Supplementary Online Appendix for "Solomonic Separation: Risk Decisions as Productivity Indicators" (Journal of Risk and Uncertainty 2013)

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

March 28, 2013

This Appendix describes the numerical results mentioned in Section 4.5 of the paper.

1 Benefit functions

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

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

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

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

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]

[1] In 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.

[2] We 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.

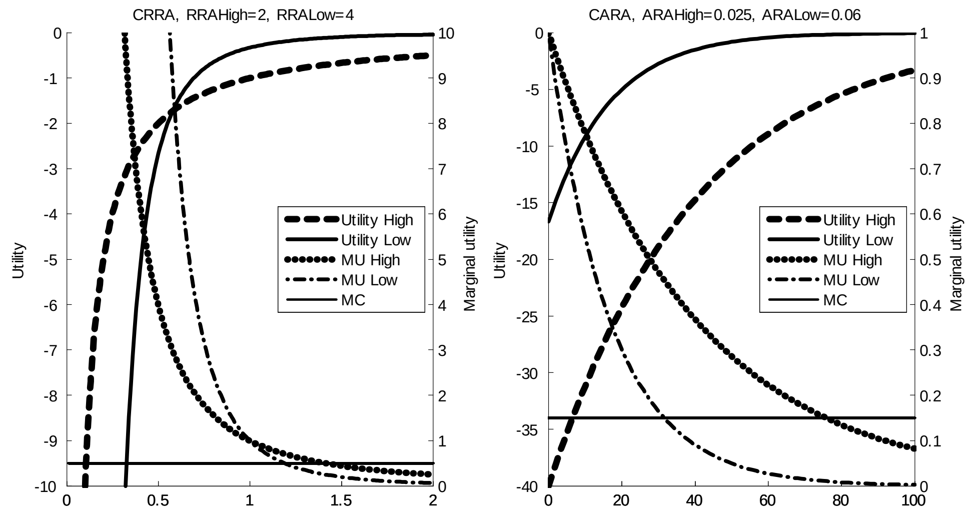

Figure S-1: Marginal benefits and costs given CRRA and CARA utility functions. RRA denotes relative risk aversion coefficients. ARA denotes absolute risk aversion coefficients.

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

2 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.

3 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}.

4 Screening through risk aversion

For CRRA benefits, the principal's problem is to

$$\max \lambda \frac{1}{1-\gamma_H} \left[\left\{ pb^{1-\gamma_H} + (1-p)\, s^{1-\gamma_H} \right\} - c_{CRRA} \left\{ pb + (1-p)\, s \right\} \right] + \quad \text{(S-3)}$$

$$+ (1-\lambda) \left[\frac{1}{1-\gamma_L} (z)^{1-\gamma_L} - c_{CRRA} z \right] \quad (1)$$

$$\text{s.t. } z^{1-\gamma_L} = pb^{1-\gamma_L} + (1-p)\, s^{1-\gamma_L}. \quad \text{(S-4)}$$

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

$$R = 1 - \frac{U_P^{FB} - U_P^{SB}}{U_P^{FB} - U_P^{fixed}}. \quad \text{(S-5)}$$

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.

5 Results for CRRA benefits

Table S-1 illustrates the results for the case of CRRA benefits when $x_{\min} = 0$. Strictly speaking, there is no solution in this case. 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 S-1: CRRA benefits

	First-Best		Second-Best				
Relative risk aversion of High	x_H^{FB}	x_L^{FB}	p	b	s	z	Recovery rate %
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.

6 Results for CARA benefits

The numerical analysis also helps to illustrate the analytical results for benefit functions of the CARA class. Table S-2 shows these results for three absolute risk aversion levels for High (0.015, 0.025, and 0.05) and 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 S-2: CARA benefits

		First-Best		Second-Best				
Absolute risk aversion of High	Lower bound	x_H^{FB}	x_L^{FB}	p	b	s	z	Recovery rate %
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 $(\lambda) = 0.5$.

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, 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 S-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.