Introduction

In a single period model, agents must decide today how much to consume and how much to save for later. In this note we take that decision as given, and assume that agents derive their utility from consumption at the end of the period, represented as is common in finance by a single good. Agents prefer more to less, but the marginal utility of each additional unit of consumption is decreasing, i.e., the last bite is never as good as the first one.

We capture this intuition by a utility function $u : \mathbb{R}^+ \to \mathbb{R}$, where $\mathbb{R}^+ = [0, \infty)$ since consumption cannot be negative. We assume throughout that u is continuous, differentiable of all orders, increasing and strictly concave,[1] that is, $u'(c) > 0$ and $u''(c) < 0$. The level of utility is not important and can even be negative. In some cases it will also be useful to consider utility functions such that $u'(0) = \infty$, that is, in starvation an extra unit of consumption provides an infinite amount of extra utility.

Two families of utility functions will recur throughout these notes. The first is **power utility**,

$$u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma}, \quad \gamma > 0, \gamma \neq 1,$$

with marginal utility $u'(c) = c^{-\gamma}$, together with **log utility**, $u(c) = \ln(c)$. Log utility is not a separate special case but the $\gamma \to 1$ limit of power utility.[2]

The second family is **exponential utility**,

$$u(c) = -e^{-ac}, \quad a > 0.$$

[1] A function $f : D \to \mathbb{R}$ is concave if

$$f(tx_1 + (1-t)x_2) \geq tf(x_1) + (1-t)f(x_2)$$

for all $x_1, x_2 \in D$ and $0 \leq t \leq 1$. The function is strictly concave if the inequality is strict.

[2] For small x, $y^x = \exp(x \ln(y)) \approx 1 + x \ln(y)$. Taking $x = 1 - \gamma$ and $y = c$,

$$\frac{c^{1-\gamma} - 1}{1 - \gamma} \approx \frac{(1-\gamma)\ln(c)}{1-\gamma} = \ln(c).$$

The error in the approximation is of order $(1-\gamma)^2$, so it vanishes as $\gamma \to 1$ and the limit is exactly $\ln(c)$.

We will see that power/log and exponential utility induce different attitudes under uncertainty.

Expected Utility and Risk Aversion

The analysis in the next two sections closely follows Chapter 1 in Ingersoll (1987). Denote by $\tilde{W}$ the wealth of the agent at the end of the period, a random variable defined over a probability space (Ω, P), where the tilde emphasizes that its value is unknown today. Since all wealth is consumed at the end of the single period, $\tilde{W}$ plays the role of end-of-period consumption, and we use $u(\tilde{W})$ interchangeably with $u(c)$ from here on.

Under uncertainty the utility of final consumption is random too, so the agent ranks gambles not by the consumption they deliver but by the **expected utility** they deliver:

$$U(\tilde{W}) = \mathrm{E}(u(\tilde{W})).$$

Thus, we have that $\tilde{W}_1 \gtrsim \tilde{W}_2$ if $U(\tilde{W}_1) \geq U(\tilde{W}_2)$, and $\tilde{W}_1 \sim \tilde{W}_2$ if $U(\tilde{W}_1) = U(\tilde{W}_2)$. Note that the expected utility of a certain level of wealth W is just $u(W)$.

Consider now a **fair gamble** $\tilde{\varepsilon}$, by which we mean a random variable with $\mathrm{E}(\tilde{\varepsilon}) = 0$ and $\mathrm{V}(\tilde{\varepsilon}) > 0$. Adding it to a certain wealth W gives $\tilde{W} = W + \tilde{\varepsilon}$, which has the same mean as W but strictly more variance. An agent is **risk-averse** if she prefers the certain wealth to the gamble, that is, if

$$u(W) > \mathrm{E}(u(W + \tilde{\varepsilon})). \tag{1}$$

Figure 1 illustrates why this holds for a strictly concave utility function. Take a gamble that pays $W - \Delta$ or $W + \Delta$ with equal probability, so that $\mathrm{E}(\tilde{W}) = W$. Expected utility is the height of the chord at W, and concavity places the chord strictly below the curve, so $\mathrm{E}(u(\tilde{W})) < u(W)$: the sure thing beats the gamble. The figure also marks two quantities defined in the next section, the certainty equivalent CE and the insurance premium $\Pi_i = W - CE$.

