Gaming and Strategic Opacity in Incentive Provision*

Florian Ederer† — Yale University

Richard Holden‡ — UNSW

Margaret Meyer§ — Oxford and CEPR

July 6, 2017

Abstract

It is often suggested that incentive schemes under moral hazard can be gamed by an agent with superior knowledge of the environment and that deliberate lack of transparency about the incentive scheme can reduce gaming. We formally investigate these arguments in a two-task moral hazard model in which the agent is privately informed about which task is less costly for him. We examine a simple class of incentive schemes that are "opaque" in that they make the agent uncertain ex ante about the incentive coefficients in the linear payment rule. Relative to transparent menus of linear contracts, these opaque schemes induce more balanced efforts, but they also impose more risk on the agent per unit of aggregate effort induced. We identify specific settings in which optimally designed opaque schemes not only strictly dominate the best transparent menu but also eliminate the efficiency losses from the agent's hidden information. Opaque schemes are more likely to be preferred to transparent ones when (i) efforts on the tasks are highly complementary for the principal; (ii) the agent's privately known preference between the tasks is weak; (iii) the agent's risk aversion is significant; and (iv) the errors in measuring performance have large correlation or small variance. (*JEL* D86, D21, L22)

*We wish to thank David Martimort (the editor) and two anonymous referees for comments and suggestions that greatly improved the paper. We are also grateful to Philippe Aghion, Sushil Bikhchandani, Simon Board, Patrick Bolton, Jeremy Bulow, Vince Crawford, Kate Doornik, Guido Friebel, Robert Gibbons, Edward Glaeser, Oliver Hart, Bengt Holmström, Ian Jewitt, Navin Kartik, Paul Klemperer, Shmuel Leshem, Phillip Leslie, Steven Lippman, Steven Matthews, John Moore, Andy Newman, In-Uck Park, Charles Roddie, Jesse Shapiro, Ferdinand von Siemens, Jason Snyder, Alexander Stremitzer, Jeroen Swinkels, and Jean Tirole for helpful discussions and suggestions, as well as to seminar participants at Bonn, Boston University, Bristol, Chicago, Columbia, Frankfurt, Georgia Tech, Gerzensee, HBS, LSE, MIT, Michigan, Oxford, Queen's, Rotterdam, UCL, UCLA, USC, and Yale. Holden acknowledges the Australian Research Council Future Fellowship FT130101159 and the Centre for International Financial Regulation at UNSW for financial support.

†Yale School of Management, Yale University, 165 Whitney Avenue, New Haven, CT 06511, U.S.A., florian.ederer@yale.edu.

‡UNSW Business School, School of Economics, Sydney, NSW 2052, Australia, richard.holden@unsw.edu.au.

§Nuffield College and Department of Economics, Oxford University, OX1 1NF, United Kingdom, margaret.meyer@nuffield.ox.ac.uk.

1 Introduction

A fundamental consideration in designing incentive schemes is the possibility of *gaming*: exploitation of an incentive scheme by an agent for his own self-interest to the detriment of the objectives of the incentive designer. Gaming can take numerous forms, among them 1) diversion of effort away from activities that are socially valuable but difficult to measure and reward, towards activities that are easily measured and rewarded; 2) exploitation of the rules of classification to improve apparent, though not actual, performance; and 3) distortion of choices about timing to exploit temporarily high monetary rewards even when socially efficient choices have not changed. Evidence of the first type of gaming is provided by Burgess, Propper, Ratto, and Tominey (2012) and Carrell and West (2010), of the second type by Gravelle, Sutton, and Ma (2010), and of the third type by Oyer (1998), Larkin (2014), and Forbes, Lederman, and Tombe (2013).[1] The costs of gaming are exacerbated when the agent has superior knowledge of the environment: this makes the form and extent of gaming harder to predict and hence harder to deter.

It has been suggested that lack of transparency—deliberate opacity about the criteria upon which rewards will be based and/or how heavily these criteria will be weighted—can help deter gaming. This idea has a long intellectual history. It dates back at least to Bentham (1830), who argued that deliberate opacity about the content of civil service selection tests would lead to the "maximization of the inducement afforded to exertion on the part of learners, by impossibilizing the knowledge as to what part the field of exercise the trial will be applied to, and thence making aptitude of equal necessity in relation to every part".[2]

More recently, responding to documented gaming of the highly transparent incentive schemes which score National Health Service organizations in England according to published lists of precisely defined performance indicators, Bevan and Hood (2004) argued in a *British Medical Journal* editorial, "What is needed are ways of limiting gaming. And one way of doing so is to introduce more randomness in the assessment of performance, at the expense of transparency" (p. 598). They invoke the "analogy [...] with the use of unseen examinations, where the unpredictability of what the questions will be means that it is safest for students to cover the syllabus" (p. 598). They reason that making it harder for hospitals to predict what performance measures will be used and how they will be weighted, coupled with hospitals' risk aversion, will reduce the hospitals' incentives for gaming. Similarly, Dranove, Kessler, McClellan, and Satterthwaite (2003) document that in the United States, report cards for hospitals "encourage providers to 'game' the system by avoiding sick

[1] Burgess et al (2012) and Gravelle et al (2010) study UK public sector organizations (an employment agency and the National Health Service, respectively), Carrell and West (2010) use data from postsecondary education, while Oyer (1998), Larkin (2013) and Forbes et al (2012) examine private sector organizations (salespeople and executives across various industries, enterprise software vendors, and airlines, respectively).

[2] Bentham, 1830/2005, Ch. IX,§16, Art 60.1.

patients or seeking healthy patients or both" and they argue that such gaming is facilitated by "risk-averse providers having better information about patients' conditions" (p. 556) than do the analysts who compile the report cards. They present evidence that the increased transparency of incentive schemes for physicians and hospitals provided by report cards increased gaming and even decreased patient and social welfare.[3]

The costs of transparency have also been discussed in the context of gaming, by law school deans, of the performance indicators used by *U.S. News* to produce its influential law school rankings. The ranking methodology is transparent and employs a *linear* scoring rule incorporating multiple performance indicators.[4] There is significant evidence that law schools deploy a range of strategies that exploit their informational advantage over *U.S. News* to increase their measured performance. Examples include cutting the number of full-time students to boost median LSAT scores and GPAs, creating make-work jobs for their own graduates to inflate the number in employment, and heavily advertising their faculty's scholarship to *U.S. News*.[5] Law scholars (e.g. Osler, 2010) have argued that greater opacity in the ranking methodology could mitigate gaming, and *U.S. News* has itself signaled its intention to move away from being "totally transparent about key methodology details".[6]

Finally, one view as to why courts often prefer standards—which are somewhat vague—to specific rules is that standards mitigate incentives for gaming. For example, Weisbach (2000) argues that vagueness can reduce gaming of taxation rules, and Scott and Triantis (2006) argue that vague standards in contracts can improve parties' incentives to fulfill the spirit of the contract rather than focusing on satisfying only the narrowly defined stipulations.

The examples discussed above suggest that "opacity" (i.e., lack of transparency) of incentive schemes can be beneficial in reducing gaming, especially when agents have superior knowledge of the environment, incentive designers care about multiple aspects of performance, and gaming takes the form of agents' focusing efforts on easily manipulable indicators. This line of argument is, however, incomplete. If agents are risk-averse, then the additional risk imposed by opaque schemes is per se unattractive to them.

[3]Relatedly, Google has experienced manipulation of its search results by some retailers. Although many retailers have been seeking greater transparency from Google about its search algorithm, Google has responded by moving in the direction of greater opacity to prevent manipulation (Structural Search Engine Optimization, Google Penalty Solutions, November 4, 2011, `http://www.re1y.com/blog/occupy-google-blog.html`). Motivated in part by this debate, Frankel and Kartik (2014) develop a signaling model of gaming in which the information conveyed by signals (e.g., prominence in search results) about agents' hidden characteristics (e.g., intrinsic relevance to the query) is "muddled" because agents are also privately informed about their gaming ability. Jehiel and Newman (2011) develop a dynamic model in which principals learn from agents' behavior about the possibilities for gaming and then choose whether to take costly measures to deter gaming.

[4]The weights in the scoring rule are quality perception (40%), selectivity (25%), placement success (20%) and faculty resources (15%) (U.S. News, March 11, 2013, `http://www.usnews.com/education/best-graduate-schools/top-law-schools/articles/2013/03/11/methodology-best-law-schools-rankings`).

[5]Law School Rankings Reviewed to Deter 'Gaming', Wall Street Journal, August 26, 2008, `http://www.usnews.com/education/blogs/college-rankings-blog/2010/05/20/us-news-takes-steps-to-stop-law-schools-from-manipulating-the-rankings`.

[6] U.S. News, May 20, 2010, `http://www.usnews.com/education/blogs/college-rankings-blog/2010/05/20/us-news-takes-steps-to-stop-law-schools-from-manipulating-the-rankings`.

Understanding when and why opaque schemes are used thus requires analyzing the tradeoff between their incentive benefits and their risk costs. The present paper provides such an analysis.

Our analysis incorporates three vital ingredients that are featured in all of our motivating examples: (i) the agent's superior information about the environment, (ii) the agent's risk aversion, and (iii) the incentive designer's need for the agent to choose a relatively balanced allocation of efforts across activities. This suite of ingredients (along with a contractual restriction to incentive schemes that are ex post linear) delivers two main messages. First, transparent incentive schemes, even when they involve menus, suffer dramatically from the problem of gaming by the agent. Second, opaque incentive schemes not only mitigate the problem of gaming but can generate a higher payoff for the principal.[7]

In our model, "opacity" corresponds to a lack of transparency about the weights on performance indicators that are used to determine rewards. Motivated by the examples discussed above, we build on Holmstrom and Milgrom's (1991) multi-task principal-agent model in which a risk-averse agent performs two tasks, which are substitutes in his cost-of-effort function, and receives compensation that is linear in his performance on each of the tasks. These linear contracts (which have been widely studied) are "transparent" in that the agent faces no uncertainty about the rate at which performance on each of the tasks is rewarded. The principal's benefit function is complementary in the agent's efforts on the two tasks; other things equal, she prefers to induce both types of agent to choose balanced efforts.[8] Into this familiar set-up, we introduce superior knowledge of the environment on the part of the agent. There are two types of agent, and only the agent knows which type he is. One type has a lower cost of effort on task 1, and the other has a lower cost of effort on task 2.[9]

The privately-informed agent games transparent incentive schemes by choosing effort allocations that are excessively (from an efficiency perspective) sensitive to his private information. In fact, we show that the agent's superior knowledge of his preferences makes it impossible for the principal, with transparent linear schemes, to induce both types of agent to exert positive efforts on both tasks, even when menus of contracts are used as screening devices. This is the sense in which transparent incentive schemes in our model suffer dramatically from the problem of gaming. One approach to mitigating gaming would be for the principal to design general (menus of) nonlinear compensation schedules. But such schedules can be

[7]The terms "opaque" and "transparent" may have alternative definitions in other contexts, but here, where we confine attention to compensation schedules that are ex post linear, an "opaque" incentive scheme will always be one that leaves the agent, when choosing efforts, uncertain about the incentive coefficients he will face, while a "transparent" scheme will be one under which the agent faces no such uncertainty.

[8]Our model, like Holmstrom and Milgrom's (1991), incorporates shocks to measured performance. These shocks are not essential for our two main messages, given our focus on contracts that are ex post linear. In fact, as discussed after Proposition 7, our findings about the benefits of opaque incentive schemes would be even stronger in the absence of such shocks. Nonetheless, it is natural to include them in the analysis; if the agent's efforts were directly observable by the principal, then the problem of moral hazard could be trivially solved by a so-called "forcing contract".

[9]The analysis would be very similar if the agent types differed with respect to the task on which they were more productive.

very complex to describe and difficult for agents to understand. Moreover, optimizing over general nonlinear contracts is difficult, especially when agents have hidden information.

Our approach is instead to explore a class of incentive schemes that is both simple and opaque. This class is simple in that, ex post, compensation is determined by one of two possible linear functions of performance measures, differing with respect to which performance measure is more highly rewarded. It is opaque in that, at the time the agent chooses his efforts, he does not know which of these two linear reward functions will be used. In the main body of the paper, we focus on one such simple, opaque scheme, which we term *ex ante randomization*. Under ex ante randomization, the principal, before the agent makes his effort choices, commits to randomizing uniformly between the two linear compensation schedules. Ex ante randomization encourages the risk-averse agent to choose relatively balanced efforts on the tasks in order to partially insure himself against the wage risk generated by the random choice of compensation schedule. The more unequal the weights on the performance measures in the two possible compensation schedules, the stronger the agent's incentive to self-insure and the more balanced his optimal efforts will be.[10]

The benefits of opaque incentive schemes in deterring gaming do, nevertheless, come at a cost: such schemes impose more risk on the agent. Given any incentive scheme involving ex ante randomization, we show that there exists a transparent contract that induces the same level of *aggregate* effort on the two tasks and imposes lower overall risk costs. Highlighting the importance of our three key model ingredients, we use this finding to prove that any opaque contract will be dominated by some transparent contract if (i) the agent has no private information about his preferences, or (ii) the agent's risk aversion is too weak for the opaque contract to induce him to choose positive efforts on both tasks, or (iii) the agents' efforts on the two tasks are not sufficiently complementary for the principal to make balanced efforts socially efficient.

Most importantly, we also identify three environments in which our simple opaque incentive schemes, with the relative weights on the performance measures chosen optimally, strictly dominate all transparent incentive schemes. In the first such setting, the agent has private information about his preferences but the magnitude of his preference across tasks is very small. The second is the case where the agent's risk aversion is very large and the variance of the shocks to outputs is very small. In the final setting, the correlation of the output shocks is very high. In each of these settings, the strict superiority of ex ante randomization over the best transparent scheme follows in the limit from the result that ex ante randomization allows the principal to

[10]In Section 7.1, we briefly discuss two other simple, opaque schemes, *interim randomization* and *ex post discretion*, which differ from ex ante randomization in the assumptions on the principal's powers of commitment. Ex post discretion is analyzed in detail in an earlier version of our paper (Ederer, Holden, and Meyer, 2014). All three such opaque schemes work in very similar ways. In particular, by making the risk-averse agent uncertain ex ante about the values of the incentive coefficients in the linear payment rule, they all provide an incentive for balancing efforts. Our findings, from the analysis of ex ante randomization, about the pros and cons of opacity are thus robust to alternative assumptions on the principal's commitment powers.

achieve a payoff arbitrarily close to what she could achieve in the absence of the agent's hidden information.

Though the results just described focus on limiting environments to prove analytically the strict dominance of optimally weighted ex ante randomization over all transparent menus, we also present more general findings about what characteristics of the environment increase the relative attractiveness of opaque schemes. We prove that as the agent becomes more risk-averse, holding the importance of risk aversion under transparent schemes fixed, the relative attractiveness of ex ante randomization increases, since the more balanced efforts chosen by the more risk-averse agent not only benefit the principal directly but also ensure lower overall risk costs. Furthermore, we show numerically that ex ante randomization is more likely to dominate the best transparent scheme when (i) the agent's privately known preference between tasks is weak, so even a small amount of uncertainty about the weights in the compensation schedule has a large effect on how balanced the agent's chosen efforts are, (ii) the agent's risk aversion is significant, so opaque schemes provide the agent with a powerful self-insurance motive for balancing efforts, (iii) efforts on the tasks are highly complementary for the principal, and (iv) the errors in measuring performance on the tasks have large correlation or small variance.

1.1 Related Literature

Our paper builds on the theoretical analyses of Holmstrom and Milgrom (1987, 1991). The first of these provides conditions in a dynamic moral hazard setting under which a linear contract is optimal. A key message of Holmstrom and Milgrom (1987) is that linear contracts are appealing because they are robust to limitations on the principal's knowledge of the contracting environment. Discussing Mirrlees's (1974) result that the first-best outcome in a hidden-action model can be approximated by a step-function (hence highly non-linear) incentive scheme, they argue "to construct the [Mirrlees] scheme, the principal requires very precise knowledge about the agent's preferences and beliefs, and about the technology he controls. The two-wage scheme performs ideally if the model's assumptions are precisely met, but can be made to perform quite poorly if small deviations in the assumptions [...] are introduced" (p. 305).[11] Motivated not only by these robustness arguments, but also by the simplicity and pervasiveness of linear contracts, we focus our analysis on compensation schedules in which, ex post, after all choices are made and random variables are realized, payments are linear functions of the performance measures.

Analyses of multi-task principal-agent models (e.g., Holmstrom and Milgrom (1991), Baker (1992)) have highlighted the inefficiencies resulting under linear contracts from an agent's ability to privately choose how to allocate his efforts across different activities. When efforts on different tasks are technological substitutes for

[11] Carroll (2015) also demonstrates an appealing robustness property of linear contracts. He shows that, in a static model with limited liability, when the principal knows some but not all of the actions available to the agent and evaluates contracts according to their worst-case performance, a linear contract is optimal.

the agent, an increase in incentives on one task will typically induce the agent not only to raise effort on that task but also to lower efforts on others, an effect termed the "effort-substitution problem". One consequence of the effort-substitution problem that is often emphasized is that to induce an agent to exert effort on tasks that are difficult to measure, it may be necessary for contracts to offer low-powered incentives on all tasks, even those that are easy to measure. The effort-substitution problem is present in our model, and we show that with transparent linear incentive schemes, the inefficiencies it generates are dramatically exacerbated when the agent is better informed than the principal about his cost function. Nevertheless, our focus is not on the implications of the effort-substitution problem for the optimal overall strength of incentives. Rather, we focus on how opaque incentive schemes can mitigate the costs of the effort-substitution problem by making the relative rewards for different tasks random. And our analysis of opaque incentive schemes focuses primarily on the optimal degree of uncertainty about relative rewards rather than on the optimal overall strength of incentives.

Like us, MacDonald and Marx (2001) analyze a principal-agent model with two tasks where the agent's efforts on the tasks are substitutes for the agent but complements for the principal, and where the agent is privately informed about his preferences. Because they restrict task outcomes to be binary, it is possible to solve for the optimal contract, and they show that the more complementary the tasks are for the principal, the more the optimal reward scheme makes successes on the tasks complementary for the agent. They do not consider ex ante randomization, and in fact, under their specific assumptions, it would have no power to mitigate gaming.[12] In our model, with a more general production technology, optimal nonlinear, non-separable contracts are prohibitively difficult to characterize, but at the same time, ex ante randomization over two linear schedules proves to be both a simple and a powerful tool for mitigating the excessive sensitivity of agents' effort allocations to their private information.

Randomization has, of course, been studied before in incentive provision. In general single-task hidden-action models allowing arbitrarily complex contracts, Gjesdal (1982) and Grossman and Hart (1983) show that exogenous randomization may be optimal, but only if the agent's risk tolerance varies with the level of effort he exerts. In our model, the agent's risk tolerance is independent of his effort level; the attractiveness of opaque incentives stems from their ability to mitigate the agency costs of *multi-task* incentive problems when compensation schedules are constrained to be ex post linear.

The potential benefits of exogenous randomization have also been explored in hidden-information models,

[12] In their model, ex ante randomization over which task to reward more highly would not generate a self-insurance motive for balancing efforts, even for a risk-averse agent. The reason is that, since effort affects the probability of good performance rather than the level of good performance, the marginal benefit to the agent of effort on a task would not be weighted by the agent's marginal utilities in the two events corresponding to the two possible compensation schedules. In our model, in contrast, the marginal benefit of effort on a task is weighted by the agent's marginal utilities. The difference in these marginal utilities is the source of the self-insurance motive for effort balance under ex ante randomization.

especially those studying the design of optimal tax schedules. Stiglitz (1982) and Pestieau, Possen, and Slutsky (1997), among others, have shown that randomization can facilitate the screening of privately-informed individuals and is especially effective when private information is multi-dimensional. In our hidden-action cum hidden-information setting, in contrast, ex ante randomization in fact eliminates the need for screening.

The costs and benefits of transparency in incentive design are also explored in Jehiel (2015) and Lazear (2006). Jehiel (2015) shows in an abstract moral hazard setup that a principal may gain by keeping agents uninformed about some aspects of the environment (e.g., how important specific tasks are). The benefits of suppressing information in relaxing incentive constraints can outweigh the costs of agents' less efficient adaptation of actions to the environment. Lazear (2006), in a model in which agents have no hidden information, explores high-stakes testing in education and the deterrence of speeding and terrorism, identifying conditions under which a lack of transparency can have beneficial incentive effects. In Lazear's analysis of testing, there is an exogenous restriction on the number of topics that can be tested, whereas in our model, even when all tasks can be measured and rewarded, we show that deliberate opacity about the weights in the incentive scheme can be desirable.[13]

The remainder of the paper proceeds as follows. Section 2 outlines our model. Section 3 studies transparent incentive schemes, while Section 4 analyzes opaque schemes. Section 5 identifies settings in which opaque schemes are dominated by transparent ones. Section 6, which is the heart of the paper, identifies environments in which optimally weighted opaque schemes dominate the best transparent one. Sections 7 and 8 contain extensions and concluding remarks. Proofs not provided in the text are in the appendix.

