



HARVARD Kennedy School
JOHN F. KENNEDY SCHOOL OF GOVERNMENT

Solomonic Separation: Risk Decisions as Productivity Indicators

Faculty Research Working Paper Series

Nolan Miller

University of Illinois at Urbana-Champaign

Alexander F. Wagner

University of Zurich

Richard J. Zeckhauser

Harvard Kennedy School

November 2012
RW12-057

Visit the **HKS Faculty Research Working Paper** series at:
<http://web.hks.harvard.edu/publications>

The views expressed in the **HKS Faculty Research Working Paper Series** are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or of Harvard University. Faculty Research Working Papers have not undergone formal review and approval. Such papers are included in this series to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Solomonic Separation: Risk Decisions as Productivity Indicators

Forthcoming, Journal of Risk and Uncertainty

Nolan Miller, Alexander F. Wagner, and Richard J. Zeckhauser*

November 19, 2012

Abstract

A principal provides budgets to agents (e.g., divisions of a firm or the principal's children) whose expenditures provide her benefits, either materially or because of altruism. Only agents know their potential to generate benefits. We prove that if the more "productive" agents are also more risk-tolerant (as holds in the sample of individuals we surveyed), the principal can screen agents and bolster target efficiency by offering a choice between a nonrandom budget and a two-outcome risky budget. When, at very low allocations, the ratio of the more risk-averse type's marginal utility to that of the other type is unbounded above (e.g., as with CRRA), the first-best is approached. – A biblical opening enlivens the analysis.

JEL-Code: D82

Keywords: Asymmetric information, capital budgeting, random mechanisms, risk aversion, screening, target efficiency

*Miller is a Professor of Finance at the University of Illinois at Urbana-Champaign. Wagner is an Assistant Professor of Finance at the Swiss Finance Institute – University of Zurich, a Research Associate at Harvard University, and a CEPR Research Affiliate. Zeckhauser is the Frank Plumpton Ramsey Professor of Political Economy at Harvard University. Office address (Zeckhauser is the corresponding author): 79 JFK St. Cambridge, MA 02138, USA. Phone: (617) 495-1174, Email: richard_zeckhauser@harvard.edu. We thank John Pratt for help with the proof of Lemma 1, Carmen Tanner and Michael Kosfeld for help with the survey, Mario Häfeli and Ramona Westermann for research assistance, and the editor, the referee, and participants in various seminars for helpful comments. Wagner thanks the Swiss Finance Institute, the NCCR FINRISK and the University of Zurich Research Priority Program "Finance and Financial Markets" for support.

1 Introduction

King Solomon, thought by some to be the wisest man who ever lived, anticipated the economists' concept of separating equilibria by about 3,000 years. In his most famous case, he proposed cutting a baby in half to separate the true mother and the false mother. The true mother said: "No, give him to the other woman," whereas the claimed mother accepted the proposed deal. Not only did Solomon perceive a difference in risk preferences – he knew the true mother would not accept even a small chance of slicing the baby in half – but he anticipated that the false mother would not figure out how to pose as the true mother. The baby was placed in the true mother's arms.

Recent archeological finds discovered a lost scroll that detailed another separation decision by Solomon, where once again he used risk taking to gauge preference intensity. That decision merits equal acclaim for its wisdom.

2 The Scroll of the Scribes

2.1 The problem

One day a wealthy man came to Solomon for advice. He observed: "I have two sons, X and Y. They are both fine boys, and help me administer my business. I do not spoil them, but they both receive an adequate income. Alas, the great sadness of my life is that they do not get along, and I must keep them apart so they do not quarrel. When I die, and fortunately my health is still good, one must get my business. The other will receive my worldly possessions, but alas the division will be unequal. The business is worth far more, and the burden to run it is not great. I can not rely on either to provide an income interest to the other.

My sons are equally capable, and I love them equally. Today, knowing what the future portends, they both spend what they receive. But I know that some people receive more pleasure from consumption expenditures than do others. I would like to leave my business to the son who receives the greater pleasure. However, when I ask them, they both say their pleasure is immense. How shall I decide?"

Solomon responded. "The day after the second new moon, bring your sons to me, and we shall resolve this problem. I have but one constraint. You must let me resolve this problem, and you must remain silent as I do so." The man agreed. The appointed day arrived, and the wealthy man

and his two sons appeared before the King.

Solomon then spoke to the sons. “Alas, the two of you do not get along. When your father passes from this Earth, his wish is that one of you receive his business, and the other his worldly possessions. You will then have no need for further contact with each other.

“But wonderful things do not come without sacrifice. You see before you a large jar with a scorpion and some leaves. One of you will place his hand in this jar for a period of time to risk his sting. The scorpion may not see your hand for a while. But even when seen, it will not look like his natural prey; it may be ignored. But should the scorpion sting, it will be intensely painful, and perhaps worse. I have a papyrus scroll for each of you. You will each go to a corner of the room and write down how many minutes you are willing to leave your hand in the jar to be the one who inherits the business.”

Solomon then explained how he would conduct this as a second-price auction, and the virtues of that method. The father was sad, because he did not want either son to risk the scorpion’s sting, but he got false succor from the second-price auction, thinking that it would lead to less time at risk. But most important, as promised, he remained silent.

The sons returned to King Solomon and their father. X had written 2 minutes on his scroll. Y had written 30 minutes. Solomon then decried: “The business shall go to Y upon your father’s death, because he is the son I have determined would reap greater benefits from the excess income that would offer. Moreover, Y need not place his arm within the scorpion’s bottle. That would be a deadweight loss, conceivably in the literal sense of that term. I was confident that neither of you would decipher this game. Just as I had no intention of dividing the baby in an earlier decision, I had no intention of forcing either of you to take a dreaded risk.”

He then said: “Unlike judges in the democracies of future centuries, I do not have time to write down and justify my opinion. But I will explain to the court scribes the principles underlying my decision, so they may be recorded and available to future generations.”

The father did not understand what had happened, but Solomon was Solomon. Thus, he knew the decision was wise. The father lived to a ripe old age. When he died a happy death, Y took his business, X the worldly possessions.

As mentioned above, the scroll of the scribes has only recently come to light. It is reproduced here, together with contemporary comments from modern scholars.

2.2 Solomon's reasoning

King Solomon observed: "My job was to find a way to identify which of two sons would derive greater utility from a substantially increased income. I have spent many years receiving my many subjects, from rich, moderate and poor circumstances. I have struggled to perceive their levels of satisfaction. I have concluded that life in moderate or poor circumstances is much the same for all. But having riches separates men. Some are possessed of exquisite taste, and turn their riches to great consumptive pleasures, both for themselves and with their celebrations for the community. Others, alas, turn riches into little of value. They purchase ostentatiously to impress, and impress no one, not even themselves.

I label these groups *connoisseurs* and *boors*. A connoisseur benefits greatly from securing riches, and this possibility is, therefore, worth making great sacrifices for. Hardly so for the boor. My test was a simple one. Son Y showed himself to be a connoisseur by his willingness to take a substantial risk to win the business; son X gave away his boorish nature when he answered a mere two minutes.

I would like to claim originality for my method, but any fairy tale king who sent suitors into battle against dragons before they could claim his daughter's hand understood the underlying principle: Any hopeful dragon slayer faced a 20% chance of death, with only an 80% chance of blissful marriage to the princess. (History is written by the victors, which is why traditional accounts suggest better odds.) The fairy tale king – anticipating von Neumann and Morgenstern – recognized the implicit requirement:

$$.8U(\text{marriage to princess}) > U(\text{status quo}) - .2U(\text{death}).$$

Only the deeply devoted would have such a utility for marriage to the princess."

Contemporary comment. We will prove the wisdom of Solomon's idea below. As we show, a principal is indeed able to use risk taking as a gauge of preference intensity to great advantage. Not only can she separate connoisseurs from boors, but for a broad class of utility functions, such as constant relative risk aversion, she can approach the first-best. We also show that the typical lottery for a connoisseur involves the risk that he receives a very bad outcome (indeed, the worst possible outcome) with a very low probability.

There is suggestive evidence that in devising this screening mechanism Solomon may have

been influenced as well by family history. His father, King David, won his first wife, Michal (not Solomon's mom), in an equivalent test gamble. As you remember, Goliath repeatedly challenged the Israelites for forty days. David, then but a humble shepherd, responded when King Saul promised a reward to he who defeats Goliath. "And it will be that the king will enrich the man who kills him with great riches and will give him his daughter and make his father's house free in Israel." (Samuel 1, 17:25). (Saul, some believe, was not looking for devotion to his daughter. He recognized David as a future power threat and perhaps was hoping he would be killed by Goliath.)

Solomon continued: "I have now sought to generalize this method to help future adjudicators. My method, like the procedure of the fairy tale kings, employs lotteries, but death-by-dragon (or by scorpion) seems a rather extreme penalty. My methods employ only risk taking with money. Some day, I am confident, a highly respected profession will develop that studies money and decisions, and employs experiments. I sought to anticipate their methods. Thus, I conducted a survey among a sample of my subjects of moderate means.

1. Among our citizens, how much pleasure would you get from a consumption of 50,000 shekels per year? Please rate yourself on a percentile basis relative to your peers.

2. Say you were given a lottery offering a 50-50 chance of 20,000 or 100,000 shekels per year. What certain amount would make you just as well off as this lottery?

As I expected, there was a strong positive correlation between the answers to the two questions. That is, if we graph absolute utility as a function of income, the steeper curves were also straighter. To check for robustness, I then varied these amounts, but found that the pattern persisted. In these experiments with many subjects I discovered that risk aversion and reported pleasure from increased consumption were negatively correlated.

I expect researchers in the far future to retrogress, and to express skepticism about the use of surveys or any attempt to gauge interpersonal comparisons of utility. But I have the extreme research advantage of having ruled for 36 years, to have met regularly with my subjects, and to be blessed with what they call wisdom. This gives me the power to detect cheap talk, and to make it expensive.

Generalization can be of one's method, or of its areas of application. I have found other areas where citizens can be induced to reveal their true assessments by subjecting them to some risk. Thus, in dispensing plots for farming for citizens turning 21, beyond assuring adequate incomes

for all and well paid employment for young people, I seek to create prosperity for the kingdom. Thus, I wish to put substantially more land in the hands of high productivity workers. The more land combined with any worker, the less per hectare he will produce, but high productivity workers both get more output from the land and trim such diminution. The distinguishing feature of high productivity workers is their ability to manage young workers effectively. Thus, beyond the initial scale, their output increases linearly with land provided. I discovered that if I offer my subjects a choice, two hectares for sure, or a lottery offering an 80% chance of five hectares and a 20% chance of 1 hectare, productivity differentials make the lottery the best choice for the high productivity workers, the certain two hectares for those of low productivity.

I hope that this generality – the ability to address two quite different problems – will help my methods to gain use in the future.”

Contemporary comment. Solomon would rule for an additional four years. His two questions have been employed in contemporary surveys by the annotators. We shall return to our survey results near the end of the paper. As indicated by Solomon, his method also helps in resource allocations, where productivity of the agent, not personal benefits, are at stake, as is the case in corporate capital budgeting. Before discussing issues in practical application, we present the theory supporting the separation method identified by Solomon, what we consider to be the primary contribution of the paper.

2.3 Related work

Private information plays a major role in many economic settings. Screening mechanisms are widely used to address the problems that arise when information is asymmetrically held.¹ This paper relates to several strands of literature, including work in capital budgeting, random mechanism design, optimal taxation, and fair division. While each of these literatures provides us with interesting insights, none directly addresses the class of problems that Solomon needed to solve, i.e., gauging the benefit that an agent, who is the source of welfare for a principal, will generate from an amount of resources. We discuss the differences in turn.

The allocation of central resources to decentralized units is the canonical business example

¹The classic reference is Akerlof (1970). Salanié (2005) provides an excellent survey of the basic methodology and many applications.

that motivates this analysis. For example, a corporation center must provide resources – e.g., capital, marketing capability, R&D support, executive time – to its operating divisions. When studying such situations, it is often assumed that divisions and the center have divergent interests (see, for example, Harris and Townsend (1981), Harris, Kriebel, and Raviv (1982) and Antle and Eppen (1985), and Stein (1997)). By contrast, in our paper agent welfare feeds positively into the principal’s welfare function – an assumption that often captures the relationship between the center and the division receiving resources. Of course, the principal would still like to correctly gauge the productivity of each division, and then fund each division to the level where a dollar of resources equals a dollar of profit. A division, however, being undercharged for such resources, would like far more resources.