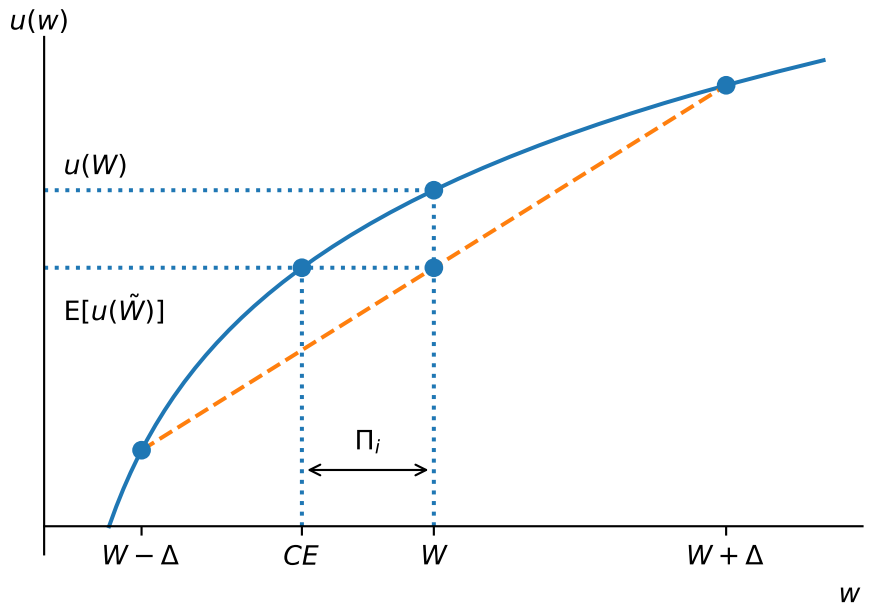


Figure 1: A concave utility function, Jensen's inequality, and the insurance premium.

The formal statement behind this argument is the following.

Property 1 (Jensen's Inequality)**.** *Let* $f : D \to \mathbb{R}$ *be a twice-continuously differentiable and strictly concave function, and* X *a random variable defined in a probability space* (Ω, P) *such that the range of* X *is contained in the domain of* f*, and* $\mathrm{V}(X) > 0$*. Then we have that*[3]

$$f(\mathrm{E}(X)) > \mathrm{E}(f(X)).$$

Jensen's inequality shows that strict concavity in u implies risk-aversion. The converse is also true: risk-aversion implies that u must be strictly concave.[4]

[3] Let $m = \mathrm{E}(X)$. Since f is strictly concave and differentiable, the tangent line at m lies strictly above the graph of f for all $x \neq m$:

$$f(x) < f(m) + f'(m)(x - m).$$

Setting $x = X$ and taking expectations,

$$\mathrm{E}(f(X)) < f(m) + f'(m)\ \mathrm{E}(X - m) = f(\mathrm{E}(X)),$$

where $\mathrm{E}(X - m) = 0$ follows from $m = \mathrm{E}(X)$, and the inequality is strict because $\mathrm{V}(X) > 0$ ensures $X \neq m$ with positive probability.

[4] Fix any x_1, x_2 in the domain of u and any $t \in (0, 1)$, and let $d = x_1 - x_2$ and $W = tx_1 + (1 - t)x_2$. Consider

Property 2 (Risk Aversion and Concavity of the Utility Function)**.** *An agent is risk-averse as defined in (1) if and only if her utility function is strictly concave.*

Strict concavity is therefore the defining mathematical feature of risk-averse preferences. It also lets us ask not just whether an agent is risk-averse but how much so, a comparison we make concrete next by asking what she would pay to get rid of a risk.

The Insurance Premium

Some risks are unavoidable: buying a car means facing the risk of an accident and costly repairs. We define the insurance premium as the maximum amount an agent is willing to pay to eliminate such a risk.