2 The Model

A principal (she) hires an agent (he) to perform a job for her. The agent's performance on the job has two distinct dimensions, which we term "tasks". Measured performance, x_j, on each task $j = 1, 2$ is verifiable and depends both on the effort devoted by the agent to that task, e_j, and on the realization of a random shock, ε_j. In particular, $x_j = e_j + \varepsilon_j$, where $(\varepsilon_1, \varepsilon_2)$ have a symmetric bivariate normal distribution with mean 0, variance σ^2, and covariance $\rho\sigma^2 \geq 0$. The efforts chosen by the agent are not observable by the principal.

As mentioned in the Introduction, our multi-task moral hazard model incorporates three key ingredients that are featured in all of our motivating examples. The first of these is that at the time of contracting, the

[13] The costs and benefits of transparency are also a focus of interest in international relations. Wikipedia defines the policy of "strategic ambiguity" as "the practice by a country of being intentionally ambiguous on certain aspects of its foreign policy [...]. It may be useful if the country has contrary foreign and domestic policy goals or if it wants to take advantage of risk aversion to abet a deterrence strategy." (`http://en.wikipedia.org/wiki/Policy_of_deliberate_ambiguity`). Multiple objectives of the principal and risk aversion of the agent are also important in our model in generating the beneficial incentive effects of opacity.

agent is better informed than the principal about his cost of exerting efforts. Specifically, with probability $\frac{1}{2}$, the agent's cost function is $c_1(e_1, e_2) = \frac{1}{2}(e_1 + \lambda e_2)^2$, in which case we will term him a type-1 agent, and with probability $\frac{1}{2}$ his cost function is $c_2(e_1, e_2) = \frac{1}{2}(\lambda e_1 + e_2)^2$, in which case he will be termed a type-2 agent. The parameter λ is common knowledge, and $\lambda \geq 1$. For each type of agent $i = 1, 2$, efforts are perfect substitutes: $\frac{\partial c_i / \partial e_1}{\partial c_i / \partial e_2}$ does not vary with (e_1, e_2).[14] Nevertheless, since $\lambda \geq 1$, the type-i agent has a preference for task i: the marginal cost of effort on task j ($j \neq i$) is λ times as large as that on task i.

The second key ingredient is the agent's risk aversion. We assume that both types of agent have an exponential von Neumann-Morgenstern utility function with coefficient of absolute risk aversion r, so the type-i agent's utility function is $U = -\exp\{-r(w - c_i(e_1, e_2))\}$, where w is the payment from the principal. The two types of agent are assumed to have the same level of reservation utility, which we normalize to zero in certainty-equivalent terms.

The third key feature of our model is that the agent's efforts on the tasks are complementary for the principal. We capture this by assuming that the principal's payoff, which consists of the benefit to her from the agent's efforts minus the payment to the agent, takes the following form:

$$\Pi = \frac{\delta \underline{e} + \overline{e}}{\delta + 1} - w,$$

where $\underline{e}$ is the smaller of the efforts on the two tasks, $\overline{e}$ is the larger of the efforts, and the parameter $\delta \in [1, \infty)$. Notice that as δ goes to ∞, the benefit to the principal goes to $\underline{e}$, so that the tasks are perfect complements for her. On the other hand, when $\delta = 1$, the principal's payoff is $\frac{1}{2}(\underline{e} + \overline{e})$, so that the tasks are perfect substitutes for her. When the agent chooses perfectly balanced efforts $\underline{e} = \overline{e} = e$, the principal's benefit is e, which is independent of δ.[15]

The relative size of δ and λ determines what allocation of effort across tasks would maximize social surplus. If $\delta > \lambda$, so the principal's desire for balanced efforts is stronger than the agent's preference across tasks, then the surplus-maximizing effort allocation involves both types of agent exerting equal effort on the two tasks. If, instead, $\delta < \lambda$, then in the socially efficient effort allocation, each type of agent focuses exclusively on his preferred task.

The principal's benefit, $\frac{\delta \underline{e} + \overline{e}}{\delta + 1}$, is assumed non-verifiable. Therefore, the only measures on which the agent's compensation can be based are x_1 and x_2. The principal chooses a compensation scheme to maximize her expected payoff, subject to participation and incentive constraints for the agent that reflect the agent's hidden information and hidden actions. We will compare incentive schemes according to the (expected)

[14] In Section 7.2, we show that our key results continue to hold when the degree of substitutability of efforts for the agent is high but imperfect.

[15] We assume throughout that difficulties of coordination would prevent the principal from splitting the job between two agents, with each agent responsible for only one dimension (task).

payoff generated for the principal.

Below we consider a variety of incentive schemes. Throughout the analysis, we restrict attention to compensation schedules in which, ex post, after all choices are made and random variables are realized, the agent's payment is a linear and separable function of the performance measures: $w = \alpha + \beta_1 x_1 + \beta_2 x_2$. We will say an incentive scheme (possibly involving menus) is *transparent* if, at the time the agent signs the contract or makes his choice from the menu, he is certain about what values of α, β_1, and β_2 will be employed in determining his pay. If, instead, even after making his choice from a menu, the agent is uncertain about the value of α, β_1, or β_2, we will say that the incentive scheme is *opaque*.

In the next section, we study transparent incentive schemes. Section 4 then analyzes the class of opaque scheme on which we focus, *ex ante randomization* (henceforth EAR). A contract with EAR specifies that with probability $\frac{1}{2}$, the agent will be compensated according to $w = \alpha + \beta x_1 + k\beta x_2$, and with probability $\frac{1}{2}$, he will be compensated according to $w = \alpha + \beta x_2 + k\beta x_1$, where the parameter $k \in (-1, 1)$.[16] Under EAR, the principal commits to employ a randomizing device to determine which of these two linear schedules will be used. Thus the agent, when choosing efforts, is uncertain about which performance measure will be more highly rewarded, and by varying the level of k, the principal can affect how much this uncertainty matters to the agent.[17]

3 Transparent Incentive Schemes

3.1 The No Hidden Information Benchmark

Suppose that the principal can observe the agent's cost type and offer each type a different contract. We will refer to this as the "no hidden information benchmark" (henceforth NHI). The NHI benchmark is important because, as we will see, there are environments in which optimally designed opaque contracts allow the principal, even in the presence of hidden information, to achieve a payoff arbitrarily close to that achievable in this benchmark.

[16] In contrast to transparent incentive schemes, in our model the performance of EAR cannot be improved by the inclusion of menus. Section B.1 of the online appendix shows that the principal's payoff under EAR is highest when she offers a single EAR contract that randomizes with equal probability between the two wage schedules.

[17] The restriction of the contracting space to ex post linear contracts is crucial to our analysis. If arbitrarily complex nonlinear contracts were available to the principal, it would be possible to show, by extending an argument of Grossman and Hart (1983), that given any contract with EAR, there would exist a nonlinear transparent contract that provides both types of agent with the same expected utility as a function of efforts as the contract with EAR and that (since the agent is risk-averse) entails a lower payment by the principal. However, this construction would necessitate a nonlinear contract that is complicated to describe and difficult to understand, whereas a contract which randomizes over two linear schedules is considerably simpler to describe and understand. This view is supported by the findings of Abeler and Jäger (2015), who show that the real-effort choices of subjects faced with complex incentive schemes are more dispersed and further from the payoff-maximizing level than those of subjects faced with simple ones.

In this setting, the optimal pair of contracts (one for each type of agent) can take one of two possible forms. The first form makes each type of agent willing to choose equal efforts on the two tasks, while the second form induces each type to exert effort only on his less costly task. The first form is a pair of contracts (C_1^{bal}, C_2^{bal}), where

$$C_1^{bal}:\ w_1 = \alpha + \beta x_1 + \lambda\beta x_2 \quad \text{and} \quad C_2^{bal}:\ w_2 = \alpha + \beta x_2 + \lambda\beta x_1,$$

with $\beta > 0$, and where the principal assigns the contract C_i^{bal} to the type-i agent. The incentive coefficients in C_i^{bal} are chosen to equate the ratio of the marginal benefits of efforts on the two tasks to the ratio of their marginal costs for type i. As stressed by Holmstrom and Milgrom (1991) and Milgrom and Roberts (1992, p.228), equalizing these ratios is necessary for a contract to induce strictly positive efforts on both tasks, an observation often referred to as the "equal compensation principle". Here, since these ratios are constant, independent of the chosen efforts, it follows that type i is indifferent over all non-negative effort pairs satisfying $\beta = e_i + \lambda e_j$. Among such effort pairs, the principal prefers type i to choose the perfectly balanced effort allocation, $e_i = e_j = \frac{\beta}{1+\lambda}$, since efforts on the tasks are complementary for the principal ($\delta > 1$).

Here, and throughout the paper, we assume that the agent, when indifferent over effort pairs, chooses the pair which maximizes the principal's payoff. This assumption is only ever relevant for transparent schemes; opaque schemes never leave the agent indifferent over effort pairs. Therefore, by assuming the best-case scenario for transparent schemes, we are strengthening our findings in Section 6 that opaque schemes can outperform transparent ones.

Figure 1 illustrates the outcomes from the contract pair (C_1^{bal}, C_2^{bal}) in the NHI benchmark when $\lambda > 1$. Since with transparent linear contracts, the cost of the risk imposed on the agent (stemming from the shocks to measured performance) is independent of the efforts chosen, each type of agent maximizes expected utility by maximizing the difference between the expected wage payment and the quadratic effort cost. For a type-1 agent, the isocost curves (shown in red) are linear in (e_1, e_2) space with slope equal to $-1/\lambda$. Under the contract C_1^{bal}, this agent's isobenefit curves (the curves of constant expected wage, one of which is shown in black) are also linear with the same slope $-1/\lambda$. Consequently, if for example the type-1 agent finds it optimal to incur a total effort cost corresponding to the isocost curve through points P and Q_1 in the figure, he is indifferent over all effort pairs on this isocost curve, since they all yield the same expected wage. Hence under our assumption on the agent's behavior when indifferent, he will choose the point P, at which $e_1 = e_2$. Symmetrically, for a type-2 agent, his isocost curves (blue) and the isobenefit curves corresponding to contract C_2^{bal} (black) are all linear with slope $-\lambda$, and since the value of β is the same in C_2^{bal} as in C_1^{bal}, the type-2 agent will also choose point P.

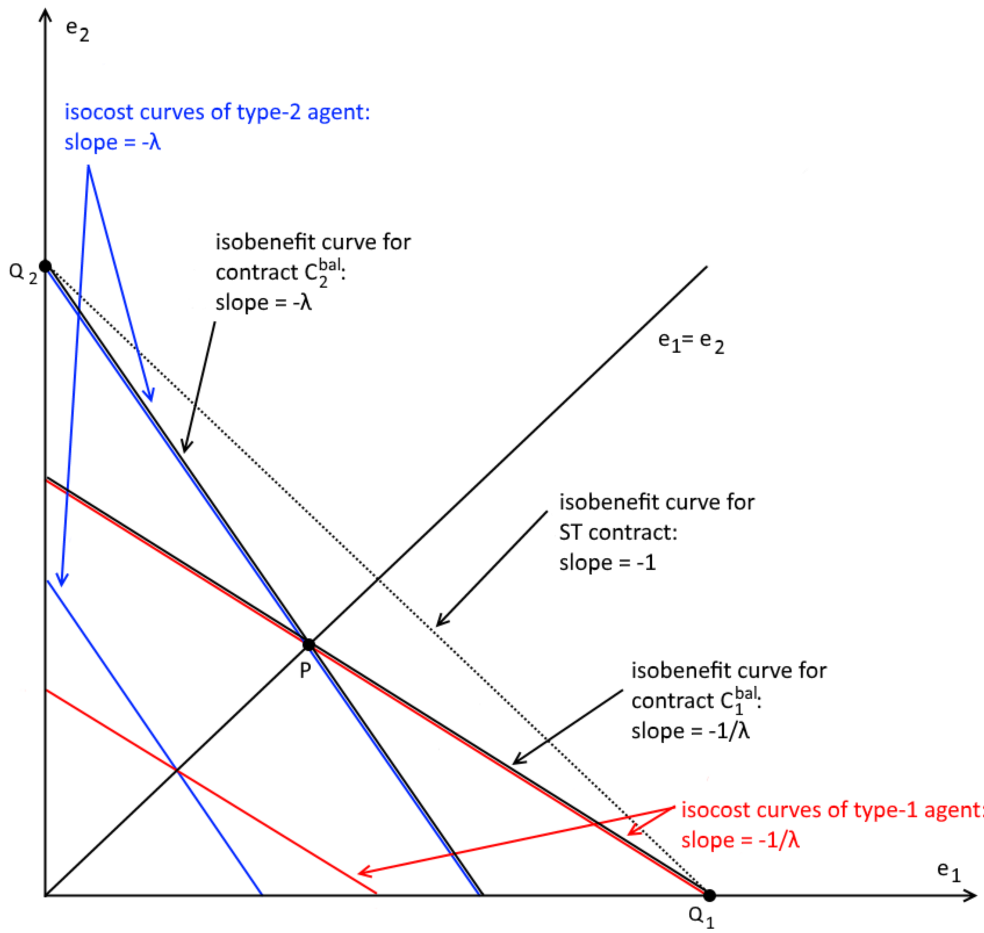

Figure 1: Isocost and isobenefit curves under transparent contracts. Contract C_1^{bal} makes the type-1 agent willing to choose point P, at which $e_1 = e_2$, and similarly the contract C_2^{bal} makes the type-2 agent also willing to choose P. The ST contract induces both types of agent to choose fully focused efforts: the type-1 agent chooses Q_1 and the type-2 agent Q_2.

Suppose that, instead of tailoring the incentive coefficients to the agent's preferences over tasks, the principal offered both types of agent a "symmetric transparent" (ST) contract

$$ST : w = \alpha + \beta x_1 + \beta x_2,$$

with β the same as in (C_1^{bal}, C_2^{bal}). Now the isobenefit curves for both types of agent would have slope -1 (one such curve is shown in Figure 1 as the dotted black line), and for both types the strictly optimal effort pair given the ST contract would be a corner solution, Q_1 for type 1 and Q_2 for type 2, corresponding to efforts fully focused on that type's less costly task. For $\lambda > 1$, the incentives provided by a symmetric transparent contract are unattractive for a principal who values effort balance: For any value of the principal's complementarity parameter δ such that $\delta > \lambda$, the principal's benefit $\frac{\delta \underline{e} + \bar{e}}{\delta + 1}$ from the fully focused effort pairs Q_1 and Q_2 is strictly below that from the perfectly balanced pair P.

In the special case where $\lambda = 1$, there is only one type of agent, and the contracts C_1^{bal} and C_2^{bal} both

reduce to the symmetric transparent contract. In this special case (and only in this case), the ST contract makes the agent indifferent between effort pairs and thus willing to choose balanced efforts $e_1 = e_2 = \frac{\beta}{2}$. Consequently, in the NHI benchmark, under our assumption on the agent's behavior when indifferent, the efforts induced by the contract pair (C_1^{bal}, C_2^{bal}), and hence the payoff received by the principal, are continuous in λ, approaching their values under the ST contract as $\lambda \to 1$.

Even though inducing perfectly balanced efforts from both types of agent, via (C_1^{bal}, C_2^{bal}), is feasible in the NHI benchmark, it is not necessarily optimal. The second type of contract pair which can be optimal in the NHI benchmark is a pair of the form

$$C_1^{foc}: \; w_1 = \alpha + \beta x_1 - \rho\beta x_2 \quad \text{and} \quad C_2^{foc}: \; w_2 = \alpha + \beta x_2 - \rho\beta x_1,$$

with $\beta > 0$, where the principal assigns C_i^{foc} to the type-i agent. Since contract C_i^{foc} has a strictly positive incentive coefficient only on x_i, this contract induces type i to exert effort only on his less costly task, task i, and for any $\lambda \geq 1$ to set $e_i = \beta$ and $e_j = 0$. Contract C_i^{foc} uses performance on task j to provide insurance for the type-i agent (without weakening his incentives on task i), by optimally exploiting the correlation between the shocks to the two performance measures.[18] Among all contract pairs that induce each type to focus only on his less costly task, pairs of the form (C_1^{foc}, C_2^{foc}) are the most attractive for the principal.[19]

In choosing, in the NHI setting, between a contract pair of the form (C_1^{bal}, C_2^{bal}) and one of the form (C_1^{foc}, C_2^{foc}), the principal faces a trade-off between the more balanced efforts induced by the former and the lower risk cost imposed by the latter. The following lemma shows that, if and only if the efforts on the two tasks are sufficiently complementary for the principal, the benefits of the balanced efforts elicited by (C_1^{bal}, C_2^{bal}) outweigh the costs of the extra risk imposed on the agent by this contract pair.

Lemma 1 *For any $\lambda \geq 1$, in the NHI benchmark, there exists a critical value of the task complementarity parameter δ in the principal's benefit function, $\delta^{NHI}(\lambda, r\sigma^2, \rho)$, increasing in each of its arguments, such that for $\delta > \delta^{NHI}$ (respectively, $\delta < \delta^{NHI}$), the principal's unique optimal contract pair has the form (C_1^{bal}, C_2^{bal}) (respectively, the form (C_1^{foc}, C_2^{foc})).*

[18] The logic here is analogous to the logic behind using relative performance evaluation to minimize an agent's exposure to risk for any given level of incentives. See, for example, Holmstrom and Milgrom (1990).

[19] Although the values of α and β could in principle be allowed to differ between C_1^{bal} and C_2^{bal} and, analogously, between C_1^{foc} and C_2^{foc}, the symmetry of the model with respect to the two types of agent makes it optimal for these values to be the same within each type of contract pair. Moreover, this symmetry also implies that it is never uniquely optimal to offer a pair of the form (C_1^{foc}, C_2^{bal}) or (C_1^{bal}, C_2^{foc}).

3.2 The General Case: Hidden Information

In the general case where $\lambda > 1$ and the agent is privately informed about his preferences across tasks, the principal can use menus of contracts as a screening device. However, the power of menus to solve the effort-substitution problem is extremely limited, as the following lemma shows.

Lemma 2 *When $\lambda > 1$, under hidden information no menu of transparent linear contracts can induce both types of agent to choose strictly positive efforts on both tasks.*

To understand Lemma 2, observe that the "equal compensation principle" has the following implication for a menu of transparent linear contracts: the only way to induce both types of agent to exert strictly positive efforts on both tasks is to induce each type to choose a contract that rewards performance on his more costly task at a rate λ times as high as it rewards performance on his less costly task. Therefore, if a menu existed which could induce both types to choose strictly positive efforts on both tasks, it would have the form

$$C_1: \ w_1 = \alpha_1 + \beta_1 x_1 + \lambda\beta_1 x_2 \quad \text{and} \quad C_2: \ w_2 = \alpha_2 + \beta_2 x_2 + \lambda\beta_2 x_1$$

and would induce the type-i agent to choose contract C_i.

We now use Figure 2, which builds on Figure 1, to explain why, no matter how $(\alpha_1, \beta_1, \alpha_2, \beta_2)$ were chosen, a menu of the form above would give at least one type of agent an incentive to select the "wrong" contract from the menu, in which case he would exert effort only on his less costly task. Suppose first that the principal sets $\beta_1 = \beta_2$ and $\alpha_1 = \alpha_2$, so that (C_1, C_2) matches the mirror-image pair (C_1^{bal}, C_2^{bal}) discussed in Section 3.1. Then if the type-1 agent were to choose C_1^{bal}, the perfectly balanced efforts of point P would maximize his expected utility. Moreover, given that C_1^{bal} and C_2^{bal} are mirror images of each other, point P would yield this agent the same expected utility under both contracts. Yet if the type-1 agent chose C_2^{bal}, under which his isobenefit curves would be more steeply negatively sloped (with slope $-\lambda$) than his isocost curves, then fully focusing his efforts by choosing point Q_1 would yield him strictly higher utility than would point P: He would incur the same overall effort cost as at P but would earn a strictly higher expected wage. Therefore, if $\beta_1 = \beta_2$ and $\alpha_1 = \alpha_2$, the type-1 agent strictly prefers to choose contract C_2^{bal} over C_1^{bal}, and symmetrically, the type-2 agent strictly prefers to choose C_1^{bal} over C_2^{bal}.[20]

The principal could, by raising β_1 and α_1 sufficiently relative to β_2 and α_2, induce the type-1 agent to choose C_1 from the menu (C_1, C_2). However, since for any β_1 and α_1, C_1 rewards task 2 more highly than task 1, the type-2 agent always derives strictly higher expected utility from C_1 than does the type-1 agent.

[20] The point Q_1 is not the type-1 agent's optimal effort choice under C_2^{bal} (he would prefer an effort pair with an even higher value of e_1, and $e_2 = 0$), but because Q_1 yields the type-1 agent higher expected utility than does P, it follows that this agent strictly prefers C_2^{bal} to C_1^{bal}.