In the literature that does assume that the center benefits directly from the productivity of the agent, the approach has been to establish truthful revelation through auditing, or a compensation scheme, or both (see, for example, Harris and Raviv (1996), Bernardo, Cai, and Luo (2001) and Bernardo, Cai, and Luo (2004)). Our mechanism, by contrast, is based solely on the capital allocations themselves.

In both Solomon’s examples and our analysis below, random mechanisms elicit private information. Other work has also used random mechanisms.² For example, risk-averse agents often take on undesirable risk to signal to others certain desirable qualities. Thus, risk-averse heads of a start-up firms tend to retain substantial undiversified stakes to assure the market of their positive views of their firms’ prospects (Leland and Pyle 1977). By contrast, in the Moselle, Degeorge, and Zeckhauser (2005) model, good types communicate their quality by choosing less risk. In a different realm, buyer risk aversion can make it in a haggling seller’s interest to employ a possibly-final offer strategy, an offer, which if rejected may be the final offer made (Miller, Piankov, and Zeckhauser 2006). Probabilistic insurance policies can be theoretically appealing (though consumers are reluctant to accept them, an observation that Wakker, Thaler, and Tversky (1997) explain by

²See Arnott and Stiglitz (1988) for a typology of situations when randomization is desirable in adverse selection and moral hazard situations. They neither characterize the shape of the optimal lottery, discuss when the first-best can be obtained, nor consider agents as a source of utility. Random contracts may, for example, be optimal in a buyer-seller framework where the price discriminating seller (the principal) wishes to extract the highest surplus possible from the buyer (the agent). High-valuation buyers envy low-valuation buyers because the latter pay lower prices in the first-best. It turns out that random contracts only dominate deterministic contracts if the high-valuation buyer is more risk-averse than the low-valuation buyer (Maskin and Riley 1989). By contrast, Solomon considered a situation where the principal and the agent have proportional utility (except for the cost of resources to the principal). He thus came to the conclusion that the principal will be able to separate the two types if and only if the connoisseur is less risk-averse.

reference to the weighting function of prospect theory). Three innovations differentiate Solomonic Separation from these methods: (1) agent productivity plays a central role, (2) risk tolerance on the resource to be allocated and productivity are correlated, and (3), the principal's payoff is strictly proportional to that of the agent. Therefore, the principal would like to avoid variability in budgets.

The optimal income tax literature has also considered the potential advantages of randomization, in the form of random taxes. High-productivity agents are endogenously richer than low-productivity agents, and thus have lower marginal utility. The problem for the government is one of redistributing from high-productivity agents to low-productivity agents, i.e., from richer to poorer, subject to the constraint that individuals choose how much to work. Randomization of outcomes for low-productivity agents can theoretically alleviate incentive constraints for high-productivity agents sufficiently that the additional scope for redistribution outweighs the immediate welfare losses from the randomization (Weiss 1976, Stiglitz 1982, Brito, Hamilton, Slutsky, and Stiglitz 1995). Hellwig (2007) shows, however, that randomization of taxes is only welfare-enhancing with a particular type of increasing risk aversion; it is undesirable with weakly decreasing risk aversion. In Solomon's problem, unlike in the optimal income tax problem, individuals differ in their utility functions, not in their income earning abilities.

In fair division problems, the social planner uses agents' relative preferences between different types of goods to divide a set of goods efficiently, while retaining envy-freeness. (See Steinhaus (1948) for the original problem formulation, Brams and Taylor (1999) for a popular book on the subject, Pratt and Zeckhauser (1990) for a case study of a practical application, and Brams (2005) and Pratt (2005) for reviews of existing division rules and developments of more sophisticated division methods.) Solomon's approach instead gauges the relative preference between different amounts of the same good, namely money (though it could be anything else), to judge how strongly the agent likes the good.³

2.4 Plan of the paper

In what follows, we aim to formalize Solomonic Separation based on risk taking. Section 3 presents the model. It considers a (female) principal who wishes to distribute funds to a (male) agent. There

³Two formal differences are that Solomon works with an expected budget and ensures envy-freeness in the screening problem between different types of one agent, while the typical fair division problem involves a fixed budget and multiple agents. Envy-freeness is not always guaranteed in fair division procedures (see, for example, the market-like Gap procedure of Brams and Kilgour (2001)).

are two agent types who differ in their level of marginal benefits (either marginal productivity or marginal utility) that they derive from funds and in their risk tolerance. The central assumption of the model is that the agents who receive higher marginal benefits are also more risk tolerant. Given this condition, the analysis shows how the principal can screen agents by offering a choice between a nonrandom budget and a risky budget. Moreover, it proves that an optimizing principal need employ only two prizes with associated probabilities in the risky budget. The benefits of using this screening method can be substantial. Indeed, when the ratio of the more risk-averse type's marginal utility to that of the other type is unbounded above for a very low payoff (e.g., as with CRRA), the first-best outcome can be approached.

Section 4 discusses the results, and investigates whether marginal benefits and risk tolerance will be positively correlated, as the model requires. It provides an intuitive argument showing this is likely to be true in a production setting. For the consumption setting, it presents supportive evidence from survey results. These results, admittedly only suggestive, find a strong positive correlation between self-reports of utility gains from windfall funds and risk tolerance. This section contains additional results. For example, it identifies conditions under which our screening methods may be feasible but not beneficial.

Section 5 concludes, and offers examples of applications.

3 Basic model

Consider a (female) principal and a (male) agent. The principal has or can acquire a stock of funds she wishes to distribute to the agent. The principal's benefit is proportional to that of the agent. That is, the agent produces benefits for the principal through output from inputs. A corporate division acting as an agent for the corporate center (the principal) is our prime example. The analysis applies as well when the principal is doling out funds for consumption (as when a child is an agent for her parents).

Principal. Let x denote a quantity of funds that the principal allocates to the agent. Let x_{\min} denote the minimum quantity of funds that the principal must give the agent. Often, it will be reasonable to assume that $x_{\min} = 0$, i.e., that the principal cannot impose a penalty on the agent, while in other applications, negative payments will be allowed. The principal has a linear cost of

funds $K(x) = cx$, and we assume for convenience that the marginal cost of funds is unity, $c = 1$. Agent i produces benefits $V_i(x)$, though some agents generate greater benefits from the principal's resources than others in a sense made precise below. The principal derives benefits from the agent's use of resources. For example, both the center's and the division's manager have better career opportunities or benefit from the company's success by way of the compensation plan in place. In particular, we assume that if the principal gives resources x to a type- i agent, she receives surplus $V_i(x) - x$. The principal's and the agent's preferences diverge because the principal takes into account both the benefits and costs of funds, whereas the agent only cares about the benefits.

Agent. Let $i \in \{H, L\}$ denote the agent's type as High (H) or Low (L). Let $V_H(x)$ and $V_L(x)$ denote the two types' utility functions. The agent maximizes expected utility. In the corporate context, utility would be the output a division can produce utilizing central resources. Diminishing marginal product would produce the equivalent of risk aversion for a division that was risk neutral. Because the model is applicable to both the production and the consumption case, we speak of benefit functions.

For both agent types, the benefit function V_i is a strictly increasing and strictly concave function, i.e., $V_i' > 0$, $V_i'' < 0$. However, agent types differ in the benefit received from being given resources and in their risk tolerance.

Regarding marginal benefits, we assume that

$$V_H'(x) > V_L'(x) \text{ for all } x \text{ where } V_L'(x) \leq 1, \quad (1)$$

which means that the principal recognizes that an additional dollar is worth more to High than it is to Low in the region of large allocations, where Low's marginal benefit is below the marginal cost of funds.⁴

Regarding attitudes toward risk, the central assumption we posit is that High is less risk-averse than Low in the sense that High's Arrow-Pratt coefficient of absolute risk aversion is smaller than Low's:

⁴Note that the quantity x denotes the change of funds relative to initial wealth. High and Low may have different current wealth levels w_H and w_L . If for benefit functions $V_H(w_H+x)$ and $V_L(w_L+x)$ Low is more risk-averse than High at x and High has a greater marginal benefit than Low at x , then there are functions of x only $U_H(x) = V_H(w_H+x)$ and $U_L(x) = V_L(w_L+x)$ such that they have the same properties. Therefore, we drop w_H and w_L for notational convenience.

$$-\frac{V_H''(x)}{V_H'(x)} < -\frac{V_L''(x)}{V_L'(x)} \text{ for all } x. \quad (2)$$

Note that for (2) to hold, we do not need High's benefit function to have a smaller curvature ($-V_H''(x) < -V_L''(x)$) everywhere because High has a higher marginal benefit.

A natural sufficient condition for (1) to hold is that $V_H'(x) > V_L'(x)$ for all permissible x . However, for some utility functions, as allocations approach the lower bound, Low's benefits begin to fall more steeply without bound than High's. Indeed, some of our results will depend critically on whether the ratio of Low's marginal benefits to High's marginal benefits is bounded when allocations decrease, i.e., on whether there exists a finite bound C such that $\lim_{x \rightarrow x_{\min}} V_L'(x)/V_H'(x) = C < \infty$. This will be noted below when required.

There is a fraction λ of High types, and a fraction $(1 - \lambda)$ of Low types in the population, where λ is common knowledge.

Mechanism. In the main analysis, the agent's type is private information, giving rise to the need to design an incentive-compatible mechanism for distributing funds to the agent. The mechanism we study capitalizes on random allocations to induce the agent to reveal his type. An allocation for a type- i agent consists of a vector of prizes $\{x_{ij}\}$, denominated in dollars, and their attached probabilities $\{p_{ij}\}$, $0 \leq p_{ij} \leq 1$, such that $\sum_{j=1}^J p_{ij} = 1$. By this notation, we mean that the allocation is a lottery that gives type i agent x_{ij} with probability p_{ij} . We denote as J_i the number of prizes in the lottery designed for type i . When employing such a mechanism, the principal maximizes expected net benefits, taking into account both the likelihood of the two types, λ and $(1 - \lambda)$, and the randomization inherent in the lottery. That is, the principal chooses $\{x_{Hj}\}$, $\{p_{Hj}\}$, $\{x_{Lj}\}$, and $\{p_{Lj}\}$ in order to maximize:

$$\lambda \left[\sum_{j=1}^{J_H} p_{Hj} (V_H(x_{Hj}) - x_{Hj}) \right] + (1 - \lambda) \left[\sum_{j=1}^{J_L} p_{Lj} (V_L(x_{Lj}) - x_{Lj}) \right].$$

Discussion of marginal assumptions for consumption case. A brief discussion of the assumptions in (1) and (2), though not required for the corporate resources example, will be helpful for the pure consumption example. While von-Neumann-Morgenstern utility uniquely defines the Arrow-Pratt coefficient for a particular decision maker, marginal utility is not uniquely de-

fined due to the ability to scale preferences using a positive linear transformation (i.e., to multiply utility by $a > 0$) without affecting them. Thus, while (2) is a reasonable assumption, strictly speaking (1) requires making an interpersonal comparison of utility and, as such, going beyond von-Neumann-Morgenstern utility. While this is true, in our case condition (1) can be interpreted as an assumption regarding the principal's view of the marginal benefits received by the two types. Thus, even within the von-Neumann-Morgenstern framework the assumption is appropriate if the principal judges the High type to be the connoisseur (and, therefore, wants to give more funds to him) and the Low type to be the boor (who thus deserves less funds).