Consider an asset worth W. Its owner can either bear the risk and end up with $W + \tilde{\varepsilon}$, or pay Π_i to insure it and keep a certain $W - \Pi_i$. She is indifferent between the two when

$$u(W - \Pi_i) = \mathrm{E}(u(W + \tilde{\varepsilon})). \tag{2}$$

The value of Π_i that solves (2) is the **insurance premium**, and the certain wealth $W - \Pi_i$ she accepts in place of the gamble is its **certainty equivalent**. In Figure 1, the certainty equivalent is found by reading $\mathrm{E}(u(\tilde{W}))$ back across to the utility curve, and the insurance premium is the horizontal distance between CE and W.

Premiums rank risk aversion only across agents facing the *same* gamble at the *same* wealth: of two such agents, the one willing to pay more for insurance is **more risk averse**. For a single agent Π_i varies with W, so her premiums at different wealth levels say nothing about a change in her risk aversion.

---

the gamble $\tilde{\varepsilon}$ that pays $(1-t)d$ with probability t and $-td$ with probability $1-t$. It is fair, since $\mathrm{E}(\tilde{\varepsilon}) = 0$, and it moves wealth to exactly x_1 or x_2, because $W + (1-t)d = x_1$ and $W - td = x_2$. A risk-averse investor dislikes it, so

$$\begin{aligned} u(W) &> \mathrm{E}(u(W + \tilde{\varepsilon})) \\ &= t\,u(x_1) + (1-t)\,u(x_2), \end{aligned}$$

which, given the definition of W, is exactly strict concavity.

Example 1. Consider an investor with logarithmic utility $u(W) = \ln W$ and current wealth $W = \$10{,}000$. There is a 10% probability that a loss event reduces wealth to $\$1{,}000$, and a 90% probability that wealth remains at $\$10{,}000$. Two insurance policies are available:

- **Plan A**: pays $\$5{,}000$ if the loss event occurs and costs $\$500$ upfront.
- **Plan B**: pays $\$8{,}000$ if the loss event occurs and costs $\$800$ upfront.

Should the investor buy Plan A, Plan B, or no insurance?

Note first that both policies are actuarially fair: the premium equals the expected payout in each case ($0.1 \times \$5{,}000 = \$500$ and $0.1 \times \$8{,}000 = \$800$), so all three alternatives leave expected wealth unchanged at

$$\mathrm{E}(\tilde{W}) = 0.9 \times \$10{,}000 + 0.1 \times \$1{,}000 = \$9{,}100.$$

Since expected wealth is the same in every case, a risk-averse investor's ranking depends only on how much dispersion remains across states, which by Jensen's inequality she wants to minimize.

Without insurance, wealth is $\$1{,}000$ in the bad state and $\$10{,}000$ in the good state, so

$$\mathrm{E}(u) = 0.1 \ln(1{,}000) + 0.9 \ln(10{,}000) = 8.9801.$$

Under Plan A, wealth net of the premium is $\$10{,}000 - \$500 = \$9{,}500$ in the good state and $\$1{,}000 - \$500 + \$5{,}000 = \$5{,}500$ in the bad state, so

$$\mathrm{E}(u) = 0.1 \ln(5{,}500) + 0.9 \ln(9{,}500) = 9.1044.$$

Under Plan B, wealth is $\$10{,}000 - \$800 = \$9{,}200$ in the good state and $\$1{,}000 - \$800 + \$8{,}000 = \$8{,}200$ in the bad state, so

$$\mathrm{E}(u) = 0.1 \ln(8{,}200) + 0.9 \ln(9{,}200) = 9.1155.$$

Translating each expected utility into a certainty equivalent, $CE = \exp(\mathrm{E}(u))$, gives

$$CE_{\text{none}} = \$7{,}943, \qquad CE_A = \$8{,}994, \qquad CE_B = \$9{,}095.$$

Since $CE_B > CE_A > CE_{\text{none}}$, the investor should choose **Plan B**. Intuitively, all three alternatives are fairly priced and share the same expected wealth, so the only thing that differs is how much each one compresses the gap between good- and bad-state wealth: no insurance leaves a $\$9{,}000$ gap, Plan A leaves a $\$4{,}000$ gap, and Plan B leaves only a $\$1{,}000$ gap. Because u is strictly concave, a risk-averse investor always prefers the fairly-priced contract that reduces this spread the most, which here is Plan B; full insurance, eliminating the gap entirely, would be even better if offered at the same fair price. □