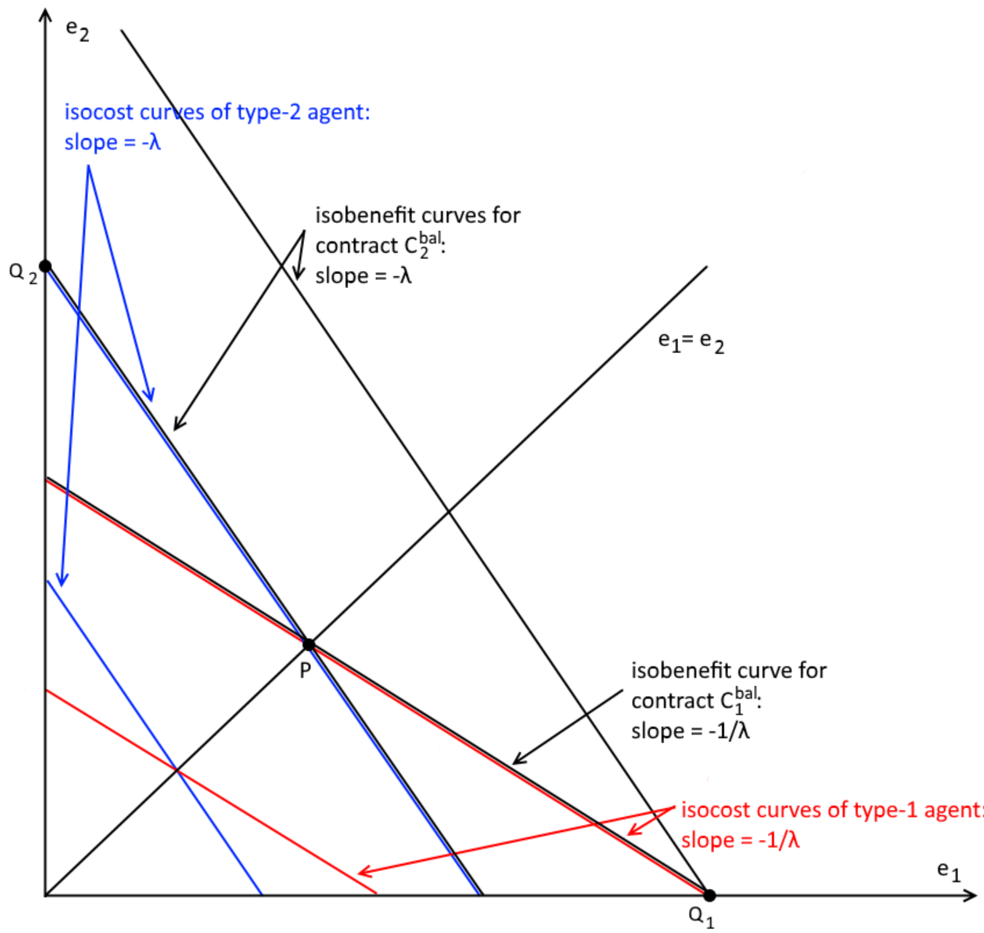

Figure 2: Graphical explanation of Lemma 2. Faced with the menu (C_1^{bal}, C_2^{bal}), if the type-1 agent were to choose C_1^{bal}, the perfectly balanced efforts of point P would maximize his expected utility. But choosing C_2^{bal} and fully focusing his efforts on task 1 (point Q_1) would yield strictly higher expected utility, since C_2^{bal} rewards task 1 more highly than does C_1^{bal}.

Thus, any adjustment in β_1 and α_1 that made the type-1 agent willing to choose C_1 would continue to induce the type-2 agent to select C_1 and would thus induce the latter agent to choose fully focused efforts.

In sum, with transparent linear contracts, which correspond in Figures 1 and 2 to linear isobenefit curves, the only way to solve the effort-substitution problem for a given type of agent is to reward more highly his more costly task. But because the bribe implicit in such a contract is even more attractive to the other type of agent, it is impossible, even with menus of transparent linear contracts, to solve simultaneously the effort-substitution and the hidden-information problems.

The above discussion makes clear that the principal might benefit from contracting instruments that generate convexity in the isobenefit curves: Since the isocost curves of the two types of agent are linear, though differently sloped, contracts that yield sufficiently convex isobenefit curves could simultaneously make interior effort choices optimal for both types of agent. We will see in Section 4 that EAR over linear contracts putting asymmetric weights on the two performance measures can generate sufficiently convex

isobenefit curves to mitigate the effort-substitution problem.

Before focusing on EAR, though, we characterize the optimal menu of transparent linear contracts and summarize some of its key properties. It follows from Lemma 2 that we can confine attention to schemes that are either an "asymmetric transparent menu" (ATM)

$$C_1^{ATM}: \ w_1 = \alpha_1 + \beta_1 x_1 - \rho\beta_1 x_2 \quad \text{and} \quad C_2^{ATM}: \ w_2 = \alpha_2 + \beta_2 x_2 + \lambda\beta_2 x_1,$$

or a "symmetric transparent menu" (STM)

$$C_1^{STM}: \ w_1 = \alpha + \beta x_1 - \rho\beta x_2 \quad \text{and} \quad C_2^{STM}: \ w_2 = \alpha + \beta x_2 - \rho\beta x_1.$$

The two mirror-image contracts in a symmetric transparent menu match the contracts (C_1^{foc}, C_2^{foc}). Since each of these contracts attaches a positive coefficient to only one performance measure, it is clear that each type of agent chooses from the menu the schedule which rewards performance on his preferred task and exerts effort only on that task. As explained in Section 3.1, the negative coefficient $-\rho\beta$ on output x_j in C_i^{STM} uses the correlation between the shocks to x_1 and x_2 to provide insurance to the type-i agent.

Under an asymmetric transparent menu of the form above, one agent (here, type 2) is induced, through appropriate choice of (α_1, α_2), to select schedule C_2^{ATM}, which leaves him indifferent over all effort pairs such that $\beta_2 = e_1 + \lambda e_2$. Given our assumption on the agent's behavior when indifferent, the type-2 agent therefore chooses the perfectly balanced effort allocation $e_1 = e_2 = \frac{\beta_2}{1+\lambda}$. The type-1 agent selects from the menu the schedule C_1^{ATM}, which induces him to choose fully focused efforts (and uses the coefficient on x_2 to provide insurance).

Relative to an STM, an ATM has the benefit of inducing one type of agent to choose balanced efforts, but it imposes more risk on that agent type and also necessitates leaving rents to the other type. Whether this benefit of an ATM outweighs these costs depends on whether δ, the strength of the principal's preference for balanced efforts, is large enough.

Proposition 1 *(i) When the agent is privately informed about his preferences, there exists a critical $\delta^{HI}(\lambda, r\sigma^2, \rho)$, increasing in each of its arguments, such that for $\delta > \delta^{HI}$, the best transparent menu for the principal is an optimally designed ATM, and for $\delta < \delta^{HI}$, her best transparent menu is an optimally designed STM.*

(ii) For all $\lambda > 1$ and for all $(r\sigma^2, \rho)$, $\delta^{HI}(\lambda, r\sigma^2, \rho) > \delta^{NHI}(\lambda, r\sigma^2, \rho)$, and as $\lambda \to 1$, $\delta^{HI} - \delta^{NHI} \to 0$.

(iii) For any $\lambda > 1$, if $\delta > \delta^{NHI}(\lambda, r\sigma^2, \rho)$, the principal is strictly worse off when hidden information is present than when it is absent.

(iv) For $\delta > \delta^{HI}(1, r\sigma^2, \rho)$, the limit as $\lambda \to 1$ of the principal's maximized payoff under hidden information

is strictly below her maximized payoff in the NHI benchmark.

The result that $\delta^{HI}(\lambda, r\sigma^2, \rho) > \delta^{NHI}(\lambda, r\sigma^2, \rho)$, for all $\lambda > 1$, says that the principal's complementarity parameter δ must be larger when hidden information is present than when it is absent for it to be optimal for her to induce balanced efforts (even from just one type of agent). The reason is the informational rent that hidden information forces the principal to leave to one agent type when offering an ATM.

Part (iii) of the proposition follows from the facts, proved in Lemmas 1 and 2, that for $\delta > \delta^{NHI}$, it is a strict optimum for the principal in the NHI benchmark to induce both types of agent to choose perfectly balanced efforts and that this outcome is infeasible under hidden information. Part (iv) shows that under hidden information, when $\delta > \delta^{HI}(1, r\sigma^2, \rho)$, the principal's maximized payoff drops discontinuously as λ is increased from 1 (where an ST contract is optimal and induces perfectly balanced efforts) to a value slightly greater than 1 (where the optimal scheme is an ATM).[21] This discontinuous drop reflects the impossibility, for even a small degree of privately-known preference across tasks, of inducing balanced efforts from both types with a transparent scheme. In contrast, in the NHI benchmark, where it is feasible to induce balanced efforts from both types for all λ, the principal's maximized payoff decreases continuously as λ is increased from 1.

4 Opaque Incentive Schemes: Ex Ante Randomization

A contract with *ex ante randomization* specifies that with probability $\frac{1}{2}$, the agent will be compensated according to $w = \alpha + \beta x_1 + k\beta x_2$, and with probability $\frac{1}{2}$, he will be compensated according to $w = \alpha + \beta x_2 + k\beta x_1$, where the key parameters are the incentive intensity $\beta > 0$ and the weighting factor $k \in (-1, 1)$.[22] Under this incentive scheme, the principal commits to employ a randomizing device to determine whether the agent's pay will be more sensitive to performance on task 1 or task 2. If the agent chooses unequal efforts on the tasks, the principal's randomization exposes the agent to extra wage risk, risk against which he can insure himself by choosing more balanced efforts. By varying k, the principal can affect how much risk the randomization per se imposes on the agent and can thereby affect the strength of the agent's incentives to balance his efforts. If k were equal to 1, the randomized scheme would collapse to the symmetric transparent (ST) contract defined in Section 3.1, which, whenever $\lambda > 1$, induces both types of agent to exert effort only on their preferred task. The smaller k is, the greater the risk imposed on the agent by the principal's randomization is, and therefore the stronger the agent's incentives are to self-insure by choosing more balanced efforts.

[21] Note that we are continuing to use our assumption that the agent, when indifferent over effort pairs, chooses the pair that maximizes the principal's payoff.

[22] The lump-sum payment α has no effect on the agent's incentives, and will always be set by the principal to make the participation constraint binding for both types of agent.

Since the two equally likely compensation schedules under EAR are mirror images with respect to the two tasks and since the cost functions of the two types of agent are also mirror images, the optimal effort choices of the two types of agent will also be mirror images. Hence we can describe both agents' optimal efforts by the same pair $(\overline{e}^{EAR}, \underline{e}^{EAR})$, where $\overline{e}^{EAR}$ denotes the effort on the agent's less costly task and $\underline{e}^{EAR}$ the effort on the agent's more costly task. Furthermore, since the principal's benefit function depends only on the minimum and maximum of the efforts on the two tasks, and not which task attracted larger effort, the principal's expected payoff under EAR can also be written as a function of $(\overline{e}^{EAR}, \underline{e}^{EAR})$.

Proposition 2 *(i) Under EAR, $k < \frac{1}{\lambda}$ is a necessary condition for each agent's optimal efforts on both tasks to be strictly positive. When EAR induces interior solutions for efforts,*
(ii) the efforts choices of each type of agent satisfy

$$\begin{aligned} \overline{e}^{EAR} + \lambda \underline{e}^{EAR} &= \frac{\beta(1+k)}{\lambda+1} && (1) \\ \exp\left[r\beta(1-k)(\overline{e}^{EAR} - \underline{e}^{EAR})\right] &= \frac{\lambda - k}{1 - k\lambda}; && (2) \end{aligned}$$

(iii) the gap in efforts, $\overline{e}^{EAR} - \underline{e}^{EAR}$, is increasing in λ, approaching 0 as $\lambda \to 1$; decreasing in $r\beta$, approaching 0 as $r\beta \to \infty$; and increasing in k, approaching 0 as $k \to -1^+$;
(iv) the principal's expected payoff under EAR, for given $\beta > 0$ and $k \in (-1, \frac{1}{\lambda})$, is

$$\begin{aligned} \Pi^{EAR}(\beta, k) = \left(\frac{\delta \underline{e}^{EAR} + \overline{e}^{EAR}}{\delta + 1}\right) - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} \\ - \frac{1}{2} r\sigma^2\beta^2(1 + 2\rho k + k^2) - \frac{1}{2r}\ln\left[\frac{(\lambda+1)^2(1-k)^2}{4(1-k\lambda)(\lambda-k)}\right]. \quad (3) \end{aligned}$$

To understand part (i), note that if $k \geq \frac{1}{\lambda}$, then for both types of agent, *whichever of the two compensation schedules is randomly selected*, the ratio of the marginal benefit of effort on the less costly task to that on the more costly task is at least as large as the corresponding ratio of the marginal costs of effort and strictly larger for one of the schedules. It follows from the "equal compensation principle", therefore, that both types of agent would optimally exert effort only on their less costly task.

Since for both types of agent, the total cost of the efforts incurred is a quadratic function of $\overline{e} + \lambda \underline{e}$, we will henceforth refer to the quantity $\overline{e} + \lambda \underline{e}$ as the *aggregate effort* exerted. Equation (1) shows the aggregate effort induced by EAR (at an interior solution for efforts), while equation (2) yields the gap between efforts on the two tasks.

To understand equation (1), note first that the sum of the marginal monetary returns to effort on the two tasks is certain to be $\beta(1+k)$ because, regardless of the outcome of the randomization, one task will be re-

warded at rate β and the other at rate $k\beta$. If optimal efforts are interior, then adding the first-order conditions for $\overline{e}$ and $\underline{e}$ must yield $\beta(1+k) = \partial c/\partial\overline{e} + \partial c/\partial\underline{e}$ for both types of agent. Since $\partial c/\partial\overline{e} + \partial c/\partial\underline{e} = (1+\lambda)(\overline{e}+\lambda\underline{e})$, this gives us (1). Note that the aggregate effort induced by EAR is independent of the agent's risk aversion.

To understand equation (2), observe that we can express each type of agent's expected utility under EAR as

$$\begin{aligned}
-\frac{1}{2}\exp&\left\{-r\left[\alpha+\frac{\beta(1+k)}{2}(\overline{e}+\underline{e})+\frac{\beta(1-k)}{2}(\overline{e}-\underline{e})-c(\overline{e},\underline{e})-\mathcal{RP}\right]\right\}\\
&-\frac{1}{2}\exp\left\{-r\left[\alpha+\frac{\beta(1+k)}{2}(\overline{e}+\underline{e})-\frac{\beta(1-k)}{2}(\overline{e}-\underline{e})-c(\overline{e},\underline{e})-\mathcal{RP}\right]\right\}\\
&=-\exp\left\{-r[b(\overline{e},\underline{e})-c(\overline{e},\underline{e})-\mathcal{RP}]\right\},
\end{aligned} \tag{4}$$

where $c(\overline{e},\underline{e}) = \frac{1}{2}(\overline{e}+\lambda\underline{e})^2$ is the cost of efforts, $\mathcal{RP} = \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2)$ is the risk premium stemming from the shocks to measured performance, and $b(\overline{e},\underline{e})$ is the certainty equivalent of the benefit, under EAR, from effort levels $(\overline{e},\underline{e})$ given by

$$b(\overline{e},\underline{e}) = \alpha + \frac{\beta(1+k)}{2}(\overline{e}+\underline{e}) - \frac{1}{r}\ln\left\{\frac{\exp[-\frac{1}{2}r\beta(1-k)(\overline{e}-\underline{e})]+\exp[\frac{1}{2}r\beta(1-k)(\overline{e}-\underline{e})]}{2}\right\}. \tag{5}$$

Observe that $b(\overline{e},\underline{e})$ is weakly less than $\alpha + \frac{\beta(1+k)}{2}(\overline{e}+\underline{e})$ which is the expected wage payment under EAR. The negative term in (5) represents the risk premium stemming from the principal's randomization over payment schedules. Under EAR, each type of agent chooses $(\overline{e},\underline{e})$ to maximize $b(\overline{e},\underline{e}) - c(\overline{e},\underline{e})$ (since $\mathcal{RP}$ is independent of efforts): an interior solution therefore equates $\frac{\partial b}{\partial\overline{e}}/\frac{\partial b}{\partial\underline{e}}$ to $\frac{\partial c}{\partial\overline{e}}/\frac{\partial c}{\partial\underline{e}}$. This yields

$$\frac{\partial b}{\partial\overline{e}}\Big/\frac{\partial b}{\partial\underline{e}} = \frac{[1+k\exp\{r\beta(1-k)(\overline{e}-\underline{e})\}]}{[k+\exp\{r\beta(1-k)(\overline{e}-\underline{e})\}]} = \frac{1}{\lambda} = \frac{\partial c}{\partial\overline{e}}\Big/\frac{\partial c}{\partial\underline{e}}, \tag{6}$$

and the central equality can be rearranged to yield equation (2).

The expression for $\frac{\partial b}{\partial\overline{e}}/\frac{\partial b}{\partial\underline{e}}$ in (6) shows that under EAR, and in contrast to transparent linear contracts, the isobenefit curves of $b(\overline{e},\underline{e})$ for a risk-averse agent are convex to the origin: As $\overline{e}-\underline{e}$, the gap in efforts on the two tasks, rises from 0, $\frac{\partial b}{\partial\overline{e}}/\frac{\partial b}{\partial\underline{e}}$ falls from 1. This convexity of the isobenefit curves reflects that when the efforts are more unequal, the wage risk from the randomization over payment schedules is greater, and thus the incentive for self-insurance (raising $\underline{e}$ relative to $\overline{e}$) is stronger.

Figure 3 illustrates the effort incentives created by EAR. With e_1 on the horizontal and e_2 on the vertical axis, isocost curves for the type-1 agent are shown in red and isocost curves for the type-2 agent in blue. Defining $b(e_1,e_2)$ by substituting e_1 for $\overline{e}$ and e_2 for $\underline{e}$ in equation (5), we plot the isobenefit curves of $b(e_1,e_2)$ in black. Since compensation under EAR is ex ante symmetric with respect to the two performance measures, these isobenefit curves are symmetric about the line $e_1 = e_2$. For each type of agent, equation (1) determines which isocost curve (corresponding to a level of aggregate effort) the chosen effort pair lies on, while the solution to equation (6) (equivalently, equation (2)) is represented, for each type, by a point of tangency

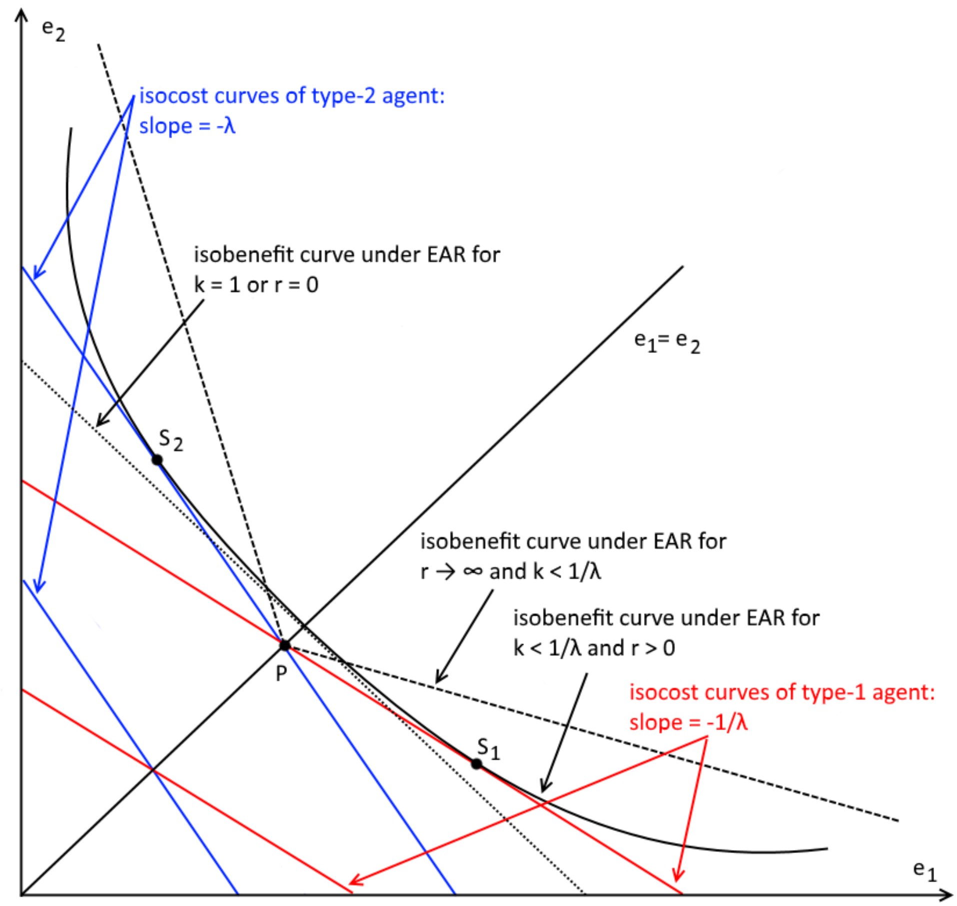

Figure 3: Under EAR with $k < 1$, the isobenefit curves for a risk-averse agent are convex to the origin and symmetric about the line $e_1 = e_2$. For the solid black isobenefit curve, each type of agent's optimal effort pair (S_1 for type 1 and S_2 for type 2) is a point of tangency between that curve and an isocost curve for that type of agent.

between that isocost curve and an isobenefit curve of $b(e_1, e_2)$. In Figure 3, the optimal effort pair for the type-1 agent is S_1, with $e_1 > e_2 > 0$, and that for the type-2 agent is, by symmetry, S_2, with $e_2 > e_1 > 0$.