3.1 First-best outcome

We begin with the case where the principal knows or can costlessly determine the agent's type. The principal, thus, can freely choose the distribution of funds that maximizes her net welfare. For this case, since all agent types are risk-averse and the principal's benefit is derived from those of the agent, the principal will never wish to introduce variability in the allocations. Because both she and the agent suffer from volatile spending, she will try to avoid it, and straightforward maximization of welfare makes variability zero. The center thus maximizes, for each agent type $i = H, L$, her benefit from this allocation minus her costs:

$$\max_{\{x_i\}} V_i(x_i) - x_i. \quad (3)$$

The solution is simply

$$V'_i(x_i^{FB}) = 1, \quad (4)$$

where the superscript FB indicates that these are first-best allocations. We also assume that $x_L^{FB} > x_{\min}$, so that we have an interior solution for the first-best. Condition (4) implies $V'_H(x_H^{FB}) = V'_L(x_L^{FB})$. Since the marginal benefit of High is greater than that of Low at all levels of funds above Low's first-best, and since benefits have diminishing marginal returns, High must receive a greater allocation than Low. Observation 1 summarizes these results.

Observation 1 (First-best outcome) *In the first-best, the principal allocates a different certainty*

amount to High and Low, and High receives more. Thus, we have

$$x_H^{FB} > x_L^{FB}. \quad (5)$$

3.2 Second-best

When the agent's type is private information, the principal must design a menu of allocations designed to maximize her payoff given that the agent chooses the allocation that maximizes his own expected utility.

3.2.1 The full problem

If the principal were to offer x_H^{FB} and x_L^{FB} to the agent, he would choose x_H^{FB} , since $x_H^{FB} > x_L^{FB}$ and both agent types prefer more funds to less. Thus, as is standard in screening problems, Low would envy High in the first-best allocation, and if the principal asked the agent to reveal his type, he would claim to be High. The best the principal can do is to offer a common menu to the agent types, consisting of a tuple $\{\{x_{ij}\}, \{p_{ij}\}\}_{i=H,L}$ with $x_{ij} \geq x_{\min}$, $J_i \in \mathbb{N}$, $0 \leq p_{ij} \leq 1$, $\sum_{j=1}^{J_i} p_{ij} = 1$. Our analysis makes use of the Revelation Principle and follows the typical approach in the screening literature.⁵ The principal's complete problem is to

$$\max_{\{x_{ij}, p_{ij}\}_{i=H,L}} \lambda \left[\sum_{j \in J_H} p_{Hj} V_H(x_{Hj}) - \sum_{j \in J_H} p_{Hj} x_{Hj} \right] + (1 - \lambda) \left[\sum_{j \in J_L} p_{Lj} V_L(x_{Lj}) - \sum_{j \in J_L} p_{Lj} x_{Lj} \right] \quad (6)$$

s.t.

$$\sum_{j \in J_L} p_{Lj} V_L(x_{Lj}) \geq \sum_{j \in J_H} p_{Hj} V_L(x_{Hj}) \quad (\text{IC-L})$$

$$\sum_{j \in J_H} p_{Hj} V_H(x_{Hj}) \geq \sum_{j \in J_L} p_{Lj} V_H(x_{Lj}). \quad (\text{IC-H})$$

Low's incentive compatibility constraint (IC-L) requires that Low cannot make himself better off by choosing the bundle the principal has designed for High. (IC-H) says the converse for High. Because the principal is giving funds to the agent, we do not consider a participation constraint

⁵As is standard in virtually all adverse selection and screening problems, we allow the principal full commitment capabilities, and permit no post-choice renegotiation.

(except through a possible ex-post requirement that $x_{ij} \geq x_{\min}$).

For the second-best, the principal can actively choose the probabilities of the lottery. In practice, the center may only be able to use natural, exogenously given uncertainty, but the full second-best remains the benchmark for desirable lottery choice.

3.2.2 Simplifying the problem

Three observations simplify solving this problem.

First, only Low's IC will be binding. That is because Low envies High in the first-best, but not vice versa. Suppose, by contradiction, that Low's IC were not binding in the second-best. Then the principal could reduce Low's allocation and give some of it to High. That would produce higher expected utility for the principal.

Second, the principal will give Low a non-random allocation. Suppose, by contradiction, that Low receives a lottery. Denote this lottery by L , and let its expected cost to the principal be K . Let C be Low's certainty equivalent for this lottery. Consider an alternative allocation of Low, in which he receives $aC + (1 - a)L$, where a is small. Low is indifferent between the two allocations. For small a , High's IC is not violated. Since the certainty equivalent, C , is less than the expected cost K , the principal is better off with the alternative allocation. Therefore, it cannot be optimal for the principal to give Low a nondegenerate lottery.

Third, the principal optimally chooses two different prizes for High, not more. This important observation dramatically simplifies the problem, because it implies that we do not have to choose a possibly very complicated distribution for the allocations, but can restrict ourselves to finding two optimal prizes with their associated two optimal probabilities. Specifically, we can show that

Lemma 1 *If there exists an N prize lottery for High that solves the principal's problem, then there also exists a two-prize lottery that solves the problem.*

Proof. See the Appendix. ■

The intuition behind this observation is that if there a solution to the principal's problem that offers Low expected utility k and has more than two prizes for High, we can write the lottery for High as a compound lottery consisting of a number of binary lotteries, each of which offers Low expected utility k . Let V_H^* denote the principal's expected utility under the original (optimal)

lottery if the agent is High. If the principal has expected utility V_H^* from the original lottery offered to High, then he has the same expected utility from the compound lottery as well, since it ultimately offers the same probabilities as the original lottery. The expected utility for the compound lottery can be written as a weighted sum of the expectation of each of the binary lotteries. Because the original lottery was optimal, the expectation of the principal for each of the binary lotteries must be V_H^* . Otherwise, there must be a binary lottery with a higher expectation of the center's utility for High, which contradicts optimality. Thus, there is a binary lottery that does at least as well as the original lottery.

3.2.3 The simplified problem

Making use of Lemma 1 and the fact that the optimal allocation for Low involves a certain (non-random) prize, we can simplify the principal's problem. Let the allocation intended for High consist of prize b with probability p and prize s with probability $1 - p$, where $b > s$.⁶ Let z denote the non-random amount of funds intended for Low. The center's simplified problem is therefore:

$$\max_{p,b,s,z} \lambda [\{pV_H(b) + (1-p)V_H(s)\} - \{pb + (1-p)s\}] + (1-\lambda) [V_L(z) - z] \quad (7)$$

$$\text{s.t. } V_L(z) = pV_L(b) + (1-p)V_L(s). \quad (\text{IC-L})$$

3.2.4 Analysis

To understand the center's optimal resource allocation policy, we begin with the case where Low's benefits fall faster than High's as allocations approach the minimum. Suppose the principal offers only a single, non-random allocation, z , that falls between the first-best allocations for the two types. Next, suppose that, in addition to offering z for sure, the principal also offers a non-degenerate lottery that yields b with probability p and s with probability $1 - p$, having expected value z . Since both types are risk-averse, neither type would choose the lottery over z for sure. But, since Low is more risk-averse than High, Low dislikes the new package even more than High. Thus, the principal can increase the expected value of the uncertain package a little bit to make it attractive for High, and she can also decrease the certain allocation a little bit without Low

⁶Mnemonicly, b denotes the big prize and s denotes the small one.

starting to prefer the variable package. Such an adjustment will benefit the principal whenever the additional cost of the risk imposed on High is not too great. Below, we show that this is often the case.

To achieve separation, the center has to take the risk that High receives less than Low with some probability. In fact, the analysis reveals an important feature of the optimal lottery that the center offers: The bad outcome for High, s^* , is set far below even what Low would have received in the first-best, but the probability of the good outcome, p^* , comes very close to unity. The principal recognizes potential efficiency losses by having states in which High (and the center) has a marginal benefit much higher than marginal cost, but because those states occur rarely, this depresses expected welfare but slightly. Moreover, the principal is able to allocate more funds to High in the good state. The intuition is that by pushing High's funds in the low state towards the lower bound, the principal makes High's package more unattractive for Low, at an increasing rate. By doing so, she obtains freedom to push up High's allocation in the good state where High has a higher marginal benefit than Low, and also to increase the probability that state occurs.

Proposition 1 summarizes this first central result of our analysis:

Proposition 1 (*Second-best approximates first-best*)

If, as allocations approach the minimum, the ratio of Low's marginal benefits to High's is unbounded above (i.e., $\lim_{x \rightarrow x_{\min}} V'_L(x) / V'_H(x) = \infty$), then it is possible to approximate the first-best arbitrarily closely.

Proof. See the Appendix. ■

If Low's marginal benefits approach infinity more quickly than High's as allocations approach the lower bound, the principal can approximate the first-best arbitrarily closely by offering High a lottery consisting of a very bad prize with small probability and something close to High's first-best with a very large probability. Low is given a non-random level of funds very close to his first-best. The condition of Proposition 1 holds, for example, with CRRA benefits, $V_i(x) = (1/(1 - \gamma_i))x^{1-\gamma_i}$, $\gamma_H < \gamma_L$, and $x_{\min} = 0$, in which case $\lim_{x \rightarrow x_{\min}} V'_L(x) / V'_H(x) = \lim_{x \rightarrow 0} x^{-\gamma_L + \gamma_H} = \infty$.

When the ratio of Low's marginal benefit's to High's is bounded as allocations decrease, the first-best cannot be approximated arbitrarily closely, but the optimal lottery retains these basic features. (This is the case, for example, for CARA benefits, $V_i(x) = -(1/r_i) * e^{-r_i x}$, $r_H < r_L$,

and $x_{\min} = 0$, in which case $\lim_{x \rightarrow x_{\min}} V'_L(x)/V'_H(x) = \lim_{x \rightarrow 0} e^{-r_L x + r_H x} = 1$.) Proposition 2 characterizes this solution when the principal finds it worthwhile to screen. (We discuss below the conditions for screening to be preferred to offering the same allocation to all agent types.) In the most general case, what can be shown is that Low receives a fixed amount greater than his first-best and High receives a lottery with prizes above and below Low's certain allocation (results 1 and 2 of the Proposition). Moreover, in result 3 of the Proposition, we provide a sufficient condition for the principal to offer High a lottery consisting of the worst available prize (x_{\min}) with low probability and a prize as close as possible to High's first-best with high probability.

Proposition 2 (Optimal lottery when the ratio of Low's marginal benefits to High's is bounded) *Suppose $\lim_{x \rightarrow x_{\min}} V'_L(x)/V'_H(x) = C < \infty$. The principal can still separate High and Low-productivity types by offering one sufficiently variable package and one package with a non-random allocation. When separating High and Low is worth it, the optimal allocations can be characterized as follows:*

1. $x_H^{FB} > b^* > z^* > \{x_L^{FB}, s^*\}$ where the term $\{\cdot\}$ means that x_L^{FB} may be greater or smaller than s^* .
2. High receives more in expectation than Low: $E_H^* = p^* b^* + (1 - p^*) s^* > z^*$
3. If Low is more downside risk-averse than High, i.e., if $-\frac{V_L'''(x)}{V_L''(x)} > -\frac{V_H'''(x)}{V_H''(x)}$, the principal chooses $s^* = x_{\min} < x_L^{FB}$.

Proof. See the Appendix. ■

Thus, three losses of utility are incurred relative to the first-best: First, Low's (non-random) allocation is greater than his first-best allocation of funds. Second, both of the prizes offered to High are smaller than his first-best allocation of funds (i.e., even in the good state, he receives strictly less than the first-best amount). Third, High may receive less than Low (the probability that the bad state is realized is non-trivial). However, these losses allow the center to screen the agents and earn a higher expected payoff than she would without employing them.

The strong result in part 3 of our Proposition 2 (which makes use of properties of utility functions shown in Chiu (2005) and Chiu (2010)) that the optimal lottery may involve the worst possible outcome for High is particularly striking also because of the condition under which it is obtained: higher downside risk aversion of Low than High. Not only is downside risk aversion (or

prudence) a plausible property.⁷ But for many familiar benefit functions (in particular for CRRA and CARA benefits), we also have that if Low is more risk-averse than High, he is also more downside risk-averse.⁸

Figure 1 plots the optimal allocations for the case of a lower bound of $x_{\min} = 0$ (a rescaling that provides an intuitive reference point without sacrificing generality). For simplicity, the figure assumes that the ratio of Low’s benefit to High’s benefit goes to unity (and it draws the x-axis at that level), but this is obviously not required.




Figure 1: Second-best screening allocation (Notation given in text.)