Example 2. Suppose the economy will be in one of two states at the end of the period: a **boom**, or a **recession**, each equally likely. A risky asset would be worth $50 in the boom and $10 in the recession, and it currently trades at $30. Two investors, A and B, are each sizing up whether it's worth buying. To keep the example simple, assume this asset is each investor's only source of wealth at the end of the period, so that end-of-period wealth W coincides with the asset's payoff and we can evaluate $u(W)$ directly at $10 and $50.

Investor A has the utility function

$$u(W) = 10 \ln(W),$$

while investor B has the linear utility function

$$u(W) = 2W + 5,$$

which makes B risk-neutral rather than risk-averse, a useful foil for A.

How much would each be willing to pay? The answer is each investor's certainty equivalent: the highest price that leaves them no worse off, in expected-utility terms, than passing on the asset.

For investor A, the expected utility of holding the asset is

$$\mathrm{E}(u(\tilde{W})) = 0.5(10 \ln(10)) + 0.5(10 \ln(50)) = 31.0730.$$

Setting this equal to the utility of a sure payoff and solving pins down A's certainty equivalent:

$$\begin{aligned} 10\ln(CE) &= 31.0730 \\ \ln(CE) &= 3.1073 \\ CE &= \exp(3.1073) = \$22.36. \end{aligned}$$

So A values the asset at only \$22.36, well short of its \$30 market price. Even though the asset pays \$30 on average, the prospect of ending up with just \$10 in a recession weighs heavily enough that A would rather not hold it.

Investor B sees things differently. Her expected utility is

$$\mathrm{E}(u(\tilde{W})) = 0.5(2 \times 10 + 5) + 0.5(2 \times 50 + 5) = 65,$$

and solving

$$\begin{aligned} 2 \times CE + 5 &= 65 \\ 2 \times CE &= 60 \\ CE &= \$30 \end{aligned}$$

gives $CE = \$30$, so B is exactly indifferent at the market price. That's no coincidence: with linear utility, B cares only about expected payoffs, so her certainty equivalent must equal the asset's expected value, $0.5 \times 10 + 0.5 \times 50 = \$30$.

The \$7.64 gap between the two certainty equivalents $(30 - 22.36)$ is the price of risk itself: the insurance premium A would pay to shed this gamble, which vanishes entirely for the risk-neutral B. □

Local Risk Aversion

A more concave utility function should induce more risk aversion. We can formalize this intuition by looking at the insurance premium for a gamble with a very small variance.

Let $\tilde{\varepsilon}$ be a fair gamble with $\mathrm{V}(\tilde{\varepsilon}) = \sigma_\varepsilon^2$, and expand both sides of (2) around W. To first order on the left,

$$u(W - \Pi_i) \approx u(W) - u'(W)\Pi_i, \tag{3}$$

and to second order on the right, using $\mathrm{E}(\tilde{\varepsilon}) = 0$ and $\mathrm{E}(\tilde{\varepsilon}^2) = \sigma_\varepsilon^2$,

$$\mathrm{E}(u(W + \tilde{\varepsilon})) \approx u(W) + u'(W)\,\mathrm{E}(\tilde{\varepsilon}) + \frac{1}{2}u''(W)\,\mathrm{E}(\tilde{\varepsilon}^2) = u(W) + \frac{1}{2}u''(W)\sigma_\varepsilon^2. \quad (4)$$

Equating (3) and (4) gives

$$\Pi_i \approx \frac{1}{2}ARA \cdot \sigma_\varepsilon^2, \qquad ARA = -\frac{u''(W)}{u'(W)}, \quad (5)$$

where *ARA* is the Arrow-Pratt **coefficient of absolute risk-aversion** (Ingersoll 1987, sec. 1.12). It is literally the dollar premium the agent pays per unit of variance, so an agent with a higher *ARA* pays more to shed any given gamble. Since *ARA* depends on W, we can also ask how risk-aversion changes with wealth: an agent with decreasing absolute risk-aversion (DARA) pays a smaller dollar premium as she gets wealthier, the pattern favored by most empirical evidence and introspection.