Equation (2) and Figure 3 both show how each type of agent's optimal degree of self-insurance against the wage risk imposed by EAR varies with the parameters of the incentive scheme and with his preferences. First, the smaller the parameter k is, the more different the two possible compensation schedules are and so the more costly the wage risk imposed by the randomization is. This is reflected in greater convexity of the isobenefit curves ($\frac{\partial b}{\partial \overline{e}} / \frac{\partial b}{\partial \underline{e}}$ falling from 1 faster as $\overline{e} - \underline{e}$ rises from 0).[23] As a consequence, reducing k strengthens the agent's incentive to self-insure by choosing more balanced efforts, and the optimal effort gap $\overline{e} - \underline{e}$ falls. As $k \to -1^+$, the self-insurance motive approaches its strongest level, and the optimal effort gap approaches 0.

Second, greater risk aversion of the agent (larger r) and a larger value of the incentive intensity β also

[23]Equation (6) shows that for $k \in (-1, 0)$, the risk imposed by the randomization makes $\frac{\partial b}{\partial \overline{e}}$ negative for $\overline{e} - \underline{e}$ sufficiently large (while $\frac{\partial b}{\partial \underline{e}}$ is always positive), so for $k \in (-1, 0)$, the isobenefit curves of $b(e_1, e_2)$ become positively sloped far enough away from the 45-degree line.

make the wage risk imposed by the randomization more costly and so, just like a smaller value of k, make the isobenefit curves more convex. As a result, the larger $r\beta$ is, the stronger the self-insurance incentives of EAR are, and thus the smaller the optimal effort gap $\overline{e} - \underline{e}$ is. In Figure 3, as $r \to \infty$, the slope of the isobenefit curves approaches $-k$ for points below the 45-degree line and $-1/k$ for points above the 45-degree line. One such curve is shown by the dashed black line. Hence for $k < 1/\lambda$, as $r \to \infty$, the optimal effort gap $\overline{e} - \underline{e}$ for each type of agent approaches 0: this corresponds to full self-insurance. Moreover, it follows from equation (1) that for each type of agent, his optimal effort pair remains on the same isocost curve as r increases, and hence as $r \to \infty$ with β held fixed, each type's optimal choice approaches point P.

If k were 1 or the agent were risk-neutral, then equation (5) and Figure 3 show that the isobenefit curves of $b(\overline{e}, \underline{e})$ would be linear with slope -1, coinciding with the isobenefit curves for an ST contract as defined in Section 3.1. One such curve is shown by the dotted black line. In either of these cases, therefore, EAR would, like an ST contract, induce fully focused efforts for any $\lambda > 1$. These cases highlight that for EAR to generate incentives for effort-balancing, it is essential that the agent face genuine wage risk and have a strict desire to insure himself against it.

Finally, the smaller the cost difference between tasks (i.e., smaller λ and thus slope of the linear isocost curves closer to -1), the less costly it is for the agent to self-insure by choosing a smaller optimal effort gap $\overline{e} - \underline{e}$). As $\lambda \to 1$, full self-insurance becomes optimal, so $\overline{e} - \underline{e}$ approaches 0.

Introducing a small amount of hidden information about the agent's preferences (raising λ from 1) has a strikingly different effect under EAR than under an ST contract. At $\lambda = 1$, an ST makes the agent indifferent over all effort pairs such that $\beta = \overline{e} + \underline{e}$ and hence willing to choose $\overline{e} = \underline{e}$. But for any $\lambda > 1$, efforts are fully focused on a single task. As a result, the principal's payoff from an ST contract drops discontinuously as λ is raised from 1. Furthermore, Proposition 1 shows that this discontinuity persists even under an optimal menu of transparent contracts. In contrast, under EAR, for any value of $k \in (-1, 1)$, both the agent's efforts and the principal's payoff are continuous in λ at $\lambda = 1$ if the agent is risk-averse. Thus EAR is more robust to the introduction of private information on the part of the agent than is the best transparent menu.[24] EAR is also more robust to uncertainty about the magnitude of λ than is a transparent menu: If the principal tries to design a transparent menu to induce one type of agent to choose balanced efforts but is even slightly wrong about the magnitude of λ, her payoff will be discontinuously lower than if she were right. The performance of EAR does not display this extreme sensitivity.[25]

[24]Even outside the exponential-normal framework we have been using, EAR induces more balanced efforts than an ST contract and is more robust to the introduction of hidden information on the agent's part as shown in Section B.2 of the online appendix.

[25]Bond and Gomes (2009) also study a multi-task principal-agent setting in which a small change in the agent's preferences can result in a large change in the behavior induced by a contract and a consequent large drop in the principal's payoff, a situation they term "contract fragility".

The effort-balancing incentives generated by EAR do, however, come at a cost in terms of the risk imposed on the risk-averse agent. As shown by equations (4) and (5), EAR imposes two distinct types of risk costs. The first is the risk stemming from the shocks to measured performance (which is the risk that would be imposed by a transparent contract of the form $w = \alpha + \beta x_1 + k\beta x_2$, or equivalently, $w = \alpha + \beta x_2 + k\beta x_1$). The second is the risk imposed by the principal's randomization over payment schedules. The first is represented by the term $\mathcal{RP}$ in (4), and the second by the negative term in (5), reflecting the amount by which $B(\overline{e}, \underline{e})$ falls short of the expected wage under EAR. Correspondingly, in the principal's payoff expression (3) in Proposition 2 the penultimate term is the risk premium stemming from the shocks and the final term is the risk premium stemming from the randomization.

We saw above how the principal, under EAR, can affect the strength of the agent's incentives for balanced efforts by varying k, the parameter representing the degree of asymmetry in the weights on the performance measures in the two possible compensation schemes. However, k also affects the level of aggregate effort induced, as equation (1) shows. To isolate the effect of k on the principal's overall payoff under EAR, holding fixed the level of aggregate effort, we define $B \equiv \beta(1+k)$ and use equations (1) and (2)) to re-express the principal's payoff (3) as a function of B and k:

$$\Pi^{EAR}(B,k) = \left(\frac{\delta\underline{e}^{EAR} + \overline{e}^{EAR}}{\delta+1}\right) - \frac{B^2}{2\left(\lambda+1\right)^2} - \frac{1}{2}r\sigma^2 B^2 \frac{1+2\rho k+k^2}{(1+k)^2} - \frac{1}{2r}\ln\left[\frac{\left(\lambda+1\right)^2\left(1-k\right)^2}{4(1-k\lambda)(\lambda-k)}\right], \quad (7)$$

where

$$\left(\frac{\delta\underline{e}^{EAR} + \overline{e}^{EAR}}{\delta+1}\right) = \frac{B}{\left(\lambda+1\right)^2} - \frac{\delta-\lambda}{\delta+1}\frac{\ln\left(\frac{\lambda-k}{1-k\lambda}\right)}{\left(\lambda+1\right)rB\left(\frac{1-k}{1+k}\right)}. \quad (8)$$

Holding B fixed and varying k identifies the effect of k on the principal's payoff from inducing any given level of aggregate effort. Equations (7) and (8) show that increasing k has three effects. First, a larger k raises the effort gap $\overline{e} - \underline{e}$ and, with B and hence aggregate effort $\overline{e} + \lambda\underline{e}$ held fixed, this larger gap lowers the principal's benefit $\frac{\delta\underline{e}+\overline{e}}{\delta+1}$ whenever $\delta > \lambda$. Second, a larger k, because it induces the agent to choose less balanced efforts, raises the cost of compensating the agent for the risk imposed by the randomization per se. This second effect of k also reduces the principal's payoff and is reflected in the final term in (7). Finally, a larger k reduces the cost (per unit of aggregate effort induced) of the risk imposed on the agent from the shocks to measured performance. This improved diversification raises $\Pi^{EAR}(B,k)$, as reflected in the second-to-last term in (7).

In general, the optimal design of a contract with EAR involves a trade-off among these three different effects. Weighting the different performance measures more equally in the two possible compensation schedules is costly in terms of effort balance and thereby in terms of the risk imposed by the randomization, but is helpful in allowing better diversification of the measurement errors. The next proposition describes

how the optimal value of k varies with several parameters of the contracting environment, holding fixed the aggregate effort to be induced, and also how the optimal k changes as the desired aggregate effort changes.

Proposition 3 *For any given level of aggregate effort to be induced, the optimal level of k under EAR is smaller (the optimal weights on the performance measures should be more unequal)*
(i) the larger is δ, given $\delta > \lambda$ (i.e., the stronger the principal's preference for balanced efforts);
(ii) the smaller is r, holding $r\sigma^2$ fixed (i.e., the less risk-averse the agent, holding fixed the importance of risk aversion under transparent contracts);
(iii) the smaller is $\sigma^2(1-\rho)$ (i.e., the lower the importance of diversification of the risk from the shocks to measured performance);
(iv) the smaller is B (i.e., the smaller the level of aggregate effort to be induced).

In Section 7.3, where we study EAR in a setting with an arbitrary number n of tasks, we show that changes in the number of randomly chosen tasks to reward have the same qualitative effects on incentives and risk as do changes in the weighting parameter k in the two-task model. The comparative statics results for the optimal number of tasks to reward are the same as those above for the optimal k.

5 When Are Transparent Incentive Schemes Preferred?

Section 4 showed that the key advantage of EAR is the effort-balancing incentives it generates for the privately informed risk-averse agent. This section demonstrates that each of the three key model ingredients we have highlighted—the agent's hidden information about his preferences, the agent's risk aversion, and the principal's desire for the agent's efforts to be balanced across tasks—is *necessary* for EAR to dominate the best transparent scheme.

Proposition 4 *For any given (β, k), with $k \in (-1, 1)$, EAR yields a strictly lower payoff for the principal than a suitably designed transparent contract, if any of the following conditions hold:*
(ia) $\lambda > 1$ and the principal knows which task the agent finds less costly;
(ib) $\lambda = 1$, so the agent finds both tasks equally costly, and $\rho < 1$;
(ii) the agent is not sufficiently risk-averse for EAR to induce positive effort on both tasks;
(iii) $\delta \leq \lambda$, so the principal's desire for balanced efforts is outweighed by the agent's preference across tasks.

Underlying each part of this proposition is the important result that, given any contract with EAR, there exists a transparent contract that induces the same level of *aggregate* effort on the two tasks and that

imposes lower overall risk costs. Under any of the conditions identified in the proposition, the benefits from the effort-balancing incentives generated by EAR are insufficient to outweigh the greater risk costs it imposes.

Suppose first that $\lambda > 1$ but the principal knows which task the agent finds less costly, so we are in the NHI benchmark. There are three respects in which EAR performs worse than what the principal can achieve in the NHI benchmark. First, in the NHI benchmark the principal can induce perfectly balanced efforts by offering the agent for whom task i is less costly a contract of the form C_i^{bal}: $w = \alpha + \beta x_i + \lambda\beta x_j$. When $\lambda > 1$, EAR, by contrast, does not induce perfectly balanced efforts, and this per se (i.e., holding aggregate effort fixed) lowers the principal's benefit. Second, given the unequal efforts chosen under EAR, the randomization over compensation schedules imposes risk costs on the agent for which the principal has to compensate him. Third, EAR must set k strictly less than $1/\lambda$ if it is to induce interior effort choices (as Proposition 2 showed), whereas with C_i^{bal}, k is effectively set equal to $1/\lambda$. That is, the ratio of the smaller to the larger incentive coefficient in C_i^{bal} is $1/\lambda$. As shown in the discussion following equation (7), when k is larger, the diversification of the risk from the shocks to measured performance is better. This risk is thus better diversified under C_i^{bal} than under EAR. Overall, therefore, EAR is dominated by a transparent contract in the NHI benchmark both with respect to effort balance and with respect to the two sources of risk costs borne by the agent.

If $\lambda = 1$, EAR induces perfectly balanced efforts, just as (under our assumption on the agent's behavior when indifferent) a symmetric transparent (ST) contract does. Nevertheless, under EAR, $k < 1$, whereas under an ST contract, $k = 1$. Hence, an ST contract that induces the same efforts as EAR achieves better diversification of the risk from the shocks and thus generates a higher overall payoff for the principal, strictly higher as long as the correlation of the shocks, ρ, is strictly less than 1.

If the agent is not sufficiently risk averse for EAR to induce positive efforts on both tasks, then an ST contract can induce exactly the same focused effort pair as EAR. Moreover, such an ST contract imposes lower overall risk costs, both by achieving better diversification of the risk from the shocks and by insulating the agent from the wage risk stemming from the randomization.

Finally, if $\delta < \lambda$, the socially efficient effort allocation involves each type of agent focusing exclusively on his preferred task. Given an EAR contract inducing strictly positive efforts on both tasks, an ST contract that induces the same aggregate effort will yield the principal a strictly higher payoff, both because of the greater social surplus from the focused efforts under the ST contract and, as in the paragraph above, because of the lower overall risk costs.

Proposition 4 is important because it highlights that each of our three key model ingredients is *necessary* for EAR to outperform the best transparent menu. The next section, which is the heart of the paper, identifies when these three key ingredients together are *sufficient* for EAR to do so.

6 When Are Opaque Incentive Schemes Preferred?

We now analytically and later numerically identify environments in which opaque schemes, when designed optimally, strictly dominate the best transparent menu. In each of the three environments for which we prove the superiority of EAR analytically, this superiority follows in the limit from our demonstration that EAR allows the principal to achieve a payoff arbitrarily close to what she could achieve if she knew the agent's preferences across tasks, as in the NHI benchmark. Hence, in these limiting environments, EAR eliminates the efficiency losses from the agent's hidden information.

Given that Proposition 4 (part (ia) and the discussion following it) identifies three distinct respects in which EAR performs worse than what the principal could achieve in the NHI benchmark, it may seem surprising that we can find environments in which optimally designed EAR yields her a payoff arbitrarily close to her NHI benchmark. The explanation is as follows. In each of our three limiting environments, EAR, with the weighting parameter k adjusted optimally, induces essentially *perfectly* balanced efforts. With perfectly balanced efforts, the first and second drawbacks of EAR relative to the NHI benchmark disappear. Furthermore, in all three environments, with k set optimally, the risk from the shocks to measured performance is as well diversified under EAR as it is in the NHI benchmark.

Even though we focus for the most part on limiting environments, our analytical results are strong in two respects. First, as explained above, not only do they show that optimally designed EAR outperforms the best transparent menu under hidden information, but also that it approximates the principal's payoff in the NHI benchmark. Second, they show that *for any level of aggregate effort to be induced,* EAR dominates the best transparent menu. Even without optimizing with respect to the overall intensity of incentives, we can be sure that in these environments (and those close to them), EAR dominates. This means, for example, that even if the benefit component of the principal's payoff were scaled up or down relative to the wage cost, the results of Propositions 5, 6, and 7 would continue to hold.

6.1 Very Weak Preferences across Tasks for the Agent

Consider first a setting in which the agent has private information about his preferences, but the magnitude of his preference across tasks is very weak. Formally, we study the case in which λ is strictly greater than but arbitrarily close to 1, which we term the limiting case as $\lambda \to 1^+$.

We saw in Section 4 that under EAR the agent's effort choices and the principal's payoff are continuous in λ at $\lambda = 1$. This robustness of EAR to the introduction of hidden information underlies the superiority of this scheme in the limiting case as $\lambda \to 1^+$, as we now show.

Proposition 2 shows that as $\lambda \to 1$, so the two tasks become equally costly, $\overline{e} - \underline{e} \to 0$ for any $k \in (-1, 1)$ under EAR. Equations (7) and (8) show how varying k affects the principal's payoff from EAR, $\Pi^{EAR}(B, k)$, holding fixed at $\frac{B}{1+\lambda}$ the level of aggregate effort induced. Whereas in general, as discussed in Section 4, increasing k has conflicting effects on $\Pi^{EAR}(B, k)$, in the limit as $\lambda \to 1$, the situation is dramatically simpler:

$$\lim_{\lambda \to 1} \Pi^{EAR}(B, k) = \frac{B}{4} - \frac{B^2}{8} - \frac{1}{2} r\sigma^2 B^2 \left(\frac{1 + 2\rho k + k^2}{(1+k)^2} \right). \tag{9}$$

Because, as $\lambda \to 1$, efforts under EAR become perfectly balanced, the risk cost imposed by the randomization tends to zero. Hence an increase in k has only one effect on $\Pi^{EAR}(B, k)$, holding B fixed: it improves the diversification of the shocks to measured performance, as reflected in the final term of (9). Thus, as $\lambda \to 1$, $\Pi^{EAR}(B, k)$ is increasing in k (strictly so for $\rho < 1$), as long as k induces interior solutions, which it does as long as $k < \frac{1}{\lambda}$. Therefore, as $\lambda \to 1$, $\Pi^{EAR}(B, k)$ is maximized, for any B, by setting k arbitrarily close to, but less than, 1 ($k \to 1^-$). With k set in this way, the principal's payoff approaches

$$\lim_{k \to 1} \lim_{\lambda \to 1} \Pi^{EAR}(B, k) = \frac{B}{4} - \frac{B^2}{8} - \frac{1}{4} r\sigma^2 B^2 (1 + \rho). \tag{10}$$

The right-hand side of (10) equals the payoff the principal would achieve, if λ were exactly equal to 1, from a symmetric transparent (ST) contract with $\beta = \frac{B}{2}$, since such a contract would induce effort $\frac{B}{4}$ on each task and generate the same diversification of the shocks as EAR does when $k \to 1^-$.[26] Thus, for any B, as $\lambda \to 1^+$, the principal's payoff under optimally weighted EAR is arbitrarily close to that from an ST contract when the agent has no preference between tasks.

For the NHI benchmark, Section 3.1 shows that the efforts and payoff from the contract pair (C_1^{bal}, C_2^{bal}) are continuous at $\lambda = 1$, where they match the efforts and payoff from the ST contract. Lemma 1 shows that as $\lambda \to 1$, a pair of the form (C_1^{bal}, C_2^{bal}) is strictly optimal for the principal as long as $\delta > \lim_{\lambda \to 1} \delta^{NHI}(\lambda, r\sigma^2, \rho)$. On the other hand, Proposition 1 shows that under hidden information, even as $\lambda \to 1^+$, the principal's maximized payoff from a transparent menu is bounded away from that in the NHI benchmark, because even for λ arbitrarily close to 1, it is impossible to induce positive efforts on both tasks from both types of agent.

The arguments in the preceding paragraphs together imply:

Proposition 5 *Consider the limiting case as $\lambda \to 1^+$. Under EAR, for any given level of aggregate effort, $\overline{e} + \lambda \underline{e}$, to be induced:*

(i) the gap in efforts, $\overline{e} - \underline{e}$, approaches 0 for any $k \in (-1, 1)$;

(ii) the optimal value of $k \to 1^-$;

[26] See equation (20) in the proof of Lemma 1 in the appendix, and set $\lambda = 1$.

(iii) with k adjusted optimally, the principal's payoff under EAR approaches her payoff in the NHI benchmark from (C_1^{bal}, C_2^{bal}). This limiting payoff equals the principal's payoff from the symmetric transparent (ST) contract at $\lambda = 1$.

Therefore, for $\delta > \lim_{\lambda \to 1} \delta^{NHI}(\lambda, r\sigma^2, \rho)$, EAR with k and β adjusted optimally strictly dominates the best transparent menu under hidden information.

The preceding analysis has another important implication. Even if $\lambda = 1$, the symmetric transparent contract leaves the agent indifferent to how total effort is split between the tasks, while under EAR, for any $k < 1$, the optimal allocation of efforts is unique. Thus, when $\lambda = 1$, with the weighting parameter k set arbitrarily close to (but less than) 1, EAR not only yields the principal a payoff arbitrarily close to the best-case payoff from the ST contract, but EAR also ensures that the agent has a *strict* preference for choosing perfectly balanced efforts.

6.2 Very Large Risk Aversion and Very Small Variance of the Shocks

Consider now the effect of increasing the agent's coefficient of absolute risk aversion r, holding fixed the value of the product $r\sigma^2$. This change has no impact on the principal's payoff from any transparent scheme, since with transparent schemes, the agent's efforts are independent of r, and the risk premium from the shocks to measured performance depends on r only via the product $r\sigma^2$. This change does, however, increase the principal's payoff under EAR, as long as EAR induces interior solutions for efforts. The reason is that, as shown by equations (1) and (2), an increase in the agent's risk aversion r, while it has no effect on the aggregate effort induced by EAR, strengthens the agent's incentive to self-insure against the wage risk from the randomization. The resulting reduction in $\overline{e} - \underline{e}$ both raises the principal's benefit, as shown in equation (8), and reduces the cost of compensating the agent for the risk from the randomization, as shown by final term in (7). As with transparent schemes, the cost of the risk imposed on the agent from the shocks to measured performance depends on r only via $r\sigma^2$, so remains unchanged.