The solid line in Figure 1, connecting the utility levels $V_L(b^*)$ and $V_L(s^*)$, allows us to track the expected utility for Low when he mimics High. By Jensen’s inequality, the center can give High a larger expected allocation than Low, that is, $E_H^* = p^*b^* + (1 - p^*)s^* > z^*$.

Figure 2 presents the solution in an alternative way. Expected utility maximization models can be represented by a two-moment decision model under some conditions (Sinn 1983, Meyer 1987). Representing the principal’s problem in mean-variance space has the advantage that we can see

⁷Many commonly used utility functions exhibit it (Gollier 2001), it helps explain actually observed risk choices such as those at the horse track (Golec and Tamarkin 1998), and prudence is a necessary condition for absolute risk aversion to be decreasing.

⁸We use a prudence-based definition of more downside risk-aversion. Liu and Meyer (2012) introduce an alternative definition based on decreasing absolute risk aversion (DARA). The definitions are closely connected (see section 4 in their paper).

immediately the presence of a screening condition. In particular, in $\mu-\sigma$ -space (with volatility on the horizontal axis), indifference curves of more risk-averse agents are more steeply sloped. Under our assumption that Low is more risk-averse, those with high productivity require fewer additional funds for any increase in fund variability than those with low productivity. Thus, the indifference curves look as in Figure 2.




Figure 2: Representation of the second-best screening allocation in $\mu-\sigma$ space.

In the first-best, zero volatility is optimal, and High receives a larger allocation than Low, as indicated by the points $H-FB$ and $L-FB$, respectively. The second-best, compared to the first-best, employs a higher certain allocation for Low (and still no volatility) and a lower expected and more risky allocation for High, as indicated by the points $H-SB$ and $L-SB$, respectively. The principal trades off giving a larger expected amount to High against having to do this with greater risk in order to keep Low from mimicking High. The second-best solution optimally balances these two objectives, such that Low is just indifferent (as can be seen by the fact that $L-SB$ and $H-SB$ lie on the same indifference curve for Low).

4 Discussion and interpretation

4.1 The correlation between marginal productivity and risk tolerance

The separation of High and Low depends on the positive correlation between marginal productivity and risk tolerance. (When High is less risk-tolerant, the best that the principal can do is to offer a

fixed budget, i.e., an identical allocation to the two types.) How plausible is the positive correlation assumption?

For the production case, it is straightforward to think of situations in which one agent (a division, for example) is more marginally productive and has a less curved production function, effectively making the agent more risk tolerant than another agent. The same factors that lead marginal product to be larger for High plausibly also lead to marginal product declining more slowly. For example, consider managerial talent. A talented manager is not only better at making workers productive, but he is also better at maintaining those productivity gains for additional workers.

The consumption case is more subtle, but economic intuition suggests that the positive correlation may be a good starting point. Suppose that two agents receive, in fact, the same marginal benefits from an incremental dollar. Then the one who is more risk-tolerant – that is, who has a benefit function that flattens out more slowly – will benefit more from a more-than-incremental windfall. To determine the presence or absence of the correlation between being risk-tolerant and being a connoisseur, we conducted surveys among Masters and Ph.D. students in Public Policy at Harvard University (the "Harvard sample"), two separate groups of undergraduate students of economics and psychology at the University of Zurich ("undergraduate sample 1" and "undergraduate sample 2"), and undergraduate students of finance at the University of Zurich (the "finance sample"). In total, 339 subjects participated. We pool data from all samples. In the regression analysis we use dummies to identify each sample.

The questionnaire consisted of two parts. (See the Appendix for details.) First, we employed a widely used format of asking for the degree of risk tolerance, entailing pairwise comparisons of a risky choice and a certain outcome. In essence, the series of questions we asked culminated in an answer to the question: Consider a fair lottery where you have a 50% chance of doubling your income, and a 50% chance of losing a certain percentage, say $x\%$ of your income. What is the highest loss x that you would be willing to incur to agree to taking part in this lottery? Similar questions are proposed by Gollier (2001), Mankiw and Zeldes (1991), and others.⁹ Our average

⁹A large literature deals with methods of eliciting risk preferences. Anderson and Mellor (2009) establish that risk preference estimates can vary greatly across elicitation methods. We opted for a simple approach because given we wanted to avoid a dependence of the results on the perhaps heterogeneous numerical skills of survey participants in our overall sample (Dave, Eckel, Johnson, and Rojas 2010). Results in Hey, Morone, and Schmidt (2009) suggest that pairwise choice methods of the type underlying our survey tend to be less noisy than alternative methods. Of course, a drawback of our approach is that given the magnitude of stakes of interest, we had to ask hypothetical questions.

finding of $x\% = 26\%$ is consistent with the finding of Barsky, Juster, Kimball, and Shapiro (1997) who report an average of $x=23\%$.

Second, we confronted participants with two (imaginary) situations. In situation 1, participants were asked to imagine that they had found a Swiss Franc 50 note (US\$50 in the Harvard sample), while in situation 2, participants were asked to imagine that they had just won in a lottery an amount of Swiss Franc 500,000 (US\$500,000 in the Harvard sample). In both situations, we asked participants to compare themselves to a peer of similar wealth and income and answer whether they would derive much greater, greater, equal, less or much less welfare from these positive events. While the average of participants answered that they would benefit about the same as their peers, there was substantial variation in the answers. Figure 3 plots the mean risk tolerance levels in the various categories and the corresponding standard errors.




Figure 3: Risk tolerance and stated marginal benefits. The panel on the left plots mean risk tolerance levels and standard errors in the five marginal benefit categories for the US\$/Swiss Franc 50 question (N=275). The panel on the right does the same for the US\$/Swiss Franc 500,000 question (N=232). Total N = 339 (some participants answered both questions). For details on the survey see the text and the Appendix.

Figure 3 illustrates the central finding of this analysis: There is a strong positive correlation between risk tolerance and a subjective feeling of being able to use windfall funds in a more effective, utility-enhancing way than one's peers.¹⁰

¹⁰We coded the answers into the five categories and created variables MU50 and MU500k, where MU denotes

We interpret this as being consistent with the assumption that High marginal utility types (connoisseurs) are more risk-tolerant. Thus, the findings are supportive of Solomon’s working hypothesis.¹¹ Clearly, this approach is limited and merely exploratory, because it relies on introspection and subjective assessments, employs solely student subjects, and does not involve actual monetary stakes.

Also, there may be unobserved heterogeneity in the population. For example, if survey participants differ substantially on wealth, then answers to our survey question on marginal utility would miss this unobserved variable. In general, we would expect poorer individuals to have higher marginal utility of wealth and to be less risk-tolerant. (Evidence in Dohmen, Falk, Huffman, and Sunde (2010) suggests, for example, that credit-constrained individuals exhibit a lower willingness to take risks.) Therefore, unobserved unequal wealth levels are likely to work against the empirical assumption of the model. Thus, our findings of a positive correlation despite the differential wealth factor suggest that the positive relationship holding wealth constant must be strong. Moreover, in many practical circumstances, we believe, the principal will have firm prior information about the wealth level of her agents, which she can take into account in her allocation decisions.

Differing decision-making abilities may produce a second realm of unobserved heterogeneity. This phenomenon will often work in favor of the assumption: More able agents find more value-generating uses of funds, and they may also avoid the excessive risk aversion well documented for small gambles (Rabin 2000).

We recognize, of course, that principals may not merely be interested in allocating resources where they offer the highest marginal utility, hence achieving the highest total utility. They may feel, for example, that less able agents should receive more, even if they will not put resources to

marginal utility. Higher numbers denote higher stated marginal benefits. Because our marginal benefits variables are ordinal, we used rank-order correlations to statistically determine its relationship with risk tolerance. Both the Spearman correlation coefficient and Kendall’s tau are positive and indicate a highly statistically significant relationship, with p-values of 0.05 for the US\$/Swiss Franc 50 question and below 0.01 for the US\$/Swiss Franc 500,000 question. We also ran regressions with risk tolerance as the dependent variable and marginal benefits as the key explanatory variable (and sample dummies as controls), obtaining the same findings. The detailed results are available on request.

¹¹The data from Solomon’s surveys, alas, are lost in history. Our surveys thus add modestly to the knowledge about the joint distribution of risk attitudes and preference intensity in the population. Other researchers, while not yet familiar with Solomon’s Lost Scroll, have also found evidence that is consistent with his observation. In particular, Dohmen, Falk, Huffman, and Sunde (2010) present evidence that individuals with higher cognitive abilities are less risk-averse. Donkers, Melenberg, and van Soest (2001) find that, controlling for income, education is positively associated with risk tolerance Grinblatt, Keloharju, and Linnainmaa (2011) show that individuals with higher IQ participate in the stock market to a greater extent. If the principal wishes to give more funds to an agent with higher cognitive abilities (who may be able to use these funds more productively for his and the principal’s benefit), this supports the effectiveness of Solomon’s screening method.

the highest value use. If promoting total utility is not the goal, then our analysis is silent.

4.2 When is separation not profitable?

Even when screening is possible, it does not always pay. To see the intuition, consider the alternative to screening, namely allocating the optimal identical amount to both types. The principal would solve

$$\max_y \lambda [V_H(y)] + (1 - \lambda) [V_L(y)] - y, \quad (8)$$

leading to a y^* that satisfies

$$\lambda V'_H(y^*) + (1 - \lambda) V'_L(y^*) = 1. \quad (9)$$

Note that as the fraction of one of the two types increases, the fixed budget allocation y^* resembles more the first-best of that type and less the first-best of the other type. In the limiting case, of course, $\lim_{\lambda \rightarrow 1} y^* = x_H^{FB}$ and $\lim_{\lambda \rightarrow 0} y^* = x_L^{FB}$. There is, however, a crucial difference between the types: Fixed budgets always give too much to Low, and too little to High. In addition, this distortion is costlier for High, and hence to the principal as well, because he has greater marginal benefits.

Of course, no allocation can outperform the first-best. We also know that the second-best approximates the first-best arbitrarily closely under the conditions we found above. Thus, the principal is always able to improve just a little bit on any proposed solution in the second-best. As the fraction of High types (λ) increases, fixed budgets become more efficient. To retain the superiority of screening over fixed budgets, as λ increases, the principal makes the lottery ever more extreme, pushing s^* further and further down towards x_{\min} . Consequently, it always pays for her to screen in this case. This result is stated in part (i) of the Proposition below.

When the principal is constrained in her ability to achieve the first-best, things turn out to be different and a fixed budget for both High and Low can become a serious competitor to second-best screening. When the fraction of High types increases, the principal may be limited in her abilities to adjust the allocations. As shown, in the typical case, she is already at the lower limit with the prize in the bad state. Then, the only way to avoid Low envying High is to increase the allocation to Low or the allocation to High in the good state (suitably adjusting the probabilities in High's lottery so that the associated risk is sufficiently unattractive for Low). It is, however, never

optimal to increase High's allocation beyond the first-best. Doing so would buy the principal a second-order gain at a first-order cost, because at that level of funds, the marginal cost is greater than the marginal benefit. Therefore, a single fixed budget may become optimal.

Conversely, as the fraction of Low types increases (λ decreases), the fixed budget strategy becomes less efficient, because, after all, High's marginal productivity is higher. Of course, with many Low and few High types, the screening solution also becomes less attractive. The Appendix shows that the first effect dominates the second. Intuitively, the loss in efficiency from screening comes from distorting High, but if the fraction of High types is small, it costs virtually nothing to distort them. In sum, if screening pays for some fraction of Low types, it also pays for all greater fractions.

A final important determinant of the decision whether to screen is whether the types are sufficiently different, compared to the costs of funds. Intuitively: Even if the fraction of Lows is high, the center may do well enough with fixed budgets if Lows and Highs are not that different and if costs are sufficiently high, implying that even in the first-best the principal would give them almost equal allocations. These insights are summarized in the following

Proposition 3 *Consider a fixed budget as an alternative to screening.*

(i) *If $\lim_{x \rightarrow x_{\min}} V'_L(x) / V'_H(x) = \infty$, screening is always profitable.*

(ii) *Suppose $\lim_{x \rightarrow x_{\min}} V'_L(x) / V'_H(x) = C < \infty$.*

a) *If Low is sufficiently risk-averse, then it pays to screen.*

b) *If screening is profitable at some fraction of Low types, it remains profitable at all greater fractions of Low types.*

Proof. See the Appendix. ■

4.3 Two time periods

While we have framed the screening problem as one of choosing the optimal pair of a non-random budget and a random allocation, it is also useful to consider the case of two periods when no learning about productivity takes place between periods. Now we assume that a bundle the principal can offer consists of a certain allocation in each of the periods, and that agents cannot shift resources between periods. Suppose utility is additive across periods.

