

Online Appendix

“How General are Risk Preferences? Choices under Uncertainty in Different Domains”

by Einav, Finkelstein, Pascu, and Cullen

Appendix A: A calibrated model of short-term and long-term disability insurance choices.

In this appendix we describe the details of the calibration exercise that we mention in the beginning of Section III. The objective of the calibration exercise is to illustrate how one could produce a benchmark for the correlation coefficients that would be produced if the data were generated by a model with completely domain-general risk preferences, but were subject to the non-linearities and discreteness transformations that arise from the ordinal coverage choice data. We focus on short- and long-term disability, which are the domains that are most similar to each other in their structure of choices and risks. This allows us to rely on a single choice model for both domains, rather than on a domain-specific model. We estimate the correlation in the simulated choices between the modal short-term disability menu (of 60%, 80%, and 100% replacement rates) and the modal long-term disability menu (of 50%, 60%, and 70% replacement rates), using the observed prices.

Our calibration exercise assumes a constant relative risk aversion (CRRA) per-period utility function, whereby the expected utility from a given disability insurance contract j (which specifies a given wage replacement rate and is associated with a given annual premium) is

$$Eu(c) = E_{\tilde{d}} \left[\left(1 - \tilde{d} + RR_j * \tilde{d} - p_j \right)^{1-\gamma} \right],$$

where \tilde{d} is the (ex-ante random) fraction of days in a year the employee claims (due to disability), RR_j is the wage replacement rate associated with coverage j , the premium p_j is measured per dollar of (annual) wage, and γ is the coefficient of relative risk aversion. The individual is maximizing expected utility over the duration under consideration, which we assume to be one year for short-term disability and four years for long-term disability (as after about four years, our claim data is truncated, although only few disability claims in the data remain active that long). We assume an annual discount factor of .95 for long-term disability.

We assume that the distribution of γ in the population to be lognormal with parameters μ and σ , such that the values of μ and σ are chosen to produce an average relative risk aversion coefficient of 3 or 0.7 (depending on the specification) and a coefficient of variation of risk aversion of 10. The coefficient of variation (of 10) matches the estimates reported by Cohen and Einav (2007). Cohen and Einav (2007) mention higher numbers of relative risk aversion, but they essentially estimate absolute risk aversion, so mapping it to this lower levels of relative risk aversion amounts to simply assuming lower relevant wealth (the simulated correlations remain about the same when we instead use an average coefficient of risk aversion of 30, and maintain the same coefficient of variation). To calibrate the distribution of risk (missed days), we use the risk realization of short- and long-term disability in the data to define eight risk groups based on demographics (using all combinations of gender, race, and employment status indicators), which produces a fairly large heterogeneity in ex ante risk across individuals. We assume a sample size identical to our baseline

sample (12,752) and for each individual in the calibrated sample we draw a risk aversion coefficient from the distribution of γ , assume that he or she knows the distribution of risks for his or her risk group, and compute the optimal coverage choice from the offered menus of short- and long-term disability coverage.

Using this model we simulated choices from the modal short- and long-term disability menus offered in the data, and correlated these choices with each other. When we calibrate the average coefficient of relative risk aversion in the sample to 3, we find that the correlation between short- and long-term disability insurance range from 0.18 to 0.55, depending on the extent of correlation in risks we allow across domains. The range is somewhat larger (0.07 to 0.53) when we calibrate the average coefficient of relative risk aversion to 0.7. These reported ranges are just below the empirical correlation of about 0.6 for these two domains, as reported in Table 3.

Appendix B: Implementation details for the model-based approach of Section III

This appendix provides additional details that underlie the baseline exercise reported in Section III (and Table 8).

Health, Dental, and Drug coverage. The risks in these three domains are measured in dollars. Therefore, for our baseline estimates, we assume a CARA utility function in these domains. That is, we use equation (3) to compute individual i 's expected utility from option j by substituting $u_i(x) = -\exp(-r_i x)$, incorporating the plan details (regarding deductible and out-of-pocket maximum) to compute $oop_j(c)$, and grouping individuals into risk categories by their coverage tier (single coverage, employee plus spouse, employee plus children, and family coverage) and randomly drawing from individuals' realization of total medical expenditure c .

Short-term and Long-term disability coverage. The risks (and premiums) in these domains are all proportional to the employee's (annual) wage. It is therefore natural to assume a CRRA utility function for these two domains. Again, we use equation (3) to compute individual i 's expected utility from option j by substituting $u_i(x) = x^{1-\gamma_i}$, assuming all individuals are grouped at the same risk, drawing the claimed disability days for each individual, and computing $oop_j(d) = (1 - RR_j)d$ where RR_j is the wage replacement rate associated with coverage j . A minor complication arises in the case of long-term disability coverage, where the data on realized risks is slightly censored (for long spells of disability absence), so we impute the full predicted absence based on the observed propensity to remain on (long-term) disability between the first and second year.