We have thus shown:

Lemma 3 *Holding $r\sigma^2$ fixed, increasing r increases the principal's payoff from EAR, as long as EAR induces interior solutions for efforts, but leaves the principal's payoff from any transparent scheme unchanged.*

It follows from Lemma 3 that the more risk-averse the agent, holding $r\sigma^2$ fixed, the more likely it is that optimally designed EAR will dominate the best transparent menu. We now consider the limiting case where r gets very large and σ^2 gets very small, with $r\sigma^2$ fixed at $R < \infty$. Proposition 2 shows that, in

this environment, for any $k \in (-1, \frac{1}{\lambda})$, $(\overline{e} - \underline{e}) \to 0$ under EAR. As the agent becomes infinitely risk-averse, it becomes optimal for him to choose perfectly balanced efforts, so providing full self-insurance against the wage risk generated by the randomization.

Under EAR, in the limit as $r \to \infty$ and $\sigma^2 = \frac{R}{r} \to 0$, both $\overline{e}$ and $\underline{e}$ approach $\frac{B}{(\lambda+1)^2}$ (as long as $k < \frac{1}{\lambda}$). As a consequence, $\Pi^{EAR}(B,k)$, as given by equations (7) and (8), simplifies to

$$\lim_{r\to\infty,\sigma^2=R/r\to 0} \Pi^{EAR}(B,k) = \frac{B}{(\lambda+1)^2} - \frac{B^2}{2(\lambda+1)^2} - \frac{1}{2}RB^2\frac{1+2\rho k+k^2}{(1+k)^2}. \tag{11}$$

Equation (11) shows that, when r gets very large and $\sigma^2 = \frac{R}{r}$ very small, the only effect on $\Pi^{EAR}(B,k)$ of increasing k, over the range $k \in (-1, \frac{1}{\lambda})$ where the induced gap in efforts $(\overline{e} - \underline{e})$ is approximately 0, is to improve the diversification of the shocks to measured performance. Hence, just as for the case where $\lambda \to 1^+$, it is optimal to set k arbitrarily close to, but less than, $\frac{1}{\lambda}$ (i.e., $k \to (\frac{1}{\lambda})^-$). Doing so generates for the principal a payoff approaching

$$\lim_{k\to(1/\lambda)^-} \lim_{r\to\infty,\sigma^2=R/r\to 0} \Pi^{EAR}(B,k) = \frac{B}{(\lambda+1)^2} - \frac{B^2}{2(\lambda+1)^2} - \frac{1}{2}RB^2\frac{\lambda^2+2\rho\lambda+1}{(\lambda+1)^2}. \tag{12}$$

The right-hand side of (12) is exactly the payoff the principal would obtain, in the NHI benchmark, from using (C_1^{bal}, C_2^{bal}) with $\beta = \frac{B}{1+\lambda}$, since this pair of contracts would induce from each type of agent effort $\frac{B}{(\lambda+1)^2}$ on each task and would impose a risk premium (from the shocks to measured performance) given by the final term.[27]

Thus as $r \to \infty$ and $\sigma^2 = \frac{R}{r} \to 0$, optimally weighted EAR allows the principal, for any B, to get arbitrarily close to her payoff in the NHI benchmark. Since, by Proposition 1, the best transparent menu under hidden information leaves the principal strictly worse off than in the NHI benchmark whenever $\delta > \delta^{NHI}(\lambda, R, \rho)$, we have proved:

Proposition 6 *Consider the limiting case where $r \to \infty$ and $\sigma^2 = \frac{R}{r} \to 0$. Under EAR, for any given level of aggregate effort, $\overline{e} + \lambda\underline{e}$, to be induced:*

(i) the gap in efforts, $\overline{e} - \underline{e}$, approaches 0 for any λ and for any $k < \frac{1}{\lambda}$;

(ii) the optimal value of $k \to \left(\frac{1}{\lambda}\right)^-$;

(iii) with k adjusted optimally, the principal's payoff under EAR approaches her payoff in the NHI benchmark from (C_1^{bal}, C_2^{bal}).

Therefore, for $\delta > \delta^{NHI}(\lambda, R, \rho)$, EAR with k and β adjusted optimally strictly dominates the best transparent menu under hidden information.

[27] See equation (20) in the appendix, and set $r\sigma^2 = R$.

6.3 Very High Correlation between the Shocks

Under EAR, as the correlation, ρ, between the shocks to measured performance on the two tasks approaches 1, diversification of the risk from these shocks becomes impossible. This is reflected in equation (7) where, as $\rho \to 1$, the risk premium due to the shocks, $\frac{1}{2} r\sigma^2 B^2 \frac{1+2\rho k+k^2}{(1+k)^2}$, approaches $\frac{1}{2} r\sigma^2 B^2$, which is independent of k. This means that although in general, as shown by equations (7) and (8), the weighting factor k under EAR has conflicting effects on the principal's payoff $\Pi^{EAR}(B,k)$, in the limit as $\rho \to 1$, $\Pi^{EAR}(B,k)$ is *decreasing* in k, for any fixed B. Lowering k strengthens the agent's self-insurance motive, leading to a reduction in $\bar{e} - \underline{e}$ which both raises the principal's benefit and reduces the risk cost of the exogenous randomization.

Hence, as $\rho \to 1$, it is optimal under EAR, for any level of aggregate effort to be induced, to set k arbitrarily close to, but larger than, -1 (that is, $k \to -1^+$), thus inducing a gap in efforts arbitrarily close to, but larger than, 0 (as shown by (2)). With k set in this way, the principal achieves under EAR a payoff arbitrarily close to[28]

$$\lim_{k \to -1^+} \lim_{\rho \to 1} \Pi^{EAR}(B,k) = \frac{B}{(\lambda+1)^2} - \frac{B^2}{2(\lambda+1)^2} - \frac{1}{2} r\sigma^2 B^2. \tag{13}$$

This limiting payoff matches what the principal would obtain, in the NHI benchmark with $\rho = 1$, from using (C_1^{bal}, C_2^{bal}) to induce perfectly balanced efforts and setting $\beta = \frac{B}{1+\lambda}$.[29] Thus, in this limiting environment as well, optimally weighted EAR yields the principal as high a payoff as in the absence of hidden information for any level of aggregate effort to be induced. Combining these results with Proposition 1 yields:

Proposition 7 *Consider the limiting case of perfect correlation of the shocks: $\rho \to 1$. Under EAR, for any given level of aggregate effort, $\bar{e} + \lambda \underline{e}$, to be induced:*

(i) the optimal value of $k \to -1^+$, and the resulting gap in efforts, $\bar{e} - \underline{e}$, approaches 0 for any λ;

(ii) with k adjusted optimally, the principal's payoff under EAR approaches her payoff in the NHI benchmark from (C_1^{bal}, C_2^{bal}).

Therefore, for $\delta > \delta^{NHI}(\lambda, r\sigma^2, 1)$, EAR with k and β adjusted optimally strictly dominates the best transparent menu under hidden information.

Analogous arguments and conclusions hold in the limiting environment where the variance σ^2 of the shocks to measured performance approaches 0, holding risk aversion r fixed. In this environment, optimally weighted EAR allows the principal to come arbitrarily close to the outcome in which perfectly balanced efforts are induced with the imposition of no risk cost at all on the agent. As a result, EAR allows the

[28] As k is lowered, the coefficient β must be raised to keep aggregate effort, which is proportional to $B \equiv \beta(1+k)$, fixed. The value of k must remain slightly larger than -1 to ensure that aggregate effort is strictly positive.

[29] See equation (20) in the appendix, and set $\rho = 1$.

principal to achieve a payoff arbitrarily close to what he would achieve in the absence of both hidden action and hidden information. For $\lambda > 1$, the best transparent menu under hidden information cannot achieve balanced efforts, even as $\sigma^2 \to 0$, so the best transparent menu is dominated in this environment by optimally designed EAR, as long as $\delta > \delta^{NHI}(\lambda, 0, \rho) = \lambda$.

6.4 Numerical Results

A general analytic characterization of the optimal values of the weighting factor k and the incentive intensity β under EAR is prohibitively complex. As shown in equations (7) and (8), k has complicated nonlinear effects on the principal's payoff, and even if k were fixed at some specified value (e.g., 0), the optimal β would be the solution to a cubic equation, because increasing β not only increases incentives for aggregate effort (equation (1)) but also strengthens the agent's self-insurance motive for balancing efforts (equation (2)).

This section uses numerical methods to optimize both the weighting factor k and the incentive intensity β under EAR. We then compare the principal's maximized payoff under EAR to that under the best transparent menu, which is characterized in Proposition 1 and is either an asymmetric transparent menu (ATM) or a symmetric transparent menu (STM). The numerical analysis demonstrates the robustness of the effects highlighted by our analyses of limiting environments. Specifically, it confirms that the benefits of EAR in inducing balanced efforts are more likely to outweigh the extra risk costs it imposes when (i) the agent's privately known preference between tasks is weak (λ is small), so even a small amount of uncertainty about the weights in the compensation schedule provides a strong impetus for effort balance, (ii) the agent is very risk-averse (r is large), so opaque schemes generate a powerful self-insurance motive for effort balance, (iii) efforts on the tasks are highly complementary for the principal (δ is high), or (iv) the errors in measuring performance on the tasks have large correlation or small variance ($\sigma^2(1-\rho)$ is small).

Figure 4a plots the regions in which EAR (black), STM (gray) or ATM (white) are optimal for the principal for different combinations of the agent's preference parameter λ and the principal's task complementarity parameter δ, holding the other parameters fixed at $r = 4$, $\sigma^2 = 0.02$, and $\rho = 0$. EAR is optimal for λ not too large and δ sufficiently large. As the agent's preference between tasks λ becomes stronger, under EAR the risk costs imposed on the agent eventually become too costly for the principal to use low values of the weighting factor k to induce balanced efforts. This is the case even when effort balance δ is very important to the principal. Between the two types of transparent menus, ATM, which induces balanced efforts from one type of agent, is optimal only when δ is sufficiently large relative to λ. Otherwise, the principal's optimal scheme is STM which, though inducing fully focused efforts from both types of agent, saves on risk costs and also avoids leaving rent to one type.

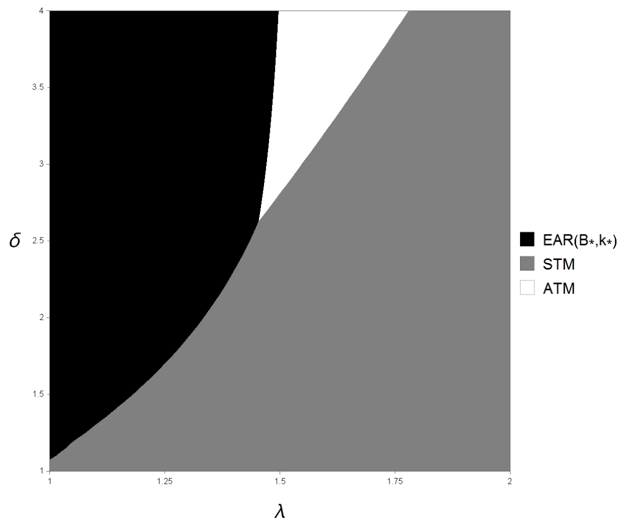

(a) $r = 4$ and $\sigma^2 = 0.02$

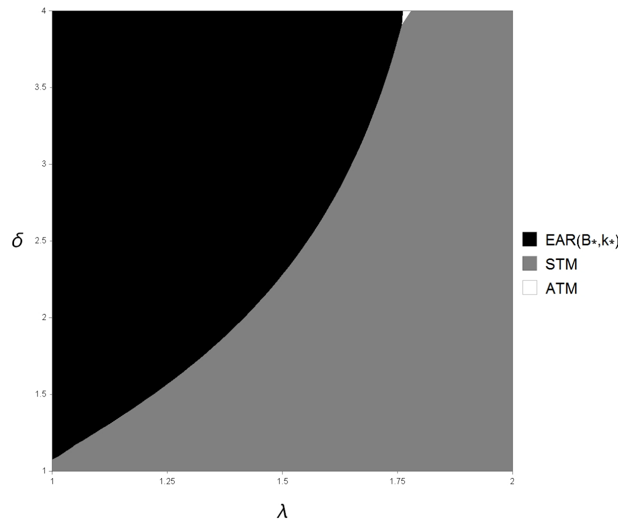

(b) $r = 8$ and $\sigma^2 = 0.01$

Figure 4: Optimal schemes for different combinations of λ and δ, with $\rho = 0$.

Figure 4b shows how the dominance regions for the three incentive schemes change as the agent's risk aversion r increases and the variance σ^2 of the shocks to the performance measures falls, holding the product $r\sigma^2$ constant (and keeping $\rho = 0$). The contrast between Figures 4a and 4b illustrates the implication of Lemma 3 in Section 6.2. Increasing the risk aversion of the agent, holding $r\sigma^2$ constant, expands the region in parameter space in which EAR outperforms the best transparent menu. As explained earlier, such a change strengthens the agent's self-insurance motive under EAR, resulting in more balanced efforts and thus a higher payoff for the principal, but has no effect on the principal's payoff under any transparent scheme. The contrast between the two figures also shows that the expansion of the dominance region for EAR occurs primarily at the expense of the dominance region for ATM.

Figure 5a plots the dominance regions for the three incentive schemes for different combinations of the agent's preference parameter λ and risk aversion r, adjusting σ^2 as r changes so as to keep $r\sigma^2$ constant at 0.08 and fixing $\delta = 2$ and $\rho = 0$. There is a critical value of λ (independent of r, since $r\sigma^2$ is held fixed) below which the best transparent menu is ATM and above which it is STM: this is the case since the principal's payoff from STM, given the focused efforts it induces, is independent of λ, while that from ATM declines with λ. Most importantly, for any λ, as implied by Lemma 3, there is a critical value of r above which EAR is superior to the best transparent menu. Figure 5a shows that this critical value of r is increasing in λ (because as λ increases, inducing balanced efforts under EAR necessitates imposing too much risk on the agent), and the critical r is very steeply increasing for (the large) values of λ for which STM dominates ATM. This is because the payoff from STM is independent of λ.

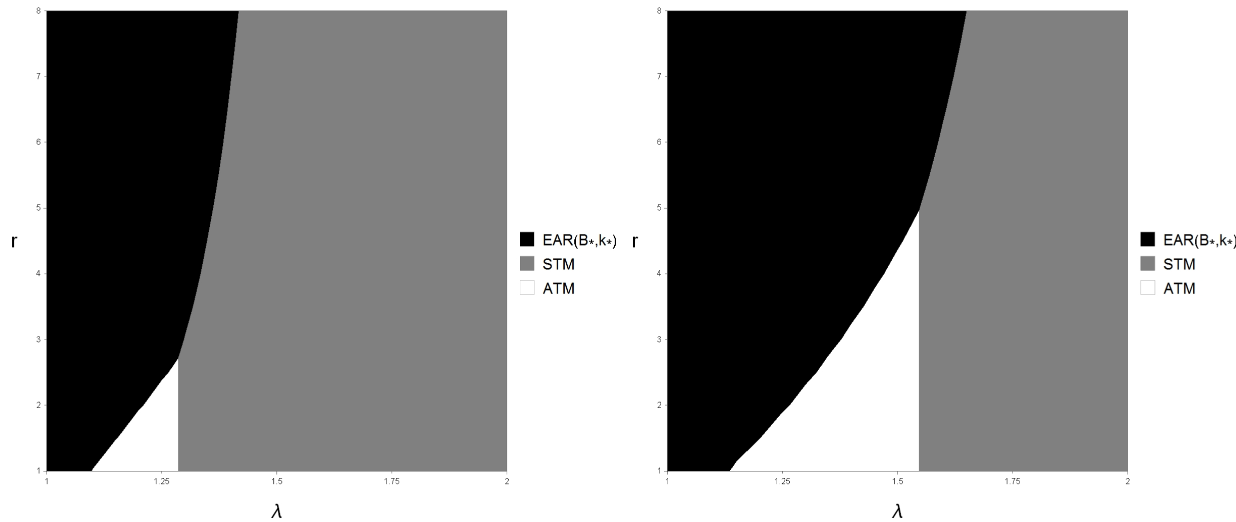

(a) $\delta = 2$ **(b)** $\delta = 3$

Figure 5: Optimal schemes for different combinations of λ and r, with $r\sigma^2 = 0.08$ and $\rho = 0$.

Figure 5b shows how the dominance regions for the three incentive schemes change as the principal's task complementarity parameter δ increases from 2 to 3, keeping everything else the same as in Figure 5a. As δ increases, the region where EAR outperforms both transparent menus expands, as does the region where ATM is the optimal scheme. The expansions in the dominance regions for EAR and ATM illustrate that these two schemes, unlike STM, achieve some degree of effort balance, which becomes more valuable for the principal as δ increases.

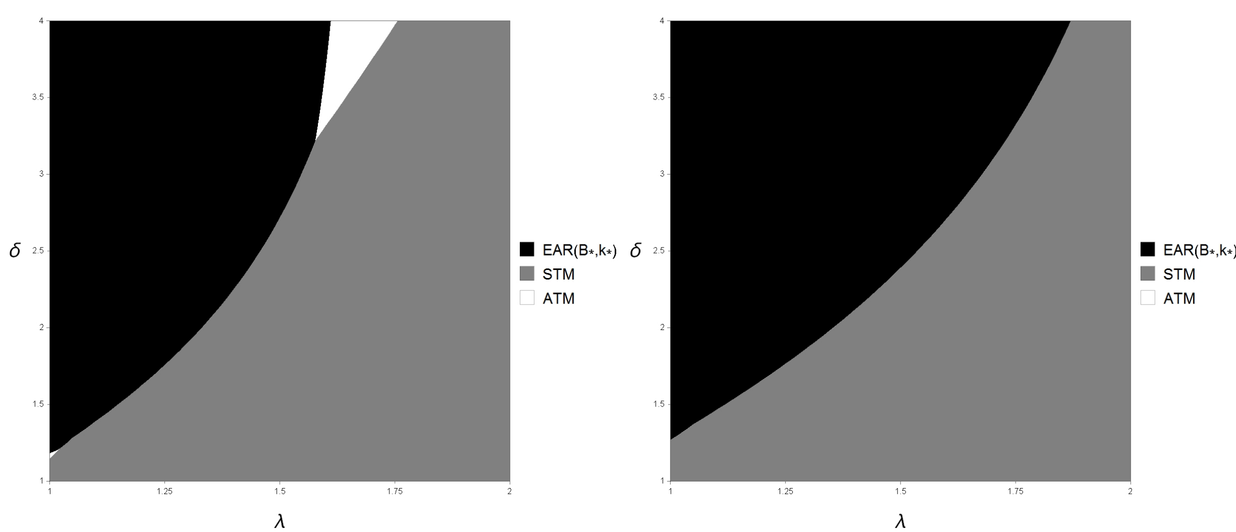

(a) $\rho = 0.5$ **(b)** $\rho = 0.9$

Figure 6: Optimal schemes for different combinations of λ and δ, with $r = 4$ and $\sigma^2 = 0.02$.

Figures 6a and 6b show that raising the correlation ρ of the shocks to the performance measures increases

the principal's payoff from EAR relative to her payoff from the transparent menus. These figures, like Figure 4a, plot the dominance regions for different combinations of λ and δ, and they use the same parameter values as Figure 4a (where we set $\rho = 0$) except that now $\rho = 0.5$ and $\rho = 0.9$. Contrasting these three figures shows that the dominance region for EAR expands with the increase in ρ. As shown in Section 6.3, under EAR, when ρ gets large, the scope for using the weighting factor k to diversify the risk due to the shocks diminishes, so the dominant effects of k on the principal's payoff in equations (7) and (8) are the ones operating through its effect on effort balance. As ρ increases, therefore, it becomes less and less costly in terms of risk to set a low value of k in order to induce a high degree of effort balance. As with the effect of increasing the agent's risk aversion (from Figure 4a to Figure 4b), the expansion here of the dominance region for EAR as ρ increases occurs primarily at the expense of the dominance region for ATM.

These numerical results demonstrate the robustness of the effects highlighted by our analytical results for the limiting environments. They confirm that EAR is more likely to dominate the best transparent menu when the agent's privately known preference between tasks is weaker, when the agent is more risk-averse, when the tasks are more complementary for the principal, and when the correlation in the performance measure shocks is higher.

7 Extensions and Robustness

7.1 Alternative Assumptions on the Principal's Commitment Powers

We have analyzed the trade-offs involved in the choice between transparent and opaque incentive schemes under the assumption that under EAR the principal can, before the agent makes his effort choices, commit to randomizing uniformly between the two compensation schedules.[30] It is natural to wonder whether opaque incentive schemes corresponding to alternative assumptions about the principal's commitment powers would change our conclusions.

Assume, instead, that the principal chooses the randomizing probability at the same time as the agent chooses efforts. We term this incentive scheme *interim randomization*. We can prove that under interim randomization, the unique (Bayes-Nash) equilibrium is exactly the same as the outcome described in Proposition 2. Therefore, all of our results on the benefits and costs of opacity continue to hold.