Begin with the case where there are two time periods of equal length, and discounting is set aside for simplicity. Let x_i be the first period allocation for type i , and let y_i be the second period allocation for that type. In the first-best, $x_H^{FB} = y_H^{FB} > x_L^{FB} = y_L^{FB}$. As before, we know immediately that there is no point in giving different allocations to Low (over time). Only High's allocation needs to be distorted in order to make it unattractive for Low. The principal's problem in this case is to

$$\max \lambda [V_H(x_H) + V_H(y_H) - (x_H + y_H)] + (1 - \lambda) [V_L(x_L) + V_L(y_L) - (x_L + y_L)] \quad (10)$$

$$\text{s.t. } 2V_L(x_L) \geq V_L(x_H) + V_L(y_H) \quad (\text{IC-L})$$

One solution that is consistent with Low's IC would be $x_H = y_H$. But since not both x_H and y_H can be greater than $x_L = y_L$, the only solution would be for all spending amounts to be equal, as in the fixed-budgets case. But it is easy to see that this cannot be optimal for the principal. Thus, we can check the possibility that $x_H > x_L = y_L > y_H$. (The case of $x_H < x_L = y_L < y_H$ is symmetric.) Doing similar calculations as before, we arrive at the following

Proposition 4 *The principal can separate high and low-productivity types by offering one sufficiently variable package and one package with constant funds over time. In particular:*

1. $x_H^{FB} > x_H^* > x_L^* = y_L^* > \{x_L^{FB}, y_H^*\}$ where the last term means that x_L^{FB} may be greater or smaller than y_H^* .
2. $x_H^* + y_H^* > x_L^* + y_L^*$.

Result 2 can also be written as $\frac{x_H^* + y_H^*}{2} > x_L^*$, indicating that the expected per-period allocation to High is bigger than the expected per-period allocation to Low.

Note that with this scheme the principal will not be able to achieve close to first-best target efficiency. In particular, this case operates as if the principal were restricted to using lotteries with equal probability on both prizes. But if time becomes divisible (preferably infinitely so), then the principal can shorten the time period spent consuming the low prize (i.e., reduce the weight on y_H), and lengthen the time period for the high prize. If the principal can choose both the time periods and the prizes, the problem becomes equivalent to the main case.

4.4 Multiple types

The analysis extends to the case where there are N agent types. Normalize types such that $V_1'(x) < V_2'(x) < \dots < V_N'(x)$ for all x . In the first-best, $x_N > x_{N-1} > \dots > x_1$.

Characterizing the second-best in the case where marginal utility of the lowest type goes to infinity faster than that of the other types as funds approach the lower bound proceeds, in principle, along the same lines as the two-type analysis. Consider, for example, three types. Only the local upward incentive compatibility constraints will bind. The lowest type receives a non-random prize. The principal needs to make sure that he does not envy the Medium type. Thus, she gives the Medium type a lottery with a high probability on Medium's first-best, and a low probability on a very low prize. Also, the center needs to assure that Medium does not envy High. Thus, High will receive a very high probability on his first-best and a low probability on an even worse prize than the bad prize for Medium. Note that the proof of the equivalence of an N -prize lottery to a two-prize lottery applies here as well.

Depending on the fractions of Lows and Mediums in the population, and depending on how tight the bounds are on the low prizes for the productive types, the principal may do far better with screening, or may do no better than with identical budgets. When Low's and Medium's ratios of marginal benefits, respectively, to the next higher type's marginal benefits are bounded, screening may still be desirable. If so, the principal achieves welfare in between the first-best and the fixed budgets allocation.

4.5 Numerical results

We can use numerical simulations to verify the analytical results and to quantify the welfare gains. The Appendix shows the results of such simulations for CRRA and CARA benefit functions.

5 Conclusion and applications

In many important contexts, a principal allocates resources to agents who then employ those resources to create benefits for both themselves and the principal. A challenge arises when, as is frequent, the principal does not know the agents' abilities to create such benefits. Such problems arise in fairy tales, within the family, the corporation, and with nonprofit and government programs.

In fairy tales, the benevolent king wishes to find out which of the suitors loves his daughter the most. Perhaps the most common everyday example is the parent distributing resources to a child, not knowing how much benefit the child will receive from them. In the business world, a corporate center allocates resources to divisions, not knowing what level of production or profitability will be reaped from those resources. In academia, deans allocate slots and funds to different departments, but are likely to be only partially informed on what benefits they will bring to the school. In the world of nonprofits, the philanthropist provides funding to a variety of endeavors, but does not know how effective each will be in promoting her causes. The common characteristic of all these cases is that the principal's benefit is strongly related, and possibly directly proportional, to those of her agents. Thus, the agent is the source of the principal's benefits. Nevertheless, a divergence of interests arises, because an agent is solely interested in his own productivity, and rarely pays fully for his own budget. The principal is concerned with the productivity of all agents.

Our central result is that a principal can successfully separate differentially productive types by using the insight that the degree to which an agent produces marginal benefits from funds can be (and plausibly frequently is) positively correlated with the source's effective tolerance for risk. If this correlation between risk tolerance and productivity holds, the optimal allocation of resources requires giving fixed funds to the less productive type, but a surprisingly extreme lottery to the productive type. The High producer demonstrates his capabilities by taking excess risk, such as accepting a very small but positive probability of receiving a small allocation of resources. When benefits to Low fall without bound and faster than benefits to High as resources decrease (as is the case for constant relative risk aversion benefit functions, which are frequently used in economic modeling), the center is able to approximate first-best expected welfare very closely, even if negative payments are not permitted. When negative payments are allowed, e.g., if individuals have endowments that the principal can cut, the principal can always get almost first-best welfare for herself and for the source. Although the extreme lottery required for this outcome may be hard to implement politically, it provides an important benchmark result for applications.

In practice, centers can use the external world as a randomization device, even without necessarily aiming to optimize the probabilities associated with the prizes. Indeed, the model implies that a center may not wish to minimize uncertainty because of the opportunities it generates for effective screening. The extreme lottery result of Proposition 2 even suggests that in many circumstances

it is beneficial to the center to have at its disposal extreme downside risks that occur with very low probability. As such, the model helps understand some arrangements in the real world, and it offers recommendations for an improved use of resources. For example, many individuals pursuing job alternatives have the choice between a risky and a safe job. The model explains why a firm (say a university) might want to have its job risky: This property offers a mechanism for assuring that the person would benefit greatly from the position. Thus, the candidate will not leave quickly for another position. As another example, family companies often encounter the paper's separation problem and sometimes use versions of the proposed solution when implementing succession plans, aiming to give the company to the most deserving heir.¹² Finally, as an example of the concrete normative implications of the model, a venture capital fund, in trying to screen out unproductive investments, may employ the method proposed here. If it thinks there is a correlation between productivity and risk tolerance, say because entrepreneurs with a better idea will have a better fallback, then it could offer entrepreneurs a risky package. Thus, the venture capital investors would not insure entrepreneurs if some extreme untoward event happened, even though efficient risk spreading would recommend they should. Note that in principle this has nothing to do with motivating effort, just assessing productivity. Of course, in real world contexts effort choice will play an additional role.

Thus, in a variety of settings, gauging the intensity of preferences based on risk taking can provide a powerful method to increase target efficiency. This screening tool will be effective whenever there is a positive correlation between marginal benefits and tolerance for risk. We argued that this is a plausible condition when central resources must be allocated to corporate divisions. Our empirical samples for the consumption case, drawn from Switzerland and the US, found our required condition satisfied: Risk tolerance and marginal benefits from resources were positively correlated. Given this straightforward condition, whether dealing with for-profit entities or individuals as beneficiaries, resource allocation can be implemented wisely, using the system of Solomon.

¹²Consider the following real case. The founder and CEO of Spedag group, a leading Swiss logistics company, had three children, A, B, and C. He wanted the child who felt he could get the most value to lead the company. He did not want to separate ownership and control. (Skill may have been an additional consideration.) He separated the original company into an operating company (OP), and a holding company (H). H initially owned everything. Son A received a personal loan from H enabling him to buy OP from it. OP also rented buildings from H. Daughter B and son C received part of the price that A paid as an advance on their inheritance, but were excluded from the business. Son A, as sole owner and CEO of OP bears substantial risk, but in expectation felt his holdings were worth significantly more than the advances paid to B and C. B and C received enough that they were happy to accept this arrangement.

References

- AKERLOF, G. (1970): “The Market for ‘Lemons,’ Qualitative Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84(3), 488–500.
- ANDERSON, L. R., AND J. M. MELLOR (2009): “Are risk preferences stable? Comparing an experimental measure with a validated survey-based measure,” *Journal of Risk and Uncertainty*, 39(2), 137–160.
- ANTLE, R., AND G. D. EPPEN (1985): “Capital Rationing and Organizational Slack in Capital Budgeting,” *Management Science*, 31(2), 163–174.
- ARNOTT, R., AND J. E. STIGLITZ (1988): “Randomization with asymmetric information,” *RAND Journal of Economics*, 19(3), 344–362.
- BARSKY, R. B., T. F. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): “Preference parameters and individual heterogeneity: An experimental approach in the health and retirement study,” *Quarterly Journal of Economics*, 112(2), 537–579.
- BERNARDO, A. E., H. CAI, AND J. LUO (2001): “Capital budgeting and compensation with asymmetric information and moral hazard,” *Journal of Financial Economics*, 61(3), 311–344.
- (2004): “Capital Budgeting in Multidivision Firms: Information, Agency, and Incentives,” *Review of Financial Studies*, 17(3), 739–767.
- BRAMS, S. J. (2005): “Fair division,” in *Oxford Handbook of Political Economy*, ed. by B. R. Weingast, and D. Wittman. Oxford University Press.
- BRAMS, S. J., AND D. M. KILGOUR (2001): “Competitive fair division,” *Journal of Political Economy*, 109(2), 418–443.
- BRAMS, S. J., AND A. D. TAYLOR (1999): *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. W. W. Norton, New York.
- BRITO, D. L., J. H. HAMILTON, S. M. SLUTZKY, AND J. E. STIGLITZ (1995): “Randomization in optimal income tax schedules,” *Journal of Public Economics*, 56(2), 189–223.
- CHIU, W. H. (2005): “Skewness Preference, Risk Aversion, and the Precedence Relations on Stochastic Changes,” *Management Science*, 51(12), 1816–1828.
- (2010): “Skewness Preference, Risk Taking and Expected Utility Maximization,” *Geneva Risk and Insurance Review*, 35(2), 108–129.
- DAVE, C., C. C. ECKEL, C. A. JOHNSON, AND C. ROJAS (2010): “Eliciting risk preferences: When is simple better?,” *Journal of Risk and Uncertainty*, 41(3), 219–243.