Determining cutoffs and defining intervals. Given a value of λ_d for domain d with J options, we partition $[0, \infty]$ into J intervals $[r_1(\lambda_d) = 0, r_2(\lambda_d)], \dots, [r_k(\lambda_d), r_{k+1}(\lambda_d)], \dots, [r_J(\lambda_d), r_{J+1}(\lambda_d) = \infty]$, such that an individual with a given distribution of risk and a risk aversion parameter in interval $[r_k(\lambda_d), r_{k+1}(\lambda_d)]$ will choose option k . For a given λ_d , a menu of options and distribution of risk, we first find the level of risk aversion $r_{k,k+1}(\lambda_d)$ such that an individual is indifferent between choices k and $k+1$, where option $k+1$ has the higher premium and higher coverage. There are a couple of cases to bear in mind:

- If a risk neutral individual prefers option $k+1$ over option k then option k is dominated and choice $k+1$ cannot be rationalized. In such a case, some of the intervals will be empty.

- For lower values of lambda, the risk might be small enough that option k should never be chosen. In the limit, if $\lambda = 0$, then only the lowest coverage option can be rationalized, and again some of the intervals are empty.
- For all other cases, we can find a cutoff value such that an individual faced with option k and $k + 1$ will choose option k for $r < r_{k,k+1}$, and option $k + 1$ if $r > r_{k,k+1}$.

Using the procedure described above, we compute $J(J-1)/2$ cutoff values for each pair of options, which define the relevant intervals of risk aversion implied by each coverage choice in the data.

401(k) choice. Here, because the decision is continuous, we slightly deviate from the description provided in the paper, and instead rely on the same exercise performed in the seminal paper of Friend and Blume (1975). As they show, one can convert one’s share invested in a risky asset α_i to one’s coefficient of relative risk aversion γ_i , by applying $\gamma_i = \frac{1}{\alpha_i} \frac{E(r_m - r_f)}{\sigma_m^2}$. Our inclusion of a domain-specific adjustment $\lambda_{401(k)}$ simply implies that $\lambda_{401(k)}$ multiplies the right-hand-side, illustrating how this manipulation frees up the level of risk aversion. We use the average return of the safe funds to compute the (monthly) risk free return $E(r_f) = 0.0036$. We aggregate all the funds in our sample invested in the risky funds to compute an estimate of the expected (monthly) return of the “risky” asset and its standard deviation, which are given by $E(r_m) = 0.0103$ and $\sigma_m = 0.0285$. Taken together, this implies that $\frac{E(r_m - r_f)}{\sigma_m^2} = 8.35$. We further assume that people who invest all their 401(k) contributions in the risky funds are at a corner solution, implying that for such individuals we obtain that $\gamma_i \in \left[0, \lambda_{401(k)} \frac{E(r_m - r_f)}{\sigma_m^2}\right]$.

Conversion between absolute and relative risk aversion. For each individual we have three intervals for their value of absolute risk aversion, based on their health, dental, and drug coverage choice, and three intervals (or a point in some cases for 401(k)) for the value of their relative risk aversion from short-term and long-term disability coverage and their 401(k) allocation. To evaluate the consistency of choices across all six domains, we need to convert the absolute risk aversion intervals to relative-risk aversion. We use another free parameter γ (which could be interpreted as converting annual income to wealth), as well as the data on the individual’s annual income, such that $RRA = ARA \times wage \times \gamma$, where RRA and ARA are the coefficients of relative and absolute risk aversion, respectively. In Appendix Tables A3(a) and A3(b) we report results from analogous specifications where we use CRRA and CARA, respectively, for all domains, and use the observed data on annual income to convert absolute risk to relative and vice versa.

