Joynt, Karen E., Jose F. Figueroa, E. John Orav, Ashish K. Jha. 2016. Opinions on the hospital readmission reduction program: Results of a national survey of hospital leaders URL `"http://www.ajmc.com/journals/issue/2016/2016-vol22-n8/opinions-on-the-hospital-readmission-reduction-program-results-of-a-national-survey-of-hospital-le`

Kahneman, Daniel, Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society* 263–291.

Klevens, R Monina, Melissa A Morrison, Joelle Nadle, Susan Petit, Ken Gershman, Susan Ray, Lee H Harrison, Ruth Lynfield, Ghinwa Dumyati, John M Townes, et al. 2007. Invasive methicillin-resistant staphylococcus aureus infections in the united states. *Jama* **298**(15) 1763–1771.

Laffont, Martimort. 2001. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Ma, Albert C., Henry Y. Mak. 2012. Information disclosure and the equivalence of prospective payment and cost reimbursement. *Working Paper* .

Maizel, Samuel R, Craig B Garner. 2014. The poor get poorer: The fate of california's hospitals under the affordable care act. *California Health Law News* **32**(3) 2015.

Mirrlees, James A. 1999. The theory of moral hazard and unobservable behaviour: Part i. *The Review of Economic Studies* **66**(1) 3–21.

Morreim, E Haavi. 1991. Gaming the system: dodging the rules, ruling the dodgers. *Archives of Internal Medicine* **151**(3) 443–447.

Porter, Michael E, Elizabeth Olmsted Teisberg. 2006. *Redefining health care: creating value-based competition on results*. Harvard Business Press.

Rizzo, Ellie. 2013. 6 stats on the cost of readmission for cms-tracked conditions. *"Becker's Infection Control & Clinical Quality"* URL `"http://www.beckershospitalreview.com/quality/6-stats-on-the-cost-of-readmission-for-cms-tracked-conditions.html"`.

Rosenthal, Meredith, R. Adams Dudley. 2006. Pay for performance: a decision guide for purchasers .

Rosenthal, Meredith, R. Adams Dudley. 2007. Pay-for-performance: Will the latest payment trend improve care? *JAMA* **297**(7) 740–744. doi:10.1001/jama.297.7.740. URL `+http://dx.doi.org/10.1001/jama.297.7.740`.

Scott, R Douglas, et al. 2009. The direct medical costs of healthcare-associated infections in us hospitals and the benefits of prevention .

Suter, Lisa G, Shu-Xia Li, Jacqueline N Grady, Zhenqiu Lin, Yongfei Wang, Kanchana R Bhat, Dima Turkmani, Steven B Spivack, Peter K Lindenauer, Angela R Merrill, et al. 2014. National patterns of risk-standardized mortality and readmission after hospitalization for acute myocardial infarction, heart failure, and pneumonia: update on publicly reported outcomes measures based on the 2013 release. *Journal of general internal medicine* **29**(10) 1333–1340.

Truffer, C.J., S. Keehan, S. Smith, J. Cylus, A. Sisko, J.A. Poisal, J. Lizonitz, M.K. Clemens. 2010. Health spending projections through 2019: the recessions impact continues. *Health Affairs* **29**(3) 522–529.

Werner, Rachel M, G Caleb Alexander, Angela Fagerlin, Peter A Ubel. 2002. The hassle factor: What motivates physicians to manipulate reimbursement rules? *Archives of Internal Medicine* **162**(10) 1134–1139.

Yaesoubi, Reza, Stephen D. Roberts. 2011. Payment contracts in a preventive health care system: A perspective from operations management. *Journal of Health Economics* **30** 1188–1196.

Yamin, Dan, Arieh Gavious. 2013. Incentives' effect in influenza vaccination policy. *Management Science* **59**(12) 2667–2686.

Zhang, Dennis J, Itai Gurvich, Jan A Van Mieghem, Eric Park, Robert S Young, Mark V Williams. 2016. Hospital readmissions reduction program: An economic and operational analysis. *Management Science* .