To see that the outcome described in Proposition 2 is *an* equilibrium under interim randomization, note that given that the two types of agent are equally likely and given that their effort profiles are mirror images,

[30] Given the power to commit to a randomizing probability, it is optimal for the principal to commit to randomize uniformly. Doing so results in the most balanced profile of effort choices, assessed ex ante, and also avoids leaving any rent to either type of agent.

the principal anticipates equal expected output on the two tasks, so is willing to randomize uniformly over the two mirror-image compensation schedules. Given that the principal randomizes uniformly, the optimal behavior for each type of agent is clearly as described in the proposition. To see that this outcome is the *unique* equilibrium, observe that if the two types of agent conjectured that the principal would choose the schedule rewarding task 1 more highly than task 2 with a probability greater than (less than) 1/2, then the ex ante expected profile of efforts would be skewed towards task 1, so the principal would strictly prefer to choose the schedule rewarding task 2 more (less) highly than task 1. Thus, the attractive properties of EAR are not crucially dependent on the principal's having the power to commit to the randomizing probability.

We also obtain qualitatively similar results for another class of opaque incentive schemes. Under a contract with *ex post discretion* (EPD), the principal, *after* observing the performance measures x_1 and x_2, chooses whether to pay the agent according to $w = \alpha + \beta x_1 + k\beta x_2$ or $w = \alpha + \beta x_2 + k\beta x_1$, where again $k \in (-1, 1)$. EPD provides the agent with the same self-insurance motive but also generates an additional incentive for effort balance: the principal's strategic ex post choice of which linear schedule to use means that the more the agent focuses his effort on his preferred task, the less likely that task is to be the more highly compensated one, so the lower the relative marginal return to that task. In an earlier version of this paper (Ederer, Holden, and Meyer, 2014), we showed that the opaque incentives resulting from EPD generate at least as great a payoff for the principal as EAR because (i) EPD induces a strictly smaller gap in efforts $\overline{e} - \underline{e}$ than EAR, while the two schemes induce the same aggregate effort $\overline{e} + \lambda\underline{e}$ and hence the same total cost of effort, and (ii) EPD imposes lower risk costs on the agent than EAR. As a result, the beneficial incentive effects of EAR are robust even if the agent suspects that the principal might deviate to EPD.

7.2 Imperfect Substitutability of Efforts for the Agent

So far we have focused on the case where efforts are perfect substitutes in the agent's cost function. Although this assumption does not qualitatively affect the performance of EAR, it simplifies the analysis of transparent schemes. We explain here that even when we allow some substitutability of efforts, transparent schemes continue to suffer dramatically from the problem of gaming by an agent with hidden information about his preferences. Consequently, it remains true that (i) if tasks are sufficiently complementary for the principal, EAR is superior to transparent menus in settings where EAR generates very strong incentives for balanced efforts, and (ii) in such settings, EAR eliminates the efficiency losses from the agent's hidden information.

Let the two equally likely types of agent have cost functions of the form

$$c(\overline{e}, \underline{e}) = \frac{1}{2}\left(\overline{e}^2 + 2s\lambda\overline{e}\underline{e} + \lambda^2\underline{e}^2\right), \tag{14}$$

where the parameter $s \in [0,1]$ measures the degree of substitutability of efforts. Perfect substitutability corresponds to $s = 1$ and no substitutability to $s = 0$.

With the cost function given in (14), the ratio of the marginal cost of effort on the agent's costlier task to that on his cheaper task is $\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}} = \frac{s\lambda\overline{e}+\lambda^2\underline{e}}{\overline{e}+s\lambda\underline{e}}$. When efforts are imperfect substitutes for the agent ($s < 1$), the isocost curves of $c(\overline{e}, \underline{e})$ are concave to the origin: Starting from perfectly balanced efforts, as the agent shifts his effort allocation towards his preferred task (increasing $\overline{e}$ and decreasing $\underline{e}$), $\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}}$ falls. However, the minimum value of this ratio, attained when $\underline{e} = 0$, is $s\lambda$. It follows that as long as $s\lambda \geq 1$ (representing a situation of high, but imperfect, substitutability), a symmetric transparent contract (for which the isobenefit curves have slope -1) still induces fully focused efforts from both types of agent, just as with perfect substitutability.

It also follows that, under hidden information, the only way with transparent contracts to induce interior efforts from both types of agent is to induce each type to choose, from a menu, a contract that rewards his costlier task at least $s\lambda$ times as highly as his cheaper task. But with $s\lambda \geq 1$, the bribe implicit in such a contract is even more attractive to the other type of agent. As a consequence, Lemma 2 continues to hold as long as $s\lambda \geq 1$, implying that it is impossible, even with menus of transparent linear contracts, to solve simultaneously the effort-substitution and the hidden-information problems.

In the NHI benchmark, on the other hand, the principal can offer each type of agent a contract of the form $w = \alpha + \beta\overline{x} + v\beta\underline{x}$ with $v \geq 1$, where $\overline{x}$ (respectively, $\underline{x}$) denotes measured performance on the preferred (respectively, other) task. The weighting factor v is a choice variable for the principal, and under the simplifying assumption that the tasks are perfect complements for her ($\delta \to \infty$), it is always optimal for her to induce each type to choose equal efforts on the two tasks, which is achieved by $v^{NHI} = \frac{\lambda(\lambda+s)}{1+s\lambda}$. This finding, combined with the generalization of Lemma 2 noted above, implies that the principal's maximized payoff from transparent menus under hidden information is bounded away from that in the NHI benchmark.

Importantly, the incentives provided by EAR are not qualitatively affected by whether efforts are imperfect or perfect substitutes for the agent. EAR continues to give the risk-averse agent an incentive to partially self-insure by choosing relatively balanced efforts on the two tasks. Interior optimal efforts under EAR satisfy

$$\frac{\partial c}{\partial \overline{e}} + \frac{\partial c}{\partial \underline{e}} = \beta(1+k) \tag{15}$$

and

$$\exp\left[r\beta(1-k)(\overline{e}-\underline{e})\right] = \frac{\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}} - k}{1 - k\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}}}. \tag{16}$$

Equation (16) generalizes (2), replacing the constant λ with the function $\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}}$ of $(\overline{e}, \underline{e})$.

Consider now the three environments studied in detail in Section 6. As $\lambda \to 1^+$ or as $r \to \infty$, $\sigma^2 \to 0$, it

follows from (16) that EAR induces perfectly balanced efforts for any $k \in (-1, \frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}})$.[31] Therefore, in these limiting cases, the only effect of increasing k is to improve the diversification of the risk from the shocks. Hence it is optimal in both environments to set k as large as possible subject to keeping efforts perfectly balanced, i.e., to take $k \to (\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}})^{-}$. Since with perfectly balanced efforts, $\frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}} = \frac{1+s\lambda}{\lambda(\lambda+s)} = 1/v^{NHI}$, it follows that as $\lambda \to 1^{+}$ or as $r \to \infty$, $\sigma^2 \to 0$, the optimal k approaches $1/v^{NHI}$. Therefore, just as in the original model, in these two limiting environments, optimally weighted EAR generates a payoff for the principal arbitrarily close to what she achieves in the NHI benchmark. In the setting where $\rho \to 1$, the weight k has no effect on diversification, so it is optimal under EAR to set k to induce perfectly balanced efforts; in this setting, too, optimally weighted EAR generates a payoff arbitrarily close to that in the NHI benchmark.

As long as $s\lambda \geq 1$, we saw above that under hidden information, the principal's maximized payoff from transparent menus is bounded away from that in the NHI benchmark. It follows, therefore, that in the environments studied in Section 6, optimally designed EAR is superior to the best transparent menu. Hence, allowing the agent's efforts on the tasks to be less than perfect substitutes in his cost function does not alter our main results.

7.3 Opaque Incentives and the Choice of How Many Tasks to Reward

We have assumed so far that the job performed by the agent has only two distinct dimensions (tasks) and that noisy measures of performance on both tasks are used in randomized incentive schemes. When, however, performance on a job has many distinct dimensions, the costs of monitoring the different dimensions may become significant. The principal can economize on monitoring costs, while still providing incentives for balanced efforts, by randomizing over compensation schedules each of which rewards only a subset of the dimensions of performance. We now study some of the trade-offs involved in the design of randomized incentive schemes in environments with many tasks. We find that reducing the number of tasks randomly selected to be rewarded, holding fixed the aggregate effort induced, has qualitatively the same effects, on the agent's incentives and on the principal's payoff, as reducing the weighting coefficient k in EAR in the two-task model. Consequently, the same factors that made it optimal in the two-task model to create significant uncertainty for the agent about the relative weights to be used make it optimal with many tasks to create significant uncertainty about which ones will actually be rewarded.

Let the job performed by the agent consist of $n > 2$ tasks, for each of which measured performance $x_j = e_j + \epsilon_j$, where $(\epsilon_1, \ldots, \epsilon_n)$ have a symmetric multivariate normal distribution with mean 0, variance σ^2, and pairwise correlation $\rho \geq 0$. Suppose there are n equally likely types of agent, with the agent of type i

[31] If $k > \frac{\partial c/\partial \underline{e}}{\partial c/\partial \overline{e}}$, (16) shows that EAR cannot induce interior solutions for efforts.

having cost function $c_i(e_1, \ldots, e_n) = \frac{1}{2}(\lambda e_i + \sum_{j \neq i} e_j)^2$. Thus, each type of agent has a particular *dislike* for exactly one of the n tasks, and λ measures the intensity of this dislike. Let the principal's payoff be given by

$$\Pi = \frac{\delta}{\delta + n - 1} \min\{e_1, \ldots, e_n\} + \frac{1}{\delta + n - 1} \left(\sum_{j=1}^{n} e_j - \min\{e_1, \ldots, e_n\} \right) - w,$$

where δ parameterizes the strength of the principal's desire for a balanced effort profile. As in the two-task model, the socially efficient effort profile is perfectly balanced whenever $\delta > \lambda$.

Consider the following family of incentive schemes with EAR, parameterized by κ, the number of tasks rewarded: Each subset of κ out of n tasks is chosen with equal probability, each task in the chosen subset is rewarded at rate β, and the lump-sum payment is always equal to α. We will not explicitly model the direct costs of generating the performance measures. Instead we will focus on the trade-off between the effects on incentives and risk of varying the number of tasks κ included in each of the possible compensation schedules. Details of the derivations are in Appendix B.3.

Denote by $\underline{e}$ each type of agent's effort on his disliked task and by $\overline{e}$ his effort on each of the other tasks. If, for a given κ and β, the agent's optimal efforts are interior, then aggregate effort $(\lambda \underline{e} + (n-1)\overline{e})$ and the gap in efforts $\overline{e} - \underline{e}$ satisfy, respectively,

$$\lambda \underline{e} + (n-1)\overline{e} = \frac{\kappa \beta}{\lambda + n - 1} \qquad \text{and} \qquad \overline{e} - \underline{e} = \frac{1}{r\beta} \ln \left[\frac{\lambda(n - \kappa)}{(n-1) - (\kappa - 1)\lambda} \right]. \tag{17}$$

Reducing κ, the number of tasks rewarded, makes the risk imposed by the randomization more costly, so strengthens the agent's incentive to self-insure. As a result, the agent's optimal effort profile is more balanced ($\overline{e} - \underline{e}$ is smaller), the smaller is the number of tasks rewarded.

Since aggregate effort is proportional to $\kappa\beta$, define $\tilde{\beta} \equiv \kappa\beta$. Using (17), we can write the principal's payoff as a function of $\tilde{\beta}$ and κ, when α is set to ensure zero rent for each type of agent:

$$\begin{aligned} \Pi(\tilde{\beta}, \kappa) = \frac{\delta \underline{e} + (n-1)\overline{e}}{\delta + n - 1} &- \frac{\tilde{\beta}^2}{2\,(\lambda + n - 1)^2} \\ &- \frac{1}{2} r \sigma^2 \tilde{\beta}^2 \frac{(1 + \rho(\kappa - 1))}{\kappa} - \frac{1}{nr} \ln \left[\frac{(n-\kappa)^{n-\kappa}(n - 1 + \lambda)^n}{n^n \lambda^\kappa \left((n-1) - (\kappa - 1)\lambda\right)^{n-\kappa}} \right], \end{aligned} \tag{18}$$

where

$$\frac{\delta \underline{e} + (n-1)\overline{e}}{\delta + n - 1} = \frac{\tilde{\beta}}{(\lambda + n - 1)^2} - \frac{(\delta - \lambda)(n-1)\kappa}{(\delta + n - 1)(\lambda + n - 1) r \tilde{\beta}} \ln \left[\frac{\lambda(n - \kappa)}{(n-1) - (\kappa - 1)\lambda} \right]. \tag{19}$$

Holding $\tilde{\beta}$ fixed and varying κ isolates the effect of changing the number of tasks rewarded, holding fixed the level of aggregate effort. Comparison of equations (18)-(19) with equations (7)-(8) reveals that changes in κ have qualitatively the same three effects on the principal's payoff in this n-task model as do variations in the weighting coefficient k in EAR in the original two-task model. Specifically, an increase in

κ, by inducing a larger gap $\overline{e} - \underline{e}$, has two negative effects: First, it lowers the principal's benefit $\underline{e} + \frac{n-1}{\delta}\overline{e}$ when aggregate effort is held fixed, as long as $\delta > \lambda$. This corresponds to the fact that (19) is decreasing in κ. Second, it raises the cost of compensating the agent for the risk imposed by the exogenous randomization (this corresponds to the fact that the term in square brackets in (18) is increasing in κ). At the same time, raising κ also improves the diversification of the risk from the shocks to measured performance. This is reflected in the fact that $\frac{1+\rho(\kappa-1)}{\kappa}$ in (18) is decreasing in κ.

Given the qualitative similarity between the role of κ in the n-task model and that of k in the two-task model, it is relatively straightforward to derive the following comparative statics results for the optimal number of tasks to reward, given any desired level of aggregate effort. Analogously with Proposition 3, the optimal number of tasks to reward is smaller, (i) the stronger the principal's preference for balanced efforts, (ii) the less risk-averse the agent (holding $r\sigma^2$ fixed), (iii) the lower the importance of diversification of the risk from the shocks to measured performance, and the smaller the level of aggregate effort to be induced.

8 Conclusion

Gaming of incentive schemes is a serious concern to incentive designers in a wide range of settings. We analyzed a principal-agent model in which the agent's superior information about the environment leads to the gaming of (menus of) transparent incentive schemes. In contrast, opaque incentive schemes not only mitigate the agent's gaming but can generate a higher overall payoff for the principal despite imposing additional risk on the agent. In general, the principal faces a trade-off between the benefits of the more efficient effort allocations induced by opaque schemes and the costs of the greater risk they impose.

We showed that opaque incentives are superior when (i) the agent's privately known preference between tasks is weak, so even a small degree of opacity in the compensation schedule has a large effect on inducing balance in the agent's efforts, or (ii) the agent's risk aversion is significant, so opaque schemes provide the agent with a powerful self-insurance motive for balancing efforts, or (iii) the principal has a strong desire for effort balance, or (iv) the errors in measuring performance on the tasks have large correlation or small variance.

We emphasize that, because of the agent's hidden information, opaque schemes can dominate transparent ones in our model even when pay can be based upon measured performance on both tasks. When costs of measurement constrain an incentive designer to base pay on only one performance measure, the attractiveness of opaque incentives over which task to measure and reward is clearly significantly enhanced relative to the best transparent contract rewarding only one task.

Our analysis suggests that even beyond the specific multi-task setting on which we have focused, opacity

of incentive schemes can be a valuable tool for incentive designers when there are restrictions on the complexity of reward schemes or when resources for monitoring agents are limited. By making agents more uncertain about the consequences of their actions for their rewards, opaque schemes can help principals to mitigate the costs of gaming by agents who are exploiting their superior information about the environment. Future research should explore the benefits of opaque incentive schemes in deterring gaming in other settings, identifying under what conditions these incentive benefits can outweigh the risk costs of opacity.

References

[1] **Abeler, Johannes and Simon Jäger**, "Complex Tax Incentives," *American Economic Journal: Economic Policy*, Aug. 2015, *7* (3), 1-28.

[2] **Baker, George P.**, "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, Jun. 1992, *100* (3), 598-614.

[3] **Baker, George, Robert Gibbons, and Kevin J. Murphy**, "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics*, Nov. 1994, *109* (4), 1125-1156.

[4] **Bentham, Jeremy**, *Constitutional Code*, Vol. 1, London: R. Heward, 1830.

[5] **Bevan, Gwyn and Christopher Hood**, "Targets, Inspections, and Transparency," *British Medical Journal*, March 2004, *328*, 598.

[6] **Burgess, Simon, Propper, Carol, Ratto, Marisa and Emma Tominey**, "Incentives in the Public Sector: Evidence from a Government Agency," *CEPR Discussion Paper 9071*, Jul. 2012.

[7] **Bond, Philip and Armando Gomes**, "Multitask principal-agent problems: Optimal contracts, fragility, and effort misallocation," *Journal of Economic Theory*, Jan. 2009, 144 (1), 175-211.

[8] **Carrell, Scott E. and James E. West**, "Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors," *Journal of Political Economy*, Jun. 2010, *118* (3), 409-432.

[9] **Carroll, Gabriel**, "Robustness and Linear Contracts," *American Economic Review*, Feb. 2015, *105* (2), 536-563.

[10] **Dranove, David, Kessler, Daniel, McClellan, Mark and Mark Satterthwaite**, "Is More Information Better? The Effects of 'Report Cards' on Health Care Providers," *Journal of Political Economy*, Jun. 2003, *111* (3), 555-588.

[11] **Ederer, Florian, Holden, Richard, and Margaret Meyer**, "Gaming and Strategic Opacity in Incentive Provision", *Cowles Foundation Working Paper*, 2014.

[12] **Forbes, Silke, Lederman, Mara and Trevor Tombe**, "Quality Disclosure Programs and Internal Organizational Practices: Evidence from Airline Flight Delays," *American Economic Journal: Microeconomics*, May 2015, *7* (2), 1-26.

[13] **Frankel, Alex and Navin Kartik**, "Muddled Information," *Working Paper*, 2014.

[14] **Gjesdal, Froystein**, "Information and Incentives: The Agency Information Problem," *Review of Economic Studies*, Jul. 1982, *49* (3), 373-390.

[15] **Gravelle, Hugh, Sutton, Matt and Ada Ma**, "Doctor Behaviour under a Pay for Performance Contract: Treating, Cheating and Case Finding?" *Economic Journal*, Feb. 2010, *120* (1), F129-F156.

[16] **Grossman, Sanford J. and Oliver D. Hart**, "An Analysis of the Principal-Agent Problem," *Econometrica*, Jan. 1983, *51* (1), 7-45.

[17] **Holmstrom, Bengt and Paul Milgrom**, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, Mar. 1987, *55* (2), 303-328.

[18] — **and** —, "Regulating Trade Among Agents," *Journal of Institutional and Theoretical Economics*, Mar. 1990, *146* (1), 85-105.

[19] — **and** —, "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,"

Journal of Law, Economics, and Organization, Jan. 1991, *7* (Special Issue: Papers from the Conference on the New Science of Organization), 24-52.

[20] **Jehiel, Philippe**, "On Transparency in Organizations," *Review of Economic Studies*, Apr. 2015, *82* (2), 736-761.

[21] **Jehiel, Philippe and Andrew F. Newman**, "Loopholes: Social Learning and the Evolution of Contract Form, " *Working Paper*, 2011.

[22] **Larkin, Ian**, "The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales," *Journal of Labor Economics*, Apr. 2014, *32* (2), 199-227.

[23] **Lazear, Edward P.**, "Speeding, Terrorism, and Teaching to the Test," *Quarterly Journal of Economics*, Aug. 2006, *121* (3), 1029-1061.

[24] **MacDonald, Glenn and Leslie M. Marx**, "Adverse Specialization," *Journal of Political Economy*, Aug. 2001, *109* (4), 864-899.

[25] **Milgrom, Paul R. and D. John Roberts**, *Economics, Organization and Management*, Englewood Cliffs: Prentice-Hall, 1992.

[26] **Mirrlees, James A.**, "Notes on Welfare Economics, Information and Uncertainty," in M. Balch, D. McFadden, and S. Wu, eds., *Essays in Equilibrium Behavior under Uncertainty*, Amsterdam: North-Holland, 1974.

[27] **Osler, Mark,** "The Case Against Transparency in Law School Rankings Methodology," *Law School Innovation*, May 4, 2010, http://lsi.typepad.com/lsi/2010/05/the-case-against-transparency-in-law-school-rankings-methodology.html.

[28] **Oyer, Paul**, "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," *Quarterly Journal of Economics*, Feb. 1998, *113* (1), 149-185.

[29] **Pestieau, Pierre, Possen, Uri M. and Steven M. Slutsky**, "The Value of Explicit Randomization in the Tax Code," *Journal of Public Economics*, Jan. 1998, 67(1), 87-103.

[30] **Prendergast, Canice**, "The Provision of Incentives in Firms," *Journal of Economic Literature*, Mar. 1999, 37 (1), 7-63.

[31] **Scott, Robert E. and George G. Triantis**, "Anticipating Litigation in Contract Design," *Yale Law Journal*, Jan. 2006, *115* (4), 814-879.