Appendix Table A1: Coverage Details for Insurance Plans

Summary of Key Coverage Details (1)		Additional details (2)
Health Insurance ^a	Deductible (In-network / out-of-network)	
Option 1 ^b	3,000 / 6,000	
Option 2	1,500 / 3,000	After satisfying the annual deductible, cost sharing is 10% in-network and 30% out-of-network for all options. All options also specify in-network and out-of-network out-of-pocket maximums, but these are rarely binding. Preventive care is covered in full under all coverage options.
Option 3	1,000 / 2,000	
Option 4	500 / 1,000	
Option 5	0 / 500	
Prescription Drug Insurance	Cost sharing for branded drugs (retail / mail order)	
Option 1	50% / 40%	All options have cost-sharing of 10% for generic (non-branded) mail order drugs and 20% for generic retail drugs. All options have a \$50 deductible (\$100 for family) and a \$50 (\$100 for mail-order) maximum per prescription.
Option 2	40% / 30%	
Option 3	30% / 20%	
Dental Insurance	Per person Deductible / Maximum annual benefit	The family deductible is double the per-person amount. Both plans fully cover preventative care, provide identical coverage for other special treatments. Oral surgery is covered at 50% under option 1 and 100% under option 2. Orthodontia is not covered under option 1 and is covered at 50% under option 2.
Option 1	50 / 1000	
Option 2	25 / 2000	
Short-Term Disability Insurance ^c	Wage replacement rate	Salary workers have 100% replacement rate for first two weeks of disability under all options; all options provide up to 26 weeks of benefits.
Option 1	mostly 60% (sometimes 40%)	
Option 2	mostly 80% (sometimes 60%)	
Option 3	mostly 100% (sometimes 80%)	
Long-Term Disability Insurance ^c	Replacement rate	All long-term disability coverage is payable after 26 weeks of disability (when the short-term disability coverage is capped).
Option 1	mostly 50%	
Option 2	mostly 60%	
Option 3	mostly 70%	

All options are shown in the ordinal ranking from more (option 1) to less risk exposure (with the possible exception of health insurance option 1; see note b and text for details). Column (1) summarizes key features of each option. Column (2) provides additional details.

^a Health insurance: deductibles are shown for the non-single coverage tier; deductibles for single coverage are half what is shown.

^b Option 1 includes a Health Reimbursement Account (HRA) in which Alcoa contributes \$1,250 in tax free money each year that the employee can use to fund eligible out of pocket health care expenses. Any balance remaining at the end of the year can be rolled over to pay for future out of pocket costs. See text for more details.

^c Short-term and Long-term disability benefits (column (1)) are proportional to the employee's wage.

Appendix Table A2: List of funds available for 401(k) allocation

Fund name (Asset Class)	Monthly return				
	Share ^a	Mean	Std. Dev.	Min.	Max.
<u>Classified (by us) as "Risk Free":</u>					
GIC/Stable Value (Fixed Income)	24.47%	0.35	0.02	0.31	0.37
Vanguard Total Bond	3.95%	0.42	0.83	-1.09	1.92
<u>All other classified as risky:</u>					
American Balanced (Balanced Equity)	10.58%	0.65	1.36	-2.34	2.89
Inv. Co. of America (Large Cap US Equity)	9.62%	0.83	1.84	-3.82	3.86
AMCAP (Large US Equity)	6.77%	0.66	2.06	-4.19	4.01
Vanguard Institutional Index (Large Cap US Equity)	9.42%	0.79	2.21	-4.18	4.43
MSDW International Equity	4.09%	1.25	2.32	-3.30	4.92
New Perspective (International Equity)	5.34%	1.49	2.72	-4.13	6.32
Putnam OTC (Mid Cap US Equity)	3.23%	1.01	3.40	-6.35	7.45
Small Cap Core (Small Cap US Equity)	0.30%	0.29	3.44	-6.95	7.90
Putnam Vista (Mid Cap US Equity)	3.71%	0.56	3.55	-8.58	6.75
MSDW Emerging Markets	2.62%	3.13	5.83	-11.69	15.03
Company (Alcoa) Stock Fund	15.90%	1.30	6.71	-8.85	16.79
<u>Benchmarks during the same period:</u>					
Risk free ^b	--	0.37	0.05	0.26	0.43
S&P 500	--	0.63	2.21	-4.40	4.33

Employee contributions to their 401(k) accounts can be made with either pre- or after-tax dollars. Employees can contribute 1-16% of eligible pay with some additional restrictions for some highly paid employees. In our sample, Alcoa usually matches 100% of pre-tax contributions, up to 6% of eligible pay. Employer (Alcoa) contributions are always invested in the company stock and can only be moved to a different fund after two years. In the 2004 data that we are using, the above 13 funds are available for contributions (sorted by the standard deviations of monthly returns). In the analysis we use as a measure of riskiness of the portfolio the share of employee contributions invested in those (two) funds that are presented as least risky. Indeed, as apparent from the table, these two funds exhibit less volatility (and mostly lower expected return). Employees also have the option to invest in a personal choice retirement account in which they have access to other funds besides the 13 funds just described. Direct contributions to this fund are not possible, only transfers, and we do not have detailed data on the composition of investments in these funds. For our analysis we only use direct employee contributions. In 2004 only about 28 percent of the sample rebalances and 24 percent of the sample changes the allocation of their contributions. The average employee contribution in the baseline sample (which restricts attention to non-zero contributions) is around \$4,600. About 40 percent of the sample has no contributions to the risk free funds, and about 17 percent invest all their contributions in the risk free funds. Just over 40 percent of the sample has some employee contributions invested in company stock. The series of returns are based on monthly returns over the 29 month period from August 2005 to December 2007, which was the longest time period for which we have consistent returns data for all funds. Returns data are from Collaborative Research Support Program (CRSP) when available, or from Hewitt (when CRSP data are not available, for the few funds that are not publicly traded).