- DOHMEN, T., A. FALK, D. HUFFMAN, AND U. SUNDE (2010): “Are risk aversion and impatience related to cognitive ability?,” *American Economic Review*, 100(3), 1238–1260.
- DONKERS, B., B. MELENBERG, AND A. VAN SOEST (2001): “Estimating Risk Attitudes using Lotteries: A Large Sample Approach,” *Journal of Risk and Uncertainty*, 22(2), 165–195.
- GOLEC, J., AND M. TAMARKIN (1998): “Bettors love skewness, not risk, at the horse track,” *Journal of Political Economy*, 106(1), 205–225.
- GOLLIER, C. (2001): *The Economics of Risk and Time*. The MIT Press, Cambridge, MA.
- GRINBLATT, M., M. KELOHARJU, AND J. LINNAINMAA (2011): “IQ and stock market participation,” *Journal of Finance*, 66(6), 2124–2164.
- HARRIS, M., C. H. KRIEBEL, AND A. RAVIV (1982): “Asymmetric Information, Incentives and Intrafirm Resource Allocation,” *Management Science*, 28(6), 604–620.
- HARRIS, M., AND A. RAVIV (1996): “The Capital Budgeting Process: Incentives and Information,” *Journal of Finance*, 52(1), 111–133.
- HARRIS, M., AND R. M. TOWNSEND (1981): “Resource Allocation Under Asymmetric Information,” *Econometrica*, 49(1), 33–64.
- HELLWIG, M. F. (2007): “The undesirability of randomized income taxation under decreasing risk aversion,” *Journal of Public Economics*, 91(3-4), 791–816.
- HEY, J. D., A. MORONE, AND U. SCHMIDT (2009): “Noise and bias in eliciting preferences,” *Journal of Risk and Uncertainty*, 39(3), 213–235.
- LELAND, H. E., AND D. H. PYLE (1977): “Informational asymmetries, financial structure, and financial intermediation,” *Journal of Finance*, 32(2), 371–387.
- LIU, L., AND J. MEYER (2012): “Decreasing absolute risk aversion, prudence and increased downside risk aversion,” *Journal of Risk and Uncertainty*, 44(3), 243–260.
- MANKIW, N. G., AND S. P. ZELDES (1991): “The Consumption of Stock Holders and Non Stock Holders,” *Journal of Financial Economics*, 29(1), 97–112.
- MASKIN, E., AND J. RILEY (1989): “Optimal multiunit auctions,” in *The economics of missing markets and information*, ed. by F. Hahn. Oxford University Press, Oxford.
- MEYER, D. J., AND J. MEYER (2005): “Relative Risk Aversion: What Do We Know?,” *Journal of Risk and Uncertainty*, 31(3), 243–262.

- MEYER, J. (1987): “Two-moment decision models and expected utility maximization,” *American Economic Review*, 77(3), 421–430.
- MILLER, N., N. E. PIANKOV, AND R. J. ZECKHAUSER (2006): “Possibly-final offers,” *Journal of Economics and Management Strategy*, 15(3), 789–819.
- MOSELLE, B., F. DEGEORGE, AND R. J. ZECKHAUSER (2005): “Conspicuous conservatism in risk choice,” *Journal of Risk and Uncertainty*, 35(1), 1–16.
- PRATT, J. W. (2005): “Fair (and not so fair) division,” *Journal of Risk and Uncertainty*, 35(3), 203–236.
- PRATT, J. W., AND R. J. ZECKHAUSER (1990): “The Fair and Efficient Division of the Winsor Family Silver,” *Management Science*, 36(11), 1293–1301.
- RABIN, M. (2000): “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica*, 68(5), 1281–1292.
- SALANIÉ, B. (2005): *The Economics of Contracts: A primer (2nd edition)*. MIT Press, Cambridge, MA.
- SINN, H.-W. (1983): *Economic Decisions under Uncertainty*. North-Holland, Amsterdam et al.
- STEIN, J. C. (1997): “Internal Capital Markets and the Competition for Corporate Resources,” *Journal of Finance*, 52(1), 111–133.
- STEINHAUS, H. (1948): “The Problem of Fair Division,” *Econometrica*, 16(1), 101–104.
- STIGLITZ, J. E. (1982): “Self-selection and Pareto efficient taxation,” *Journal of Public Economics*, 17(2), 213–240.
- WAKKER, P. P., R. H. THALER, AND A. TVERSKY (1997): “Probabilistic Insurance,” *Journal of Risk and Uncertainty*, 15(1), 7–28.
- WEISS, L. W. (1976): “The desirability of cheating incentives and randomness in the optimal income tax,” *Journal of Political Economy*, 89(6), 1343–1352.

6 Proofs

Lemma 1 *If there exists an N prize lottery for High that solves the principal's problem, then there also exists a two-prize lottery that solves the problem.*

We will, more generally, prove the following result:

Lemma A-1. *Let f and g be functions. Let g be monotonically increasing. Let $L = (p_1, \dots, p_N; x_1, \dots, x_N)$ be an N prize lottery with $x_{n+1} > x_n$ such that L solves:*

$$\max E_L f(x)$$

$$\text{s.t. } E_L g(x) \geq k.$$

Then there exists a two-prize lottery that also solves the above maximization problem.

Note that in Lemma A-1, we do not make any assumptions about f . In the model, $f(x) = V_L(x) - Kx$, i.e., the principal's objective function. The parameter k will come from the certain prize offered to Low.

Proof. We will construct a set of binary lotteries B_1, \dots, B_N such that $L = \sum_{i=1}^N \pi_i B_i$ and $E_{B_i} g(x) = k$.

Step 1: Choose the largest and smallest prizes, x_N and x_1 . Note that $g(x_N) \geq k \geq g(x_1)$. Let α_1 be the probability on x_N such that a binary lottery on x_1 and x_N has g -expectation k :

$$\alpha_1 g(x_N) + (1 - \alpha_1) g(x_1) = k.$$

Let $q_1 = \min \left\{ \frac{p_1}{\alpha_1}, \frac{p_N}{1 - \alpha_1} \right\}$. Then $L = (1 - q_1) L_1 + q_1 B_1$, where B_1 is the binary lottery that puts weight α_1 on x_N and $(1 - \alpha_1)$ on x_1 , and $L_1 = \frac{1}{1 - q_1} (p_1 - q_1(1 - \alpha_1), p_2, \dots, p_{N-1}, p_N - q_1 \alpha_1)$. Note that by construction, L_1 puts weight zero on either x_1 or x_N .

Step 2: Iterate Step 1 applied to L_1 using the largest and smallest prizes to which L_1 assigns positive weight. This proceeds after at most N steps. (Note, if there is a prize x_j on L such that $g(x_j) = k$, then one of the binary lotteries may be degenerate.)

The result is a set of binary lotteries B_1, \dots, B_N such that $L = \sum_{i=1}^N \pi_i B_i$, where each B_i has g -expectation k (i.e., $\alpha_i g(x_{i1}) + (1 - \alpha_i) g(x_{i2}) = k$). By assumption, L maximizes the expectation of f subject to $E_L g(x) \geq k$. But $E_L f(x) = \sum_{i=1}^N \pi_i E_{B_i} f(x)$. Since the f -expectation of L is a

convex combination of the expectations of the B_i 's, this implies that the expectation of each of the B_i 's must equal the expectation of f :

$$E_{B_i} f(x) = E_L f(x).$$

If not, then there must be a j such that $E_{B_j} f(x) < E_L f(x)$. But, in that case there must be i such that $E_{B_i} f(x) > E_L f(x)$. This contradicts the assumption that L was optimal. Thus, if there is a lottery L (with finite prizes) that solves the above problem, there is a two-prize lottery that also solves the problem. ■

Proposition 1 (Second-best approximates first-best)

If, as allocations approach the minimum, the ratio of Low's marginal benefits to High's is unbounded above (i.e., $\lim_{x \rightarrow x_{\min}} V_L'(x) / V_H'(x) = \infty$), then it is possible to approximate the first-best arbitrarily closely.

Proof. We will use the utility offered to Low as the control variable. Let $v_L^{FL} = V_L(x_L^{FB})$ be Low's utility when he chooses his first-best allocation. Similarly, define $v_L^{FH} = V_L(x_H^{FB})$ as Low's utility when he chooses High's first-best allocation.

We will also use the fact that since High is less risk-averse than Low, there exists a convex function $g(\cdot)$ such that $V_H(x) = g(V_L(x))$. Note that, by definition, $g(v_L^{FH}) = V_H(x_H^{FB})$.

Consider a lottery of the following form. With probability p , the lottery offers utility v_L^{FH} (and, thus, High's first-best allocation) to Low. With probability $(1 - p)$, the lottery offers utility m . We will let p go to unity, thus considering a lottery that almost surely allocates High's first-best funds to whoever chooses the lottery. At the same time, we will let m become very small, thus considering a prize close to the lower bound x_{\min} .

Let the lottery keep Low indifferent between v_L^{FL} for sure and the lottery:

$$v_L^{FL} = p v_L^{FH} + (1 - p) m.$$

By convexity, High strictly prefers the lottery to Low's certain allocation:

$$g(v_L^{FL}) < p g(v_L^{FH}) + (1 - p) g(m).$$

Consider a sequence of lotteries $(p_\varepsilon, m_\varepsilon) = \left(\varepsilon, \frac{v_L^{FL} - \varepsilon v_L^{FH}}{(1-\varepsilon)}\right)$. By construction, these lotteries satisfy $v_L^{FL} = p v_L^{FH} + (1-p)m$ and offer Low his first-best utility. Note that $m_\varepsilon \rightarrow -\infty$ as $\varepsilon \rightarrow 1$.

To prove the result, we must show that in the limit as $\varepsilon \rightarrow 1$, these lotteries also offer High his first-best utility. This is indeed true:

$$\begin{aligned} p_\varepsilon g(v_L^{FH}) + (1-p_\varepsilon) g(m_\varepsilon) &= \\ p_\varepsilon g(v_L^{FH}) + (1-p_\varepsilon) m_\varepsilon (g(m_\varepsilon)/m_\varepsilon) &= \\ \lim(p_\varepsilon g(v_L^{FH})) + \lim((1-p_\varepsilon) m_\varepsilon) \lim((g(m_\varepsilon)/m_\varepsilon)) &= \\ g(v_L^{FH}) + (v_L^{FL} - v_L^{FH}) * 0 &= g(v_L^{FH}), \end{aligned}$$

where the last line follows from the assumption that $\lim_{x \rightarrow x_{\min}} V_L'(x)/V_H'(x) = \infty$ and the fact that, according to L'Hôpital's rule, $\lim_{x \rightarrow x_{\min}} V_L(x)/V_H(x) = \lim_{x \rightarrow x_{\min}} V_L'(x)/V_H'(x)$. Since Low earns certain utility v_L^{FL} and High's expected utility approaches $V_H(x_H^{FB})$, the principal's welfare approaches her first-best. ■

Proposition 2 (Optimal lottery when the ratio of Low's marginal benefits to High's is bounded) *Suppose $\lim_{x \rightarrow x_{\min}} V_L'(x)/V_H'(x) = C < \infty$. The principal can still separate High and Low-productivity types by offering one sufficiently variable package and one package with a non-random allocation. When separating High and Low is worth it, the optimal allocations can be characterized as follows:*

1. $x_H^{FB} > b^* > z^* > \{x_L^{FB}, s^*\}$ where the term $\{\cdot\}$ means that x_L^{FB} may be greater or smaller than s^* .
2. High receives more in expectation than Low: $E_H^* = p^* b^* + (1-p^*) s^* > z^*$
3. If Low is more downside risk-averse than High, i.e., if $-\frac{V_L'''(x)}{V_L''(x)} > -\frac{V_H'''(x)}{V_H''(x)}$, the principal chooses $s^* = x_{\min} < x_L^{FB}$.

Proof. The principal's problem is to

$$\max_{p,b,s,z} \lambda [\{pV_H(b) + (1-p)V_H(s)\} - \{pb + (1-p)s\}] + (1-\lambda) [V_L(z) - z] \quad (11)$$

$$\text{s.t. } V_L(z) = pV_L(b) + (1-p)V_L(s). \quad (\text{IC-L})$$

Denote with ρ the Lagrange multiplier on (IC-L).

1. The proof that $x_H^{FB} > b^* > z^* > \{x_L^{FB}, s^*\}$, where the last term means that x_L^{FB} may be greater or smaller than s^* , will proceed in three steps.

Step 1: $x_H^{FB} > b^* > s^*$: Taking the derivatives of the Lagrangian with respect to b and s , we obtain (assuming for the time being that $p > 0$)

$$\lambda (V'_H(b^*) - 1) = \rho V'_L(b^*)$$

$$\lambda (V'_H(s^*) - 1) = \rho V'_L(s^*).$$

Since (IC-L) is binding (see the main text), $\rho > 0$. Thus, the right-hand side of both first-order conditions is positive. Recall that the first-best requires $V'_H(x_H^{FB}) - K = 0$. But this implies, by concavity of V_H , that in the second-best, we must have $b^*, s^* < x_H^{FB}$.