[32] **Stiglitz, Joseph**, "Self-selection and Pareto Efficient Taxation," *Journal of Public Economics*, Mar. 1982, *17* (2), 213–240.

[33] **Weisbach, David A.**, "An Efficiency Analysis of Line Drawing in the Tax Law," *Journal of Legal Studies*, Jan. 2000, *29* (1), 71-97.

A Omitted Proofs

Proof of Lemma 1. Consider first the pair of contracts (C_1^{bal}, C_2^{bal}). Under our assumption on the agent's behavior when indifferent over effort pairs, C_i^{bal} induces agent i to choose $e_i = e_j = \frac{\beta}{1+\lambda}$, yielding each type i a certainty equivalent of

$$ACE_i(C_i^{bal}) = E(w_i) - c_i(e_1, e_2) - \frac{1}{2}r\sigma^2 var(w_i) = \alpha + \beta^2 - \frac{\beta^2}{2} - \frac{1}{2}r\sigma^2\beta^2(1 + 2\rho\lambda + \lambda^2).$$

The principal will set α to satisfy each type's participation constraint with equality, and her expected payoff from each type, as a function of β, will be

$$\Pi^{bal}(\beta) = \frac{\beta}{1+\lambda} - \frac{\beta^2}{2} - \frac{1}{2}r\sigma^2\beta^2(1 + 2\rho\lambda + \lambda^2). \tag{20}$$

With β chosen optimally, the resulting maximized payoff is

$$\Pi^{bal} = \frac{1}{2(1+\lambda)^2\left[1 + r\sigma^2(1 + 2\rho\lambda + \lambda^2)\right]}. \tag{21}$$

This payoff is continuous as $\lambda \to 1$.

Now consider the pair of contracts (C_1^{foc}, C_2^{foc}). C_i^{foc} induces type i to choose $e_i = \beta$ and $e_j = 0$. The principal

will set α to satisfy each type's participation constraint with equality, and her expected payoff from each type, as a function of β, will then be

$$\Pi^{foc}(\beta) = \frac{\beta}{\delta+1} - \frac{\beta^2}{2} - \frac{1}{2} r\sigma^2\beta^2\left(1-\rho^2\right).$$

With β chosen optimally, the resulting maximized payoff is

$$\Pi^{foc} = \frac{1}{2(\delta+1)^2[1+r\sigma^2(1-\rho^2)]}. \tag{22}$$

Comparison of the expressions for Π^{bal} and Π^{foc} shows that there is a critical value of δ,

$$\delta^{NHI}(\lambda, r\sigma^2, \rho) \equiv (\lambda+1)\left[\frac{1+r\sigma^2(1+2\rho\lambda+\lambda^2)}{1+r\sigma^2(1-\rho^2)}\right]^{\frac{1}{2}} - 1, \tag{23}$$

above (below) which $\Pi^{bal} > (<)\, \Pi^{foc}$. It is straightforward to verify that δ^{NHI} is increasing in each of its arguments. ■

Proof of Lemma 2. For a transparent menu of linear contracts to induce both types of agent to exert strictly positive efforts on both tasks, it is necessary that each type be induced to choose a contract that equates the (constant) ratio of the marginal benefits of efforts on the tasks to the (constant) ratio, for that type, of the marginal costs. Therefore, if such a menu existed, it would have the form

$$C_1 : \; w_1 = \alpha_1 + \beta_1 x_1 + \lambda\beta_1 x_2 \quad \text{and} \quad C_2 : \; w_2 = \alpha_2 + \beta_2 x_2 + \lambda\beta_2 x_1,$$

and would induce agent i to choose C_i.

Let $ACE_i(C_j)$ denote the certainty equivalent achieved by agent i from selecting contract C_j and choosing efforts optimally. For agent 1 to be willing to choose C_1 requires $ACE_1(C_1) \geq ACE_1(C_2)$, and the analogous self-selection constraint for agent 2 is $ACE_2(C_2) \geq ACE_2(C_1)$. Now for all $\lambda > 1$, $ACE_2(C_1) > ACE_1(C_1)$, since agent 1's certainty equivalent from contract C_1 equals that which he would obtain from focusing all his effort on task 1 (which is one of his optimal effort allocations), whereas agent 2's certainty equivalent from C_1 equals that which he would obtain from focusing all his effort on task 2 (which is his unique optimal effort choice), and task 2 is more highly rewarded than task 1 in contract C_1. Similarly, for all $\lambda > 1$, $ACE_1(C_2) > ACE_2(C_2)$. If $ACE_1(C_1) \geq ACE_2(C_2)$, then $ACE_2(C_1) > ACE_1(C_1)$ implies that $ACE_2(C_1) > ACE_2(C_2)$, so the self-selection constraint for agent 2 would be violated. If, instead, $ACE_1(C_1) < ACE_2(C_2)$, then $ACE_1(C_2) > ACE_2(C_2)$ implies that $ACE_1(C_1) < ACE_1(C_2)$, so the self-selection constraint for agent 1 would be violated. Therefore, there is no way to choose $(\alpha_1, \beta_1, \alpha_2, \beta_2)$ so that the menu above induces both types of privately-informed agent to choose the contract that would make each willing to choose perfectly balanced efforts. Hence perfectly balanced efforts from both types of agents cannot be achieved.

Furthermore, since faced with a menu of transparent linear contracts, an agent either is willing to exert perfectly balanced efforts or strictly prefers fully focused efforts, this argument also shows that it is not possible for the principal to induce both types of agent to exert strictly positive efforts on both tasks. ■

Proof of Proposition 1.
Part (i): Consider first an STM, consisting of the contract pair

$$C_1^{STM} : \; w_1 = \alpha + \beta x_1 - \rho\beta x_2 \quad \text{and} \quad C_2^{STM} : \; w_2 = \alpha + \beta x_2 - \rho\beta x_1.$$

Agent i strictly prefers contract C_i to contract C_j and, having chosen C_i, will then set $e_i = \beta$ and $e_j = 0$. This STM generates the same outcome, for each type of agent, as the principal achieves in the NHI benchmark setting from the contract pair (C_1^{foc}, C_2^{foc}). Therefore, the principal's maximized payoff from an STM, Π^{STM}, is given by the expression in (22). Compared to an STM, an ST contract would, for all $\lambda > 1$, also induce fully focused efforts from both agent types but would impose a larger risk premium and hence generate a lower payoff for the principal.

Now consider an ATM, consisting of the contract pair

$$C_1^{ATM} : \; w_1 = \alpha_1 + \beta_1 x_1 - \rho\beta_1 x_2 \quad \text{and} \quad C_2^{ATM} : \; w_2 = \alpha_2 + \beta_2 x_2 + \lambda\beta_2 x_1.$$

In this menu, C_i is the contract intended for agent i. If agent 2 chooses C_2, he would be indifferent over all effort

pairs such that $\beta_2 = e_1 + \lambda e_2$. Given our assumption on the agent's behavior when indifferent, agent 2 chooses the perfectly balanced effort allocation $e_1 = e_2 = \frac{\beta_2}{1+\lambda}$. If, instead, agent 2 chooses C_1, he would set $e_1 = \frac{\beta_1}{\lambda^2}$ and $e_2 = 0$. If agent 1 chooses C_1, he would set $e_1 = \beta_1$ and $e_2 = 0$, whereas if he chooses C_2, he would set $e_1 = \lambda\beta_2$ and $e_2 = 0$.

The certainty equivalents that each of C_1 and C_2 offers to each type of agent are:

$$ACE_1(C_1) = \alpha_1 + \frac{(\beta_1)^2}{2} - \frac{1}{2}r\sigma^2(\beta_1)^2\left(1-\rho^2\right) \; ; \; ACE_1(C_2) = \alpha_2 + \frac{(\lambda\beta_2)^2}{2} - \frac{1}{2}r\sigma^2(\beta_2)^2\left(\lambda^2 + 2\rho\lambda + 1\right);$$

$$ACE_2(C_2) = \alpha_2 + \frac{(\beta_2)^2}{2} - \frac{1}{2}r\sigma^2(\beta_2)^2\left(\lambda^2 + 2\rho\lambda + 1\right) \; ; \; ACE_2(C_1) = \alpha_1 + \frac{(\beta_1)^2}{2\lambda^2} - \frac{1}{2}r\sigma^2(\beta_1)^2\left(1-\rho^2\right).$$

Since the principal is equally likely to be facing each type of agent, her problem is to choose $(\alpha_1, \beta_1, \alpha_2, \beta_2)$ to maximize

$$\frac{1}{2}\left[\frac{\beta_1}{\delta+1} - \alpha_1 - (\beta_1)^2\right] + \frac{1}{2}\left[\left(\frac{\beta_2}{1+\lambda}\right) - \alpha_2 - (\beta_2)^2\right],$$

subject to participation and self-selection constraints for both types of agent:

$$\begin{array}{lclcl} ACE_2(C_2) & \geq & 0 & \text{and} & ACE_2(C_2) \geq ACE_2(C_1), \\ ACE_1(C_1) & \geq & 0 & \text{and} & ACE_1(C_1) \geq ACE_1(C_2). \end{array}$$

Since for all $\lambda > 1$ we have $ACE_1(C_2) > ACE_2(C_2)$, agent 1's participation constraint will not bind, and hence agent 1 earns an "information rent".

For the two self-selection constraints to be satisfied simultaneously, it is necessary that $\beta_1 \geq \lambda\beta_2$. For given (β_1, β_2), it is optimal for the principal to set α_2 so agent 2's participation constraint binds and to set α_1 so agent 1's self-selection constraint binds. Then the constraint $\beta_1 \geq \lambda\beta_2$ is both necessary and sufficient for agent 2 to be willing to choose C_2. We may then restate the principal's problem as

$$\max_{\beta_1,\beta_2}\left\{\begin{array}{l} \frac{1}{2}\left[\frac{\beta_1}{\delta+1} - \frac{(\beta_1)^2}{2} - \frac{1}{2}r\sigma^2(\beta_1)^2\left(1-\rho^2\right) - \left(\lambda^2-1\right)\frac{(\beta_2)^2}{2}\right] \\ +\frac{1}{2}\left[\left(\frac{\beta_2}{1+\lambda}\right) - \frac{(\beta_2)^2}{2} - \frac{1}{2}r\sigma^2(\beta_2)^2\left(\lambda^2+2\rho\lambda+1\right)\right] \end{array}\right\} \qquad \text{s.t.} \quad \beta_1 \geq \lambda\beta_2.$$

There exists a $\hat{\delta}$ such that the constraint $\beta_1 \geq \lambda\beta_2$ will be binding at the optimum if and only if $\delta \geq \hat{\delta}$. If $\delta < \hat{\delta}$, then the principal's maximized payoff from this "unconstrained" ATM (ATMU) is

$$\Pi^{ATMU} = \frac{1}{4(\delta+1)^2}\left[\frac{1}{1+r\sigma^2(1-\rho^2)} + \frac{(1+\delta)^2}{(1+\lambda)^2}\frac{1}{\lambda^2 + r\sigma^2(\lambda^2+2\rho\lambda+1)}\right],$$

whereas if $\delta \geq \hat{\delta}$, then his maximized payoff from the "constrained" ATM (ATMC) is

$$\Pi^{ATMC} = \frac{\left(\lambda^2+\lambda+\delta+1\right)^2}{8(\delta+1)^2(1+\lambda)^2\left\{\lambda^2 + r\sigma^2\left[\left(1-\frac{\rho^2}{2}\right)\lambda^2 + \rho\lambda + \frac{1}{2}\right]\right\}}.$$

It remains to determine whether an ATM (unconstrained or constrained) or an STM is optimal. It can be checked that the crucial comparison is between Π^{STM} and Π^{ATMU} and furthermore that if

$$\delta < \delta^{HI}(\lambda, r\sigma^2, \rho) \equiv (\lambda+1)\sqrt{\frac{\lambda^2 + r\sigma^2(\lambda^2+2\rho\lambda+1)}{1+r\sigma^2(1-\rho^2)}} - 1,$$

then the best STM dominates the best ATM, whereas if $\delta > \delta^{HI}(\lambda, r\sigma^2, \rho)$, then the best ATM dominates the best STM. We have $\delta^{HI} < \hat{\delta}$. This proves part (i).

Part (ii): This is easily confirmed algebraically.

Part (iii): This is proved in the second paragraph of the text following the statement of the proposition.

Part (iv): For $\delta > \delta^{HI}(1, r\sigma^2, \rho) = \delta^{NHI}(1, r\sigma^2, \rho)$, the limit as $\lambda \to 1$ of the principal's maximized payoff in

the NHI benchmark is the limit as $\lambda \to 1$ of Π^{bal}, as given in equation (21). Under hidden information, when $\delta > \delta^{HI}(1, r\sigma^2, \rho)$, the principal's best transparent menu for λ sufficiently close to 1 is an ATM. We know that $\Pi^{ATMC} \leq \Pi^{ATMU}$, and it is easy to confirm algebraically that for $\delta > \delta^{HI}(1, r\sigma^2, \rho)$,

$$\lim_{\lambda \to 1} \Pi^{ATMU} < \lim_{\lambda \to 1} \Pi^{bal}.$$

■

Proof of Proposition 2.

Parts (i) and (ii): For each type of agent, let $\overline{e}$ (respectively, $\underline{e}$) denote effort on his less costly (respectively, more costly) task, and define $\overline{x}$ and $\underline{x}$ analogously. Under EAR, with probability $\frac{1}{2}$, $w = \alpha + \beta\overline{x} + k\beta\underline{x}$, in which case we let $\overline{EU}$ denote an agent's expected utility, and with probability $\frac{1}{2}$, $w = \alpha + \beta\underline{x} + k\beta\overline{x}$, in which case we denote expected utility by $\underline{EU}$.

Recall that $k \in (-1, 1)$. Each agent's unconditional expected utility under EAR is

$$\begin{aligned}
\frac{1}{2}\overline{EU} + \frac{1}{2}\underline{EU} &= -\frac{1}{2}E\exp\left\{-r\left[\alpha + \beta\overline{x} + k\beta\underline{x} - \frac{1}{2}(\overline{e} + \lambda\underline{e})^2\right]\right\} - \frac{1}{2}E\exp\left\{-r\left[\alpha + \beta\underline{x} + k\beta\overline{x} - \frac{1}{2}(\overline{e} + \lambda\underline{e})^2\right]\right\} \\
&= -\frac{1}{2}\exp\left\{-r\left[\alpha + \beta\overline{e} + k\beta\underline{e} - \frac{r}{2}\sigma^2\beta^2(1 + 2\rho k + k^2) - \frac{1}{2}(\overline{e} + \lambda\underline{e})^2\right]\right\} \\
&\quad - \frac{1}{2}\exp\left\{-r\left[\alpha + \beta\underline{e} + k\beta\overline{e} - \frac{r}{2}\sigma^2\beta^2(1 + 2\rho k + k^2) - \frac{1}{2}(\overline{e} + \lambda\underline{e})^2\right]\right\} \qquad (24)
\end{aligned}$$

Hence the first-order conditions for interior solutions for $\overline{e}$ and $\underline{e}$, respectively, are

$$\begin{aligned}
\frac{1}{2}\left[\beta - (\overline{e} + \lambda\underline{e})\right]\overline{EU} + \frac{1}{2}\left[k\beta - (\overline{e} + \lambda\underline{e})\right]\underline{EU} &= 0 \\
\frac{1}{2}\left[k\beta - \lambda(\overline{e} + \lambda\underline{e})\right]\overline{EU} + \frac{1}{2}\left[\beta - \lambda(\overline{e} + \lambda\underline{e})\right]\underline{EU} &= 0.
\end{aligned}$$

These first-order conditions can be rewritten as

$$\beta\overline{EU} + k\beta\underline{EU} = (\overline{e} + \lambda\underline{e})(\overline{EU} + \underline{EU}) \qquad (25)$$

$$k\beta\overline{EU} + \beta\underline{EU} = \lambda(\overline{e} + \lambda\underline{e})(\overline{EU} + \underline{EU}). \qquad (26)$$

Equations (25) and (26) in turn imply

$$\overline{EU} + k\underline{EU} = \frac{k}{\lambda}\overline{EU} + \frac{1}{\lambda}\underline{EU}.$$

If $k \in [\frac{1}{\lambda}, 1)$, then the left-hand side of this equation strictly exceeds the right-hand side, so in this case interior solutions for efforts cannot exist. This proves Part (i).

Adding the first-order conditions (25) and (26) and rearranging yields equation (1). Using (1) to substitute for aggregate effort $(\overline{e} + \lambda\underline{e})$ in (25) yields, after a little algebra, $(\lambda - k)\overline{EU} + (k\lambda - 1)\underline{EU} = 0$, which simplifies to equation (2).

Part (iii): Solving (2) for $\overline{e} - \underline{e}$ yields $\overline{e} - \underline{e} = [\ln(\frac{\lambda - k}{1 - k\lambda})]/[r\beta(1 - k)]$. For $k \in \left(-1, \frac{1}{\lambda}\right)$ and $\lambda > 1$, therefore, $\overline{e} - \underline{e}$ is greater than 0, increasing in λ and k, and decreasing in r. $(\overline{e} - \underline{e}) \to 0$ as $\lambda \to 1$, $k \to -1^+$, or $r \to \infty$.

Part (iv): Using (1) and (2) to substitute into (24), and then simplifying, allows us to express each type of agent's expected utility under EAR as

$$\frac{1}{2}\overline{EU} + \frac{1}{2}\underline{EU} = -\exp\left\{-r\left[\alpha + \beta\left(\overline{e} + k\underline{e}\right) - \frac{\beta^2\left(1 + k\right)^2}{2\left(\lambda + 1\right)^2} - \frac{1}{2}r\sigma^2\beta^2\left(1 + 2\rho k + k^2\right) - \frac{1}{r}\ln\left(\frac{1 + \frac{\lambda - k}{1 - k\lambda}}{2}\right)\right]\right\}.$$

Since both types receive the same expected utility, it is optimal for the principal to set α to ensure that their participation constraints bind. Setting α in this way (so that the whole expression in square brackets above is equal

to 0), the principal's expected payoff, for given (β, k), can be simplified to equation (3) as follows:

$$\begin{aligned}
\Pi^{EAR}(\beta,k) &= \frac{\delta\underline{e}+\overline{e}}{\delta+1} - \alpha - \frac{1}{2}\beta(\overline{e}+k\underline{e}) - \frac{1}{2}\beta(\underline{e}+k\overline{e}) \\
&= \frac{\delta\underline{e}+\overline{e}}{\delta+1} + \frac{1}{2}\beta(1-k)(\overline{e}-\underline{e}) - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2\left(1+2\rho k+k^2\right) - \frac{1}{r}\ln\left(\frac{1+\frac{\lambda-k}{1-k\lambda}}{2}\right) \\
&= \frac{\delta\underline{e}+\overline{e}}{\delta+1} - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2) - \frac{1}{2r}\ln\left[\frac{(\lambda+1)^2(1-k)^2}{4(1-k\lambda)(\lambda-k)}\right],
\end{aligned}$$

where the final line uses (2). ■

Proof of Proposition 3. Define $B \equiv \beta(1+k)$ and note, from (1), that aggregate effort $\overline{e}+\lambda\underline{e}$ is proportional to B. Using (1), (2), and $\beta = \frac{B}{1+k}$ to substitute into (3) yields (7) and (8) in the text. To prove the claims regarding the effect of varying δ, r (with $r\sigma^2$ fixed), or $\sigma^2(1-\rho)$ on the optimal level of k, we use (7) and (8) to examine the sign of the cross-partial derivative of $\Pi^{EAR}(B,k)$ with respect to k and the relevant parameter, holding B and hence aggregate effort fixed. For Part (iv), we examine the sign of the cross-partial derivative of $\Pi^{EAR}(B,k)$ with respect to k and B.

Part (i): Only the second term on the right-hand side of (8) generates a non-zero value of $\frac{\partial^2\Pi}{\partial\delta\partial k}$. As long as $\delta > \lambda$, $\frac{\partial^2\Pi}{\partial\delta\partial k} < 0$, so the optimal k decreases as δ increases.

Part (ii): With $r\sigma^2$ held fixed, only the second term on the right-hand side of (8) and the fourth term in (7) vary as r increases. Examining these terms shows that $\frac{\partial^2\Pi}{\partial r\partial k} > 0$, so as r decreases (holding $r\sigma^2$ fixed), the optimal k decreases.

Part (iii): $\frac{\partial\Pi}{\partial k}$ depends on σ^2 and ρ only via the third term in (7), and $\frac{\partial\Pi}{\partial k}$ is increasing in $\sigma^2(1-\rho)$, so the optimal k decreases as $\sigma^2(1-\rho)$ decreases.

Part (iv): $\frac{\partial^2\Pi}{\partial B\partial k} > 0$, so as the B to be induced decreases, the optimal k decreases. ■

Proof of Proposition 4.

Part (ia): Given what will be shown in parts (ib), (ii), and (iii), it suffices to focus here on the case where $\delta > \lambda > 1$ and where EAR, for the given (β,k) with $k \in (-1,1)$, induces interior optimal efforts. We will show that EAR yields a strictly lower expected payoff for the principal than a suitably designed contract of the form C_i^{bal}, as defined in Section 3.1.