^a We compute the share of dollars contributed to each fund out of total 401(k) contributions made by all employees in our baseline sample.

^b For the risk free benchmark we use the CRSP three month Fama Risk Free Rates series, which are derived from average lending and borrowing rates.

Appendix Table A3(a): Model-based results using a CRRA utility function in all domains

	Obs.	All domains	All insurance domains	Three CARA domains	Three CRRA domains
	(1)	(2)	(3)	(4)	(5)
1 Minimum overlap	11,898	5%	26%	35%	10%
2 Baseline specification	11,898	31%	37%	54%	68%
3 Restricted: $\lambda = \gamma = 1$	11,898	9%	31%	44%	61%
4 Restricted: flexible λ only on 401(k); $\gamma = 1$	11,898	28%	--	--	61%
Results for different demographics groups:					
5 Females	2,666	32%	40%	56%	70%
6 Males	9,232	30%	36%	53%	69%
7 Over 55 years old	1,533	35%	43%	61%	69%
8 Under 35 years old	2,356	28%	34%	56%	64%
9 Higher wage	3,074	22%	27%	49%	64%
10 Lower wage	2,848	40%	47%	61%	76%
Alternative specifications:					
11 Discretize 401(k)	11,898	31%	--	--	68%
12 Restricted: $\gamma = 1$	11,898	--	--	--	--
13 Alternative definition of income	11,898	31%	36%	54%	68%
Alternative samples:					
14 Housing Subsample	4,054	29%	35%	53%	67%
15 House Equity < \$50,000	1,305	32%	37%	55%	69%
16 Housing Equity \$50,000-\$150,000	1,453	30%	36%	53%	68%
17 Housing Equity > \$150,000	1,296	26%	31%	50%	64%
18 Maxed out 401(k) contributions	9,394	29%	35%	53%	68%
19 Did not max out 401(k) contributions	2,504	35%	42%	58%	70%

This table replicates Table 8 in the paper, but instead of using different utility functions for different domains and only then mapping between absolute and relative risk aversion, here we use a CRRA utility function for all domains by using the observed annual income to map the absolute (dollar) risks in Health, Drug, and Dental insurance to relative risks. Everything else parallels Table 8; see Table 8's notes for more details.

Appendix Table A3(b): Model-based results using a CARA utility function in all domains

	Obs.	All domains	All insurance domains	Three CARA domains	Three CRRA domains
	(1)	(2)	(3)	(4)	(5)
1 Minimum overlap	11,898	5%	26%	35%	10%
2 Baseline specification	11,898	26%	37%	55%	57%
3 Restricted: $\lambda = \gamma = 1$	11,898	8%	31%	44%	57%
4 Restricted: flexible λ only on 401(k); $\gamma = 1$	11,898	26%	--	--	57%
Results for different demographics groups:					
5 Females	2,666	27%	35%	48%	61%
6 Males	9,232	26%	34%	53%	57%
7 Over 55 years old	1,533	31%	42%	60%	61%
8 Under 35 years old	2,356	22%	32%	54%	50%
9 Higher wage	3,074	17%	28%	48%	46%
10 Lower wage	2,848	36%	47%	63%	70%
Alternative specifications:					
11 Discretize 401(k)	11,898	26%	--	--	57%
12 Restricted: $\gamma = 1$	11,898	--	--	--	--
13 Alternative definition of income	11,898	26%	35%	52%	57%
Alternative samples:					
14 Housing Subsample	4,054	25%	35%	54%	55%
15 House Equity < \$50,000	1,305	28%	38%	55%	57%
16 Housing Equity \$50,000-\$150,000	1,453	26%	37%	55%	58%
17 Housing Equity > \$150,000	1,296	21%	31%	51%	50%
18 Maxed out 401(k) contributions	9,394	25%	34%	52%	56%
19 Did not max out 401(k) contributions	2,504	31%	42%	61%	63%

This table replicates Table 8 in the paper, but instead of using different utility functions for different domains and only then mapping between absolute and relative risk aversion, here we use a CARA utility function for all domains by using the observed annual income to map the relative risk (in short- and long-term disability insurance and in 401(k)). to dollars. Everything else parallels Table 8; see Table 8's notes for more details.