Now, if we had $b^* = s^*$, (IC-L) would collapse into $V_L(z^*) = V_L(b^*)$, and, thus, $z^* = b^*$. Thus, the principal would implement fixed budgets, a contradiction to the Proposition's premise. For the condition for screening to be worth it, see Corollary 1.

If screening pays, we must have $b^* \neq s^*$, which implies $p^* \neq 0$ and $p^* \neq 1$ (justifying the assumption we made at the beginning of this step).

Step 2: $z^* > x_L^{FB}$: See the text for the argument that the center will not introduce any risk into Low's allocation. Taking the derivative of the Lagrangian with respect to x_L , we obtain

$$(1 - \lambda) (V'_L(z^*) - 1) = -\rho V'_L(z^*).$$

Since $\rho > 0$, the right-hand side is negative, and since $V'_L(x_L^{FB}) - 1 = 0$, we have, by concavity of V_L , $z^* > x_L^{FB}$.

Step 3: $b^* > z^* > s^*$: Clearly this must be the case, because otherwise High (Low) would receive more than Low (High) for sure and the according IC would be violated.

2. From (IC-L), we have $V_L(z^*) = EV_L(b^*)$. Jensen's inequality implies $EV_L(b^*) < V_L(E_H^*)$. Thus, $z^* < E_H^*$.

3. Assume now that we also have $-\frac{V''_L(x)}{V'_L(x)} > -\frac{V''_H(x)}{V'_H(x)}$. The proof that $s^* = x_{\min}$ will proceed in 6 steps. (Per our assumptions, $x_L^{FB} > x_{\min}$, which proves the inequality.)

Step 1: Consider lottery $L = (b, p; s, 1 - p)$, a candidate optimum with $s > x_{\min}$. Consider an alternate lottery $L' = (y, q; x, 1 - q)$, where $b > y > s > x$ and $q > p$ such that $pb + (1 - p)s = qy + (1 - q)x$ and $pU_H(b) + (1 - p)U_H(s) = qU_H(y) + (1 - q)U_L(x)$. That is, L and L' have the same expected value and same expected utility for High. It is easy to show that such L and L' exists.

Step 2: According to Chiu (2010), Bernoulli lotteries are always (generalized) skewness comparable. Moreover, distribution F is more skewed to the right than distribution G if and only if F puts greater probability on the small prize than G does. Hence, L is more skewed to the right than L' . Another way to say this is that L' has more downside risk than L .

Step 3: Theorem 2a in Chiu (2010) establishes that preferences over skewness comparable lotteries can be represented by a utility function that depends only on the mean (μ), variance (σ^2) and centralized third moment (m^3) of the distribution. That is, there exists a function $U(\mu, \sigma^2, m^3)$ that represents preferences, where $U(\mu, \sigma^2, m^3) = \int u(x) dF(x)$. Furthermore, $U(\mu, \sigma^2, m^3)$ is increasing in μ and m^3 and decreasing in σ^2 if and only if $u' > 0$, $u'' < 0$, and $u''' > 0$.

Step 4: Since High is indifferent between L and L' , and $m_L^3 > m_{L'}^3$, it must be that $\sigma_L^2 > \sigma_{L'}^2$. Otherwise, L would have less risk and less downside risk, and High would necessarily prefer L to L' .

Step 5: Corollary 2 (Chiu 2005): Suppose $u'' < 0$, $v'' < 0$, and $u''' > 0$, and F and G are strongly skewness comparable, $\mu_F = \mu_G$, $\sigma_F > \sigma_G$, and $m_F > m_G$. If $\frac{-v'''(x)}{v''(x)} \geq \frac{-u'''(x)}{u''(x)}$ and $\int u(y) dG(y) = \int u(y) dF(y)$, then $\int v(y) dG(y) \leq \int v(y) dF(y)$.

Applying this in our context, it says that the assumption the Low is more downside risk-averse than High, $-\frac{V_L'''(x)}{V_L''(x)} > -\frac{V_H'''(x)}{V_H''(x)}$, implies that Low prefers L to L' .

Step 6: Whenever $s > x_{\min}$ it is possible to construct lottery L' that keeps High as well off as has the same expected value as L . Since Low prefers L to L' , it is incentive compatible for the Principal to offer L' to High and amount z' to Low, where $z' < z$. Since z is greater than Low's first best, the principal does better offering L' and z' than offering L and z . Hence L cannot be optimal. ■

Proposition 3 (Profitability of screening)

Consider a fixed budget as an alternative to screening.

(i) If $\lim_{x \rightarrow x_{\min}} V'_L(x) / V'_H(x) = \infty$, screening is always profitable.

(ii) Suppose $\lim_{x \rightarrow x_{\min}} V'_L(x) / V'_H(x) = C < \infty$. (a) If Low is sufficiently risk-averse, then it pays to screen. (b) If screening is profitable at some fraction of Low types, it remains profitable at all greater fractions of Low types.

Proof. Part (i) is trivial: Since we know under the stated condition, the second-best approximates the first-best arbitrarily closely, the principal can always find a way to do better than with a fixed budget.

Part (ii)(a): The principal's problem is to

$$\begin{aligned} & \max_{p,b,s,z} \lambda [\{pV_H(b) + (1-p)V_H(s)\} - \{pb + (1-p)s\}] \\ & + (1-\lambda)[V_L(z) - z] \end{aligned} \tag{12}$$

s.t.

$$V_L(z) = V_L(ce(b, s, p)). \tag{IC-L}$$

where $ce(b, s, p) = pV_L(b) + (1-p)V_L(s)$, i.e., Low's certainty equivalent for the lottery $(b, p; s, (1-p))$.

Consider the lottery with $b = x_H^{FB}$, $s = x_{\min}$ and $z = x_L^{FB}$. Three features of this lottery are worth noting. First, this lottery gives the principal less than the second-best welfare (because it deviates from the lottery we showed to be optimal in Proposition 3). Second, note that as Low becomes infinitely risk-averse, $ce(b, x_{\min}, p)$ goes to zero. Hence, for Low sufficiently risk-averse, for any p there is a level of risk aversion of Low such that this lottery satisfies incentive compatibility. Third, assume that the proposed lottery is incentive-compatible. Then, as $p \rightarrow 1$, the optimal value of the principal's objective function approaches the first-best.

Construct the following sequence: $(\frac{1}{n}, V_L^n)$ where $\frac{1}{n}$ is the probability of the x_{\min} prize, and V_L^n is a utility function for Low with the following properties:

$$\frac{dV_L^n(x_L^{FB})}{dx} = 1$$

and V_L^{n+1} is more risk-averse than V_L^n and

$$V_L^n \left(ce \left(b, x_{\min}, 1 - \frac{1}{n} \right) \right) \leq V_L(x_L^{FB}).$$

Thus, $\frac{1}{n}$ represents a sequence of probabilities going to zero, and V_L^n represents a sequence of increasingly risk-averse utility functions for Low chosen so that Low prefers his first-best amount for sure to the lottery intended for High.

By construction, this sequence is incentive compatible, and the value of the objective function converges to the first-best as $n \rightarrow \infty$. Finally, since no allocation mechanism can do better than the first-best, and since this lottery performs worse than the actual second-best lottery characterized in Proposition 3, we know that if Low is sufficiently risk-averse, it pays to screen.

Part (ii)(b): Call the optimized value of the principal's objective function in the above problem $OF(\lambda)$. Note that $OF(\lambda)$ is continuous by, for example, Berge's theorem. It is optimal to screen when there exist p, b, s , and z such that $OF(\lambda)$ is greater than the optimal non-screening contract.

The optimal non-screening contract (yielding a single fixed budget) maximizes $\lambda V_H(y) + (1 - \lambda) V_L(y) - y$. Let $y(\lambda)$ denote the y that solves this problem for a particular λ , and let $V^{NS}(\lambda)$ denote the principal's maximal utility in the no-screening problem. Note that $y(0) = x_L^{FB}$, $y(1) = x_H^{FB}$, and $y(\lambda)$ is increasing in λ . Since $V_H(y) > V_L(y)$, and increasing λ puts more weight on V_H and increases the optimal y , $V^{NS}(\lambda) = \lambda V_H(y(\lambda)) + (1 - \lambda) V_L(y(\lambda)) - y(\lambda)$ is increasing in λ . In other words, $V^{NS}(\lambda)$ represents the "bar" for when it is optimal to screen, and this bar decreases as λ decreases, i.e., the greater the fraction of Lows. The more Lows there are, the easier it is for screening to be optimal, for a given welfare derived from screening. It is true that the welfare for the principal in the screening case also decreases when λ decreases. But the first effect dominates the second. To see this, suppose it is optimal to screen when $\lambda = \hat{\lambda}$. This means that the value of the constrained is such that $OF(\hat{\lambda}) > V^{NS}(\hat{\lambda})$. Since $OF(\hat{\lambda})$ is continuous, $OF(\lambda) > V^{NS}(\hat{\lambda})$ for λ close to $\hat{\lambda}$. For $\lambda < \hat{\lambda}$, $V^{NS}(\lambda) < V^{NS}(\hat{\lambda})$. Hence for $\lambda < \hat{\lambda}$ we have $OF(\lambda) > V^{NS}(\hat{\lambda}) > V^{NS}(\lambda)$. This completes the proof. ■

Appendix

November 19, 2012

This Appendix describes the details of the survey discussed in Section 4.1 of the paper as well as the numerical results mentioned in Section 4.5 of the paper.

A. Survey

As explained in the paper, we conducted surveys among Masters and Ph.D. students in Public Policy at Harvard University (the "Harvard sample"), two separate groups of undergraduate students of economics and psychology at the University of Zurich ("undergraduate sample 1" and "undergraduate sample 2"), and undergraduate students of finance at the University of Zurich (the "finance sample"). The undergraduate sample 1 and 2 participants were paid a fee of Swiss Franc 15 for filling out the questionnaire. In addition to filling out the survey, the undergraduate sample students had an opportunity to earn another Swiss Franc 15 in a non-related cognitive task designed for another project. The order of tasks was randomized and there was no statistically significant difference in answers depending on the order. The other subjects only filled out the questionnaire. The finance sample was collected through an online survey, the others were collected by paper surveys. In the finance sample, 63% of the students answered only the question in situation 1, while 37% of the students only answered situation 2.

For the analysis shown in the paper, we only use the questionnaires from individuals who answered the risk aversion questions consistently (which was the vast majority of individuals). Thus, we did not use data of individuals who accepted a lottery where $x = 30\%$, but rejected one where $x = 10\%$. We also did not use data of individuals who claimed their x was 23% but still accepted a lottery with $x = 50\%$. Including the data of these (very few) individuals does not materially affect the results.

In the survey, the currency of the respective country was used. In the US surveys, this was the US\$. In the Swiss surveys, this was the Swiss Franc.

1. Please think about your total yearly income. Please put a checkmark next to the following elements which form part of your definition of income.

- a) Salary and other labor income, such as bonuses
- b) Interests, dividends, rental income, etc.
- c) Financial support from family members
- d) Stipends
- e) Other, please list here _____

2. Please consider your total yearly income after taxes. There is no "correct" definition; please keep as comprehensive a definition of "income" in mind as possible.

Let us suppose that you have the chance to double this income with a 50% probability. But to obtain this chance, you also have to accept the possibility that your income decreases by x% with 50% probability. That is, we are flipping a coin. Assume, for example, that your income is \$US/Swiss Franc 30,000 and $x = 20\%$. Then, after the coin is flipped, you either have an income of \$US/Swiss Franc 60,000 or your income is reduced to \$US/Swiss Franc 24,000. Both events are equally likely, since heads and tails occur with equal probability.