The factor of W in the second coefficient appears once we look at proportional rather than dollar gambles. For $\tilde{\varepsilon} = W\tilde{\delta}$, where $\tilde{\delta}$ is a percentage shock with $\mathrm{E}(\tilde{\delta}) = 0$ and $\mathrm{V}(\tilde{\delta}) = \sigma_\delta^2$, substituting into (5) gives

$$\Pi_i \approx \frac{1}{2}ARA \cdot W^2\sigma_\delta^2 = \frac{1}{2}RRA \cdot W\sigma_\delta^2, \qquad RRA = -\frac{u''(W)}{u'(W)}W,$$

where *RRA* is the **coefficient of relative risk-aversion**. Dividing by W, the fraction of wealth the agent sacrifices to eliminate the proportional gamble is $\Pi_i/W \approx \frac{1}{2}RRA \cdot \sigma_\delta^2$. This is why an agent with constant relative risk-aversion (CRRA) keeps the same fraction of wealth in risky assets regardless of how wealthy she is, a portfolio-choice result we state here for intuition and derive formally in a later note.

Example 3. Take

$$u(W) = \frac{W^{1-\gamma} - 1}{1 - \gamma}$$

Then $u'(W) = W^{-\gamma}$ and $u''(W) = -\gamma W^{-\gamma-1}$, implying that

$$RRA = -\left(\frac{-\gamma W^{-\gamma-1}}{W^{-\gamma}}\right)W = \gamma$$

Power utility is an example of a function that exhibits constant relative risk-aversion. □

Example 4. For exponential utility $u(W) = -e^{-aW}$ with $a > 0$, we have $u'(W) = ae^{-aW}$ and $u''(W) = -a^2 e^{-aW}$, so

$$ARA = \frac{a^2 e^{-aW}}{ae^{-aW}} = a.$$

The coefficient of absolute risk-aversion is constant (equal to a) and independent of W. Exponential utility therefore exhibits **constant absolute risk-aversion** (CARA). A direct implication is that the insurance premium $\Pi_i \approx \frac{1}{2}a\sigma_\varepsilon^2$ does not depend on the agent's wealth: a billionaire and a middle-class worker with the same CARA parameter a would pay identical dollar amounts to insure against the same gamble. This is analytically convenient but empirically implausible, since most people become less risk-averse in absolute terms as they accumulate wealth. □

Figure 2 puts these two examples side by side. The power-utility agent's ARA falls as wealth grows (DARA), while the exponential-utility agent's ARA stays flat at a no matter how wealthy she becomes (CARA). At low wealth the CRRA agent can be far more risk-averse than the CARA agent; at high wealth the ordering flips, since the CRRA agent's ARA keeps shrinking while the CARA agent's does not.

The Risk Premium

The insurance premium asks how much an agent would pay to shed a gamble she already bears. The mirror-image question is how much she must be paid to accept one she does not bear. The answer is the **compensatory risk premium** Π_c, the amount that, added to the fair gamble, leaves her indifferent between the certain wealth W and the sweetened gamble:

$$u(W) = \mathrm{E}\left(u(W + \Pi_c + \tilde{\varepsilon})\right). \tag{6}$$

Although Π_i and Π_c measure aversion to the same gamble at the same wealth, they are generally *not* equal, since they price risk from opposite sides. The insurance premium is the natural quantity for casualty or liability insurance; the compensatory premium is the natural one in