Using equations (1) and (2) in Proposition 2 to substitute for $\overline{e}^{EAR}$ and $\underline{e}^{EAR}$ in equation (3), we have

$$\begin{aligned}
\Pi^{EAR}(\beta,k) &= \frac{\beta(1+k)}{(\lambda+1)^2} - \frac{\delta-\lambda}{\delta+1}\frac{\ln\left(\frac{\lambda-k}{1-k\lambda}\right)}{(\lambda+1)r\beta(1-k)} - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} \\
&\quad -\frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2) - \frac{1}{2r}\ln\left[\frac{(\lambda+1)^2(1-k)^2}{4(1-k\lambda)(\lambda-k)}\right] \\
&< \frac{\beta(1+k)}{(\lambda+1)^2} - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2).
\end{aligned}$$

The inequality follows from the assumptions that $\delta > \lambda > 1$ and $k > -1$ and the fact, proved in Part (i) of Proposition 2, that $k < \frac{1}{\lambda}$ is a necessary condition for EAR to induce interior optimal efforts.

If the principal knows which task the agent finds less costly, so we are in the NHI benchmark, the principal can induce the agent to choose perfectly balanced efforts by offering the type-i agent a contract of the form C_i^{bal}: $w = \alpha + \beta^{bal} + \lambda\beta^{bal}$ for some β^{bal}. By choosing $\beta^{bal} = \frac{\beta(1+k)}{\lambda+1}$, the principal can induce with C_i^{bal} the same aggregate effort as under EAR for the given values of β and k. Using $\beta^{bal} = \frac{\beta(1+k)}{\lambda+1}$ and equation (20), we can write the principal's payoff under C_i^{bal} as

$$\Pi^{bal}(\beta^{bal}) = \frac{\beta(1+k)}{(\lambda+1)^2} - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\frac{\beta^2(1+k)^2}{(\lambda+1)^2}(1+2\rho\lambda+\lambda^2).$$

Hence

$$\begin{aligned}\Pi^{bal}(\beta^{bal}) - \Pi^{EAR}(\beta,k) &> \frac{1}{2}r\sigma^2\beta^2(1+k)^2\left[\frac{(1+2\rho k+k^2)}{(1+k)^2} - \frac{(1+2\rho\lambda+\lambda^2)}{(\lambda+1)^2}\right] \\ &\geq 0,\end{aligned}$$

where the second inequality follows since $k < 1/\lambda$ and since $\frac{(1+2\rho k+k^2)}{(1+k)^2}$ is decreasing in k and equals $\frac{(1+2\rho\lambda+\lambda^2)}{(\lambda+1)^2}$ for $k = 1/\lambda$. The second inequality is strict for $\rho < 1$.

Part (ib): We will show that for $\lambda = 1$ and any given (β, k) with $k \in (-1, 1)$, EAR yields a weakly lower expected payoff for the principal than a suitably designed symmetric transparent (ST) contract, of the form defined in Section 3.1, and a strictly lower expected payoff if $\rho < 1$.

For $\lambda = 1$, aggregate effort under EAR is $\bar{e}^{EAR} + \lambda\underline{e}^{EAR} = \frac{\beta(1+k)}{2}$, and $\bar{e}^{EAR} = \underline{e}^{EAR} = \frac{\beta(1+k)}{4}$. Hence, for $\lambda = 1$, equation (3) simplifies to

$$\Pi^{EAR}(\beta,k) = \frac{\beta(1+k)}{4} - \frac{1}{8}\beta^2(1+k)^2 - \frac{1}{2}r(\sigma)^2\beta^2(1+2\rho k+k^2). \tag{27}$$

Consider now an ST contract with coefficient β^{ST} chosen to induce the same level of aggregate effort as under EAR for the given values of β and k: $\beta^{ST} = \frac{\beta(1+k)}{2}$. Since $\lambda = 1$, $\bar{e}^{ST} = \underline{e}^{ST} = \frac{\beta(1+k)}{4}$, so the ST contract also induces exactly the same effort levels on each task as EAR. The principal's payoff under the ST contract is

$$\Pi^{ST}(\beta^{ST}) = \frac{\beta^{ST}}{2} - \frac{1}{2}\left(\beta^{ST}\right)^2 - r\sigma^2\left(\beta^{ST}\right)^2(1+\rho) = \frac{\beta(1+k)}{4} - \frac{1}{8}\beta^2(1+k)^2 - \frac{1}{4}r\sigma^2\beta^2(1+k)^2(1+\rho). \tag{28}$$

Subtracting (27) from (28) yields

$$\Pi^{ST}(\beta^{ST}) - \Pi^{EAR}(\beta,k) = \frac{1}{2}r\sigma^2\beta^2[(1+2\rho k+k^2) - \frac{(1+\rho)}{2}(1+k)^2] = \frac{1}{4}r\sigma^2\beta^2(1-k)^2(1-\rho).$$

Hence with $\lambda = 1$ and $\rho < 1$, $\Pi^{ST}(\beta^{ST}) - \Pi^{EAR}(\beta,k) > 0$. If $\rho = \lambda = 1$, then $\Pi^{ST}(\beta^{ST}) - \Pi^{EAR}(\beta,k) = 0$.

Part (ii): When EAR induces a corner solution for efforts (so $\underline{e}^{EAR} = 0$), the first-order condition (25) for $\bar{e}^{EAR}$ reduces to:

$$\exp\left\{r\beta\bar{e}^{EAR}(1-k)\right\} = \frac{\beta - \bar{e}^{EAR}}{\bar{e}^{EAR} - k\beta}. \tag{29}$$

Since the left-hand side of (29) is strictly greater than 1 for $k < 1$, (29) implies that $\bar{e}^{EAR} < \frac{\beta(1+k)}{2}$. When EAR induces each type of agent to choose the corner solution $\left(\bar{e}^{EAR}, 0\right)$, each type's expected utility can be written as

$$\frac{1}{2}\overline{EU} + \frac{1}{2}\underline{EU} = -\exp\left\{-r\left[\alpha + \beta\bar{e} - \frac{1}{2}\bar{e}^2 - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2) - \frac{1}{r}\ln\left(\frac{1+\exp\{r\beta(1-k)\bar{e}\}}{2}\right)\right]\right\}.$$

The principal optimally sets α so that both types' participation constraints bind (i.e., so that the whole expression in square brackets above is 0). Setting α in this way, the principal's expected payoff, for given (β, k), can be simplified as follows:

$$\begin{aligned}\Pi^{EAR}(\beta,k) &= \frac{\bar{e}}{\delta+1} - \alpha - \frac{1}{2}\beta\bar{e} - \frac{1}{2}\beta k\bar{e} \\ &= \frac{\bar{e}}{\delta+1} + \frac{1}{2}\beta(1-k)\bar{e} - \frac{1}{2}\bar{e}^2 - \frac{1}{2}r\sigma^2\beta^2\left(1+2\rho k+k^2\right) - \frac{1}{r}\ln\left(\frac{1+\exp\{r\beta(1-k)\bar{e}\}}{2}\right) \\ &= \frac{\bar{e}}{\delta+1} - \frac{1}{2}\bar{e}^2 - \frac{1}{2}r\sigma^2\beta^2\left(1+2\rho k+k^2\right) - \frac{1}{2r}\ln\left(\frac{[1+\exp\{r\beta(1-k)\bar{e}\}]^2}{4\exp\{r\beta(1-k)\bar{e}\}}\right) \\ &< \frac{\bar{e}}{\delta+1} - \frac{1}{2}\bar{e}^2 - \frac{1}{2}r\sigma^2\beta^2\left(1+2\rho k+k^2\right),\end{aligned}$$

where the inequality follows since $\exp\left\{r\beta\bar{e}^{EAR}(1-k)\right\} > 1$.

Consider now an ST contract with incentive coefficient β^{ST} chosen to induce the same effort pair $\left(\bar{e}^{EAR}, 0\right)$ as under EAR for the given values of β and k: $\beta^{ST} = \bar{e}^{EAR}$. The principal's payoff under this ST contract is

$$\Pi^{ST}\left(\beta^{ST}\right) = \frac{\bar{e}^{EAR}}{\delta+1} - \frac{1}{2}\left(\bar{e}^{EAR}\right)^2 - r\sigma^2\left(\bar{e}^{EAR}\right)^2(1+\rho). \tag{30}$$

Therefore,

$$\begin{aligned}
\Pi^{ST}\left(\beta^{ST}\right) - \Pi^{EAR}(\beta, k) &> \frac{1}{2}r\sigma^2\left[\beta^2(1+2\rho k+k^2) - 2\left(\bar{e}^{EAR}\right)^2(1+\rho)\right] \\
&> \frac{1}{2}r\sigma^2\beta^2\left[(1+2\rho k+k^2) - \frac{(1+\rho)}{2}(1+k)^2\right] \\
&= \frac{1}{4}r\sigma^2\beta^2(1-k)^2(1-\rho) \\
&\geq 0,
\end{aligned}$$

where the second strict inequality follows from the fact that $\bar{e}^{EAR} < \frac{\beta(1+k)}{2}$.

Part (iii): We will show that, when $\lambda \geq \delta \geq 1$, with at least one of these inequalities strict, then for any (β, k) with $k \in (-1, 1)$ such that EAR induces strictly positive efforts on both tasks, EAR yields a strictly lower expected payoff for the principal than a suitably designed ST contract.

Starting from equation (3) in Proposition 2, we can write

$$\begin{aligned}
\Pi^{EAR}(\beta, k) &= \left(\frac{\delta\underline{e}^{EAR} + \bar{e}^{EAR}}{\delta+1}\right) - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2) - \frac{1}{2r}\ln\left[\frac{(\lambda+1)^2(1-k)^2}{4(1-k\lambda)(\lambda-k)}\right] \\
&< \left(\frac{\delta\underline{e}^{EAR} + \bar{e}^{EAR}}{\delta+1}\right) - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2) \\
&\leq \left(\frac{\lambda\underline{e}^{EAR} + \bar{e}^{EAR}}{\delta+1}\right) - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2) \\
&= \frac{1}{\delta+1}\frac{\beta(1+k)}{\lambda+1} - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - \frac{1}{2}r\sigma^2\beta^2(1+2\rho k+k^2).
\end{aligned}$$

The first inequality follows from the assumptions that $\lambda > 1$ and $k > -1$ and the fact that $k < \frac{1}{\lambda}$ is necessary for EAR to induce interior optimal efforts. The second inequality follows since $\lambda \geq \delta$, and the final equality follows from equation (1) in Proposition 2.

Consider now an ST contract with incentive coefficient β^{ST} chosen to induce the same aggregate effort as under EAR for the given values of β and k: $\beta^{ST} = \frac{\beta(1+k)}{1+\lambda}$. Since $\lambda > 1$, the ST scheme induces $\bar{e} = \beta^{ST}$, $\underline{e} = 0$, and the principal's payoff under this ST contract is

$$\Pi^{ST}(\beta^{ST}) = \frac{1}{\delta+1}\frac{\beta(1+k)}{\lambda+1} - \frac{\beta^2(1+k)^2}{2(\lambda+1)^2} - r\sigma^2\frac{\beta^2(1+k)^2}{(\lambda+1)^2}(1+\rho).$$

Hence

$$\begin{aligned}
\Pi^{ST}(\beta^{ST}) - \Pi^{EAR}(\beta, k) &> \frac{1}{2}r\sigma^2\beta^2\left[(1+2\rho k+k^2) - 2\frac{(1+k)^2}{(\lambda+1)^2}(1+\rho)\right] \\
&> \frac{1}{2}r\sigma^2\beta^2\left[(1+2\rho k+k^2) - \frac{(1+k)^2}{2}(1+\rho)\right] \\
&= \frac{1}{4}r\sigma^2\beta^2(1-k)^2(1-\rho) \\
&\geq 0.
\end{aligned}$$

The second strict inequality follows since $\lambda > 1$. ■

B Online Appendix: Not for Publication

B.1 Menus of Opaque Incentive Schemes

This section shows that the performance of EAR cannot be improved by the use of menus. Consider the following incentive-compatible menu of two incentive schemes, each involving randomization. For $k \in (-1, 1)$, Scheme $i \in \{1, 2\}$, intended for the agent who prefers task i, specifies that with probability $p \in (\frac{1}{2}, 1)$, $w = \alpha + \beta x_i + k\beta x_j$, and with probability $1 - p$, $w = \alpha + \beta x_j + k\beta x_i$. As $p \to 1/2$, the two schemes become identical, so the menu reduces to EAR.

The value of p has no effect on aggregate effort. However, as p rises, each type of agent faces less uncertainty about his compensation schedule, hence has weaker incentives to self-insure by balancing his effort choices, so the induced effort gap $\overline{e} - \underline{e}$ rises. In this respect, a larger p mirrors the effect of a larger weighting parameter k. Nevertheless, there is a crucial difference between p and k. An increase in k improves the diversification of the risk from the shocks to measured performance. However, because, regardless of the value of p, the agent is ultimately paid either $\alpha + \beta x_1 + k\beta x_2$ or $\alpha + \beta x_2 + k\beta x_1$, changes in p have no effect on the diversification of this risk.

In consequence, whereas Proposition 3 and Section 6 showed that the weighting factor k is a valuable instrument in the design of opaque schemes, we have the following negative conclusion for the role of p: If a symmetric menu of randomized schemes with parameters (β, k, p) induces interior solutions for efforts, then as long as $\delta > \lambda$, the principal's payoff will be increased by lowering p to 1/2, thus replacing such a menu by EAR as analyzed in Section 4. Hence the principal's payoff from EAR cannot be augmented by the use of menus.

B.2 Beyond the Exponential-Normal Model

Our findings that opaque incentive schemes induce more balanced efforts than symmetric transparent ones and do so in a way more robust to hidden information of the agent, apply even outside the exponential-normal framework. Let the measurement technology remain $x_i = e_i + \varepsilon_i$, but now let $(\varepsilon_1, \varepsilon_2)$ have an arbitrary symmetric joint density. Let each type of agent's utility be $U(w - c(\overline{e}, \underline{e}))$, with $U(\cdot)$ an arbitrary strictly concave function and $c(\overline{e}, \underline{e})$, as in (14), reflecting imperfect substitutability of efforts.

Under EAR, interior optimal effort choices for each type of agent satisfy

$$\frac{\partial c}{\partial \overline{e}} + \frac{\partial c}{\partial \underline{e}} = \beta(1 + k) \qquad \text{and} \qquad \frac{E\left[U'(\cdot) I_{\{\underline{x} \text{ is more highly rewarded}\}}\right]}{E\left[U'(\cdot) I_{\{\overline{x} \text{ is more highly rewarded}\}}\right]} = \frac{\frac{\partial c / \partial \underline{e}}{\partial c / \partial \overline{e}} - k}{1 - k \frac{\partial c / \partial \underline{e}}{\partial c / \partial \overline{e}}}.$$

The second equation is a generalized version of (2) and shows that just as for the exponential-normal model, EAR gives the risk-averse agent an incentive to choose more balanced efforts to partially self-insure against the risk stemming from the uncertainty about which payment schedule will ultimately be used.

Nevertheless, we can show that whenever the symmetric transparent contract induces interior efforts, EAR does as well, and effort choices under EAR are more balanced than under the ST contract. Moreover, when efforts are perfect substitutes for the agent ($s = 1$), as λ increases from 1, $\overline{e}^{EAR}/\underline{e}^{EAR}$ increases continuously from 1, whereas $\overline{e}^{ST}/\underline{e}^{ST}$ jumps from 1 to ∞. Thus, even outside the exponential-normal framework, EAR provides stronger incentives for effort balance and is more robust to hidden information.

B.3 Ex Ante Randomization and the Choice of How Many Tasks to Reward

In Section 7.3 we discussed the trade-offs involved in the design of randomized incentive schemes in environments with many tasks. In this section we provide the derivations for our results.

Consider an EAR scheme in which each subset of κ out of n tasks is chosen with equal probability, and each task in the chosen subset is rewarded at rate β. Since this scheme is symmetric with respect to all n tasks and since each type of agent's preferences are symmetric with respect to each of his $n - 1$ "non-disliked" tasks, each agent's optimal effort profile can be described by $\underline{e}$, his effort on his disliked task, and by $\overline{e}$, his effort on each of the other tasks. If the task that an agent dislikes is included (respectively, not included) in the chosen subset, denote his (conditional) expected utility by $\underline{EU}$ (respectively, $\overline{EU}$). For any given task, the number of subsets that include it is $\binom{n-1}{\kappa-1}$, while

the number that do not is $\binom{n}{\kappa} - \binom{n-1}{\kappa-1} = \binom{n-1}{\kappa}$. Hence each type of agent's unconditional expected utility is

$$\frac{\binom{n-1}{\kappa}}{\binom{n}{k}}\overline{EU} + \frac{\binom{n-1}{\kappa-1}}{\binom{n}{k}}\underline{EU}.$$

We focus on the case where optimal efforts are interior.

The aggregate effort exerted by an agent is $\lambda\underline{e} + (n-1)\,\overline{e}$, which we define as A. To find the optimal level of A, we equate the sum over all tasks of the expected marginal monetary returns to effort to the sum over all tasks of the marginal cost of effort. Formally, this corresponds to adding the first-order conditions for effort on each of the n tasks. This yields $\kappa\beta = (n-1+\lambda)A$, so the optimal level of $A = \frac{\kappa\beta}{n-1+\lambda}$. To derive the optimal value of $\overline{e} - \underline{e}$, we need the first-order condition for $\underline{e}$, which is

$$\binom{n-1}{\kappa-1}[\beta - \lambda A]\,\underline{EU} + \binom{n-1}{\kappa}[-\lambda A]\,\overline{EU} = 0, \tag{31}$$

since the net marginal monetary return to $\underline{e}$ is $\beta - \lambda A$ if the subset of rewarded tasks includes the agent's disliked one and is $-\lambda A$ otherwise. Substituting for the optimal value of A in (31) and rearranging yields

$$\overline{e} - \underline{e} = \frac{1}{r\beta}\ln\left[\frac{\lambda\,(n-\kappa)}{n-1-(\kappa-1)\,\lambda}\right].$$

A necessary condition for interior solutions is $k-1 \leq \frac{n-1}{\lambda}$. Each type of agent's unconditional expected utility is given by

$$EU = -\frac{\binom{n-1}{\kappa-1}}{\binom{n}{\kappa}}\exp\left\{-r\left[\alpha + \beta\left((\kappa-1)\,\overline{e} + \underline{e}\right) - \frac{1}{2}\frac{\kappa^2\beta^2}{(\lambda+n-1)^2} - \frac{1}{2}r\sigma^2\beta^2\kappa\,(1+\rho(\kappa-1))\right]\right\}$$
$$- \frac{\binom{n-1}{\kappa}}{\binom{n}{\kappa}}\exp\left\{-r\left[\alpha + \beta\kappa\overline{e} - \frac{1}{2}\frac{\kappa^2\beta^2}{(\lambda+n-1)^2} - \frac{1}{2}r\sigma^2\beta^2\kappa\,(1+\rho(\kappa-1))\right]\right\}.$$

The principal will optimally set α to ensure that the participation constraint binds for each type of agent. With α set in this way, and using the expressions for each type of agent's optimal choices of A and $\overline{e} - \underline{e}$, the principal's expected payoff as a function of β and κ can be simplified to

$$\Pi(\beta,\kappa) = \frac{\delta\underline{e} + (n-1)\overline{e}}{\delta+n-1} - \frac{\kappa^2\beta^2}{2\,(\lambda+n-1)^2}$$
$$- \frac{1}{2}r\sigma^2\beta^2\kappa\,(1+\rho(\kappa-1)) - \frac{1}{nr}\ln\left[\frac{(n-\kappa)^{n-\kappa}(n-1+\lambda)^n}{n^n\lambda^\kappa\left((n-1)-(\kappa-1)\lambda\right)^{n-\kappa}}\right], \tag{32}$$

where

$$\frac{\delta\underline{e} + (n-1)\overline{e}}{\delta+n-1} = \frac{\kappa\beta}{(\lambda+n-1)^2} - \frac{(\delta-\lambda)(n-1)}{(\delta+n-1)(\lambda+n-1)r\beta}\ln\left[\frac{\lambda(n-\kappa)}{(n-1)-(\kappa-1)\lambda}\right]. \tag{33}$$

Using $\tilde{\beta} = \kappa\beta$ to substitute for β in the above payoff expression yields expressions (18) and (19) in the text. Verifying the claims in the final paragraph of Section 7.3 regarding the effect of varying δ, r (with $r\sigma^2$ fixed), or $\sigma^2(1-\rho)$, on the optimal value of κ requires signing the cross-partial derivative of $\Pi(\tilde{\beta},\kappa)$ in (18) with respect to κ and the relevant parameter, holding $\tilde{\beta}$ fixed. We can show that $\frac{\partial^2\Pi}{\partial\delta\partial\kappa} < 0$, $\frac{\partial^2\Pi}{\partial r\partial\kappa} > 0$, and $\frac{\partial^2\Pi}{\partial(\sigma^2(1-\rho))\partial\kappa} > 0$, from which the claims follow. The final claim follows from the fact that $\frac{\partial^2\Pi}{\partial\tilde{\beta}\partial\kappa} > 0$.