Are you willing to accept this opportunity and this risk, if

a) $x = 10\%$, i.e., your income doubles with 50% probability and decreases by 10% with 50% probability

YES – NO

b) $x = 90\%$, i.e., your income doubles with 50% probability and decreases by 90% with 50% probability

YES – NO

c) $x = 20\%$, i.e., your income doubles with 50% probability and decreases by 20% with 50% probability

YES – NO

d) $x = 50\%$, i.e., your income doubles with 50% probability and decreases by 50% with 50% probability

YES – NO

e) $x = 30\%$, i.e., your income doubles with 50% probability and decreases by 30% with 50% probability

YES – NO

f) $x = 40\%$, i.e., your income doubles with 50% probability and decreases by 40% with 50% probability

YES – NO

3. Assume now that you can choose the maximum loss level that you are willing to accept, i.e., you can choose x between 0% and 100%. In other words, you will double your income, or lose x% of it, both with the same probability. How high would you set x% and still accept to play this "lottery"?

x = _____ %

4. Please imagine that you find a \$US/Swiss Franc 50 note on the street. It is impossible to identify the owner, and it is, therefore, completely acceptable and morally unobjectionable that you keep the \$US/Swiss Franc 50. Think about your average peer who earns about the same amount of money as you do, and is approximately equally wealthy. Would you say that, relative to this average peer, you benefit

- a lot more
- more
- equally
- less
- a lot less

from this additional amount of money?

5. Please imagine now that a very distant relative, whom you have never heard of dies and happens to leave you \$US/Swiss Franc 500,000. Think about your average peer who earns about the same amount of money as you do, and is approximately equally wealthy. Would you say that, relative to this average peer, you benefit

- a lot more
- more
- equally
- less
- a lot less

from this additional amount of money?

B. Numerical results

For concreteness, we consider power (CRRA) and CARA benefit functions. For CRRA benefits,

$$V_H(x) = \frac{x^{1-\gamma_H}}{1-\gamma_H}, \quad V_L(x) = \frac{x^{1-\gamma_L}}{1-\gamma_L}, \quad (1)$$

with $\gamma_H < \gamma_L$, implying that High has greater marginal benefits than Low for all $x > 1$. For CARA,

$$V_H(x) = -(1/r_H)e^{-r_H x}, \quad V_L(x) = -(1/r_L)e^{-r_L x}, \quad (2)$$

with $r_H < r_L$, implying that High has greater marginal benefits than Low for all x .

To make our analysis applicable to both the corporate and the consumption cases, we choose parameters for the curvature of the benefit functions that imply typical values for risk aversion. For CRRA, values in the range of 1 and 5 seem reasonable.¹ Calibration is more difficult for CARA benefits, but we choose a range between 0.01 and 0.1. We also choose marginal costs for the principal to be a meaningful constant for the two respective cases: $c_{CRRA} = 0.5$ and $c_{CARA} = 0.15$.²




Figure B-1: Marginal benefits and costs given CRRA and CARA utility functions. RRA denotes relative risk aversion coefficients. ARA denotes absolute risk aversion coefficients.

The left panel in Figure B-1 plots benefits (left scale) and marginal benefits and marginal costs (right scale) for the case of CRRA benefits. The right panel does the same for CARA. Note that while High is appropriately labeled High because his marginal utility at Low's first-best is higher than Low's, it is nonetheless the case that Low's marginal utility is higher at very low levels of funds. In the case of CRRA, Low's benefits fall faster without bound than High's as funds go to zero. By contrast, with CARA benefits, marginal utility is bounded from above for non-negative funds, and utility is bounded from below. Thus, we expect almost first-best welfare for the principal in

¹In this numerical application we do not distinguish between consumption and wealth. See Meyer and Meyer (2005) for a discussion of how risk aversion estimates can be compared across these two cases.

²We had set $c = 1$ for the general analysis above. For CRRA, since the marginal utilities of High and Low cross exactly at unity, if we had $c = 1$, the first-best would be to give both types the same fixed budget. If we had $c > 1$, High should receive less than Low, making the terminology of High and Low inappropriate.

the case of CRRA, while only some of the welfare loss can be recaptured in the case of CARA and non-negative payments.

B.1 First-best

In what follows, we document the setup for CRRA benefits. To conserve space, we do not report the analogous setup for CARA utility. Under perfect information, the principal maximizes $\frac{(x_i)^{1-\gamma_i}}{1-\gamma_i} - c_{CRRA}x_i$ for type i . Thus, $x_i^{FB} = \left(\frac{1}{c_{CRRA}}\right)^{\frac{1}{\gamma_i}}$. Under our assumptions, this implies $x_H^{FB} > x_L^{FB}$. Denote the corresponding maximized expected utility of the principal (net of costs of funds) in the first-best by U_P^{FB} . In Figure B-1, the first-best allocation can be seen where the marginal benefits and marginal cost schedules intersect.

B.2 Fixed budgets

As discussed earlier, when information about productivity is private, one option for the principal is to allocate identical funds, \bar{y} , to both types. That is, the principal does not screen. The corresponding maximized expected welfare of the principal (net of costs of funds) is U_P^{fixed} .

B.3 Screening through risk aversion

For CRRA benefits, the principal's problem is to

$$\begin{aligned} \max \lambda \frac{1}{1-\gamma_H} [\{pb^{1-\gamma_H} + (1-p)s^{1-\gamma_H}\} - c_{CRRA}\{pb + (1-p)s\}] + (1-\lambda) \left[\frac{1}{1-\gamma_L} (z)^{1-\gamma_L} - c_{CRRA}z \right] \\ \text{s.t. } z^{1-\gamma_L} = pb^{1-\gamma_L} + (1-p)s^{1-\gamma_L}. \end{aligned}$$

This problem can only be solved numerically. Denote the corresponding maximized expected utility of the principal (net of costs of funds) by U_P^{SB} . We define the recovery rate R as

$$R = 1 - \frac{U_P^{FB} - U_P^{SB}}{U_P^{FB} - U_P^{fixed}}.$$

The closer this number is to 100%, the more powerful screening through risk aversion is in terms of allowing the principal to recapture the welfare losses when only fixed budgets are available and the agent types cannot be distinguished.

B.4 Results for CRRA benefits

Table B-1 illustrates the results for the case of CRRA benefits when $x_{\min} = 0$. Strictly speaking, there is no solution in this case, as for any candidate solution, the principal can always improve by moving High's good allocation closer to the first-best, High's bad allocation further down, the probability on High's good allocation closer to unity, and Low's allocation further down. (Thus,

the results obtained in numerical optimization depend on the sensitivity level one allows for the optimization algorithm.)

The table holds all parameters except High’s risk aversion fixed at the values given in the notes to the table. As predicted by the analysis, the principal implements an extreme lottery for High, which puts almost probability one on High’s first-best, and an almost zero probability on a very low, almost zero allocation. Low receives ever so slightly more than his first-best. For example, when $\gamma_H = 1.5$, the principal chooses $p = 0.9999987$, $b = 1.587379$ (which is only a little bit smaller than $x_H^{FB} = 1.5874$), $s = 0.00271$, and $z = 1.1893$ (which is only a little bit larger than $x_L^{FB} = 1.1892$).

Table B-1: CRRA benefits

Relative risk aversion of High	First-Best		Second-Best				Recovery rate %
	x_H^{FB}	x_L^{FB}	p	b	s	z	
1.00	2.000	1.189	~1	~2.000	~0	~1.189	~100%
1.50	1.587	1.189	~1	~1.587	~0	~1.189	~100%
2.00	1.414	1.189	~1	~1.414	~0	~1.189	~100%
2.50	1.320	1.189	~1	~1.319	~0	~1.189	~100%
3.00	1.260	1.189	~1	~1.260	~0	~1.189	~100%
3.50	1.219	1.189	~1	~1.219	~0	~1.189	~100%

Notes: Relative risk aversion of Low (γ_L)= 4, marginal cost of funds to the principal (c_{CRRA})= 0.5, proportion of High types in the population (λ)= 0.5. The approximate numbers for second-best allocations are merely illustrative. The problem, strictly speaking, has no solution since more extreme values improve the outcome.

The welfare consequences of using this screening mechanism are striking: Screening through risk taking achieves virtually the same expected welfare as the first-best. The possibility of a low allocation for High is so remote that even a risk-averse principal is in expectation almost equally well off in the second-best as in the first-best.

B.5 Results for CARA benefits

The numerical analysis also helps to illustrate the analytical results for benefit functions of the CARA class. Table B-2 shows these results for three absolute risk aversion levels for High (0.015, 0.025, and 0.05) for a given risk aversion of Low (0.06). In addition, the table presents the results for different lower bounds. If there is no lower bound, the first-best can be approximated arbitrarily closely. What is more interesting is to study what happens when there is a lower bound. With CARA utility, higher risk aversion of Low implies that Low is also more downside risk-averse than High. Thus, the model predicts that the principal optimally gives just the lower bound to High in the bad state. The simulation results confirm this prediction. Even with that threat, however, the principal cannot approximate first-best welfare as closely as in the CRRA case. The recovery rate with non-negative payments in the CARA case when High is sufficiently different from Low (0.015 vs. 0.06) is still substantial at close to two thirds of the welfare differential between the first-best and fixed budget. By contrast, when High has risk aversion 0.025, only 27% can be recovered.

Things improve dramatically when the principal can threaten to take away initial endowments with even a very small probability. As the rows with lower bounds of -10 and -100 show, the numerical analysis confirms that the center achieves results closer and closer to the first-best.

Table B-2: CARA benefits

Absolute risk aversion of High	Lower bound	First-Best		Second-Best				Recovery rate %
		x_H^{FB}	x_L^{FB}	p	b	s	z	
With $c_{CARA} = 0.15$								
0.015	0.00	126.48	31.62	0.9545	125.95	0.00	51.33	63.26%
0.015	-10.00	126.48	31.62	0.9680	126.12	-10.00	47.20	70.92%
0.015	-100.00	126.48	31.62	0.9996	126.47	-100.00	32.24	98.80%
0.015	-inf	126.48	31.62	~1	~126.48	--inf	~31.62	~100%
0.025	0.00	75.88	31.62	0.9502	71.03	0.00	46.03	27.43%
0.025	-10.00	75.88	31.62	0.9673	72.30	-10.00	43.80	37.73%
0.025	-100.00	75.88	31.62	0.9997	75.69	-100.00	32.71	94.16%
0.025	-inf	75.88	31.62	~1	~75.88	--inf	31.62	~100%
0.050	0.00	37.94	31.62	1.0000	34.76	nA	34.76	0.00%
0.050	-10.00	37.94	31.62	1.0000	34.76	nA	34.76	0.00%
0.050	-100.00	37.94	31.62	1.0000	34.76	nA	34.76	0.00%
0.050	-inf	37.94	31.62	~1	37.94	--inf	~31.62	~100%
With $c_{CARA} = 0.0015$								
0.050	0.00	130.05	108.37	1.0000	121.38	nA	121.38	0.00%
0.050	-10.00	130.05	108.37	0.9999	122.42	-10.00	120.56	0.72%
0.050	-100.00	130.05	108.37	0.9999	127.41	-100.00	114.47	36.20%
0.050	-inf	130.05	108.37	~1	~130.05	--inf	~108.37	~100%

Notes: Absolute risk aversion of Low (r_L) = 0.06. Proportion of High types in the population (λ) = 0.5.

Importantly, and perhaps surprisingly at first, the principal and High benefit if negative payments are allowed, while Low suffers. The reason is that when the principal is constrained in her design of High's lottery, she has to distort Low's allocation and give him too much. For example, in the Table, when High has risk aversion of 0.015 and the parameters are as described, Low receives 51.33 for sure when no negative payments are allowed, but only 31.62 (effectively the first-best) when negative payments (to High) are allowed. By contrast, even though High has to accept the possibility of a large negative payment, he is, in expectation, better off than without the possibility of negative payments. The intuition is that when another may envy you, opening yourself up to a comparatively cheap (your cost to his) penalty is likely to be desirable.

If the two types are too similar (such as when one has risk aversion of 0.05 and the other has 0.06, as shown in the third panel in Table B-2), then screening is not worthwhile even when the principal can take away up to 100. Only with even larger negative payments in the bad state can the first-best be approached. However, note that the recovery rate in this case is not a very good measure of the outperformance of screening over fixed budgets. When the types are very similar, fixed budgets do reasonably well, unless the cost of funds is such that in the first-best the principal

would like to give very different allocations to High and Low. This is illustrated in the final panel. When the cost of funds is significantly smaller, screening starts to pay, first very modestly with moderate negative payments and more so with higher negative payments. Similarly (not shown), screening pays when the fraction of Lows increases sufficiently.