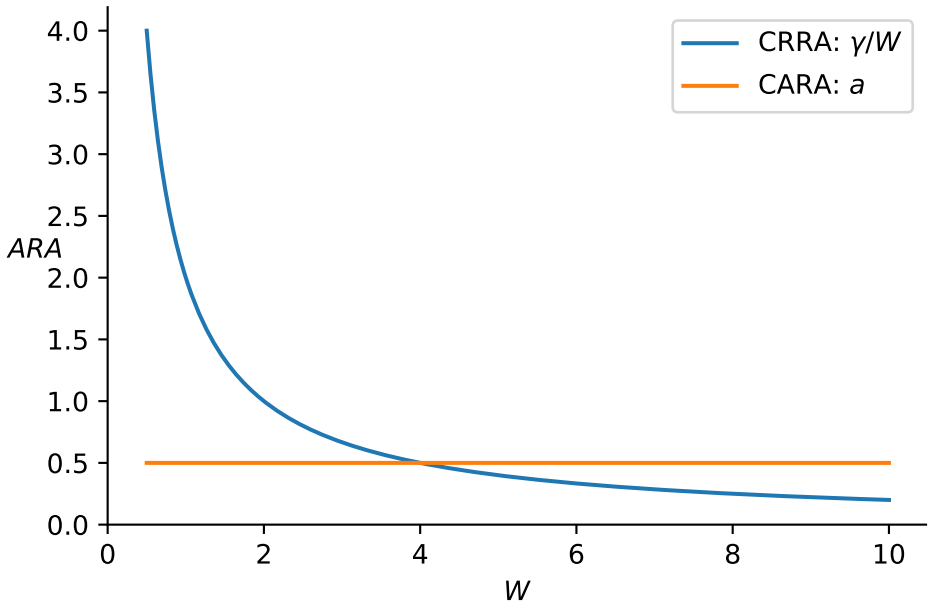


Figure 2: Absolute risk-aversion for CRRA (power) versus CARA (exponential) utility.

finance, since it corresponds to the extra expected return an investor must be offered to hold a risky asset instead of a riskless one.

For a small gamble, however, the two nearly coincide. The same second-order Taylor argument applied to (6) gives

$$\Pi_c \approx \frac{1}{2} ARA \cdot \mathrm{V}(\tilde{\varepsilon}),$$

which is exactly the approximation for Π_i in (5). Conceptually distinct, the two premiums are governed by the same local curvature of u at W.

We can connect this back to Example 2 by taking the asset's expected payoff, $\$30$, as the reference wealth W and its $\pm\$20$ deviations as the fair gamble $\tilde{\varepsilon}$. Investor A's willingness to pay only $\$22.36$ for the asset says precisely that she would demand an insurance premium of $\Pi_i = 30 - 22.36 = \$7.64$ to shed this risk, while the risk-neutral investor B has $\Pi_i = \$0$.

Practice Problems

These problems give you a chance to practice the concepts introduced in this chapter. Try to solve each one on your own before expanding the solution.

Solutions to all problems can be found at https://lorenzonaranjo.com/class-materials/investments/.

Problem 1 (The Insurance Decision Across Wealth Levels)**.** Clara Bennett, a graduate student, has a logarithmic utility function $u(W) = \ln W$, and her current wealth is $\$5{,}000$.

a. Clara faces a 50/50 chance of either winning or losing $\$1{,}000$. She can buy insurance for $\$125$ that would entirely eliminate this risk. Should she buy the insurance or take the gamble?
b. Suppose Clara took the gamble in part (a) and lost, which reduced her wealth to $\$4{,}000$. If she is now offered the same gamble and the same insurance for $\$125$, should she buy the insurance this time?

Problem 2 (Maximum Willingness to Pay for Insurance)**.** Ines Kowalski owns Kowalski Woodworks, a furniture workshop that makes up nearly all of her net worth. She faces the following risks: a 10% chance that a fire will reduce her net worth to $\$1$, a 10% chance that a fire will lower it to $\$50{,}000$, and an 80% chance that nothing will happen, maintaining the business's value at $\$100{,}000$. Her utility function is logarithmic, given by $u(W) = \ln W$. She is considering an insurance policy that would provide a payout of $\$99{,}999$ in the first scenario, $\$50{,}000$ in the second, and nothing in the third. What is the maximum amount she would be willing to pay for this insurance policy?

Problem 3 (A Fair Lottery and Concavity)**.** You are offered the possibility to participate at the following lottery:

Gain	Probability
2	0.5
0	0.5

The cost of participating at the lottery is 1 unit of consumption. If you choose not to participate, you keep your unit of consumption.

a. Is this a fair gamble?
b. Show that the decision not to participate at the gamble implies that your utility function is concave.

Problem 4 (Mean-Preserving Spreads and Risk Aversion)**.** Consider gamble A:

Gain	Probability
-2	0.09
4	0.30
10	0.40
16	0.21

From A, we can construct another gamble B by adding white noise to a number of outcomes. Indeed, we can replace the outcome 4 by the gamble A':

- 3 with probability 1/2
- 5 with probability 1/2

with $\mathrm{E}(A') = 4$. In the same manner, we can replace outcome 16 with gamble A'':

- 12 with probability 1/3
- 18 with probability 2/3

with $\mathrm{E}(A'') = 16$.

Show formally why any risk averse individual prefers gamble A to gamble B.

Problem 5 (Divergent Beliefs and the Incentive to Litigate)**.** A distant relative in Europe has recently passed away, leaving behind an estimated fortune of $\$1{,}000{,}000$. This has left two grieving but competing close relatives: Peter, who currently has no wealth, and Paul, who has

$\$10{,}000$. With the will missing, they can pursue legal action, but the winner will incur legal costs amounting to 10% of the inheritance. Both Peter and Paul share the same utility function:

$$U(W) = W^{1/2}.$$

Peter and Paul both agree that Peter has a 60% chance of winning the million, while Paul has a 40% chance. The judge cannot issue a split decision; the entire amount must go to one of them.

a. Should Peter and Paul take their dispute to court, or is there a mutually beneficial agreement they could reach instead?
b. Does the conclusion change if the heirs disagree on the probabilities? For instance, what if Peter believes he has an 80% chance of winning, while Paul thinks he has an 80% chance of winning?

Problem 6 (Quadratic Utility and Increasing Risk Aversion)**.** Consider an investor with quadratic utility

$$u(W) = W - \frac{b}{2}W^2, \qquad b > 0, \quad W < \frac{1}{b},$$

where the restriction $W < 1/b$ (the **satiation point**) is needed to keep $u'(W) > 0$.

a. Derive general expressions for *ARA*(W) and *RRA*(W).
b. Take $b = 0.0001$, so satiation occurs at $W = \$10{,}000$. Compute *ARA* and *RRA* at $W = \$1{,}000$ and again at $W = \$4{,}000$.
c. Is quadratic utility DARA or IARA? Contrast your answer with the log-utility investor of Exercise 1, whose willingness to insure the same $\$1{,}000$ gamble *increased* once her wealth fell from $\$5{,}000$ to $\$4{,}000$.

Problem 7 (How Good Is the Local Approximation?)**.** Marco Delgado has log utility $u(W) = \ln W$ and wealth $W = \$100{,}000$. His employer offers him a bonus scheme that adds or subtracts $\$10{,}000$ from his pay with equal probability.

a. Compute the exact insurance premium Marco would pay to eliminate this risk.
b. Compute the approximation in (5) and compare it with the exact premium.

Problem 8 (Premiums as a Fraction of Wealth)**.** Nadia Haddad has power utility with $\gamma = 2$ and wealth of $\$500{,}000$. Owen Fitzgerald has power utility with $\gamma = 5$ and wealth of $\$50{,}000$. Each faces a gamble that raises or lowers wealth by 10% with equal probability.

a. Use the local approximation to find the fraction of wealth, and the dollar amount, each would pay to eliminate the gamble.
b. Nadia would pay four times as many dollars as Owen. Does that make her the more risk-averse of the two?
c. Compare their coefficients of absolute and relative risk aversion.

Problem 9 (Compensation for Taking a Risk)**.** A startup asks Yusuf Adeyemi to invest in a venture whose outcome will add or subtract $\$20{,}000$ from his wealth with equal probability. Yusuf has exponential utility $u(W) = -e^{-aW}$ with $a = 0.00005$.

a. Use (6) to find the compensatory risk premium, the expected gain the startup must add to the venture for Yusuf to accept it.
b. Compare the exact premium with the local approximation.
c. Show that, for exponential utility, the answer does not depend on Yusuf's wealth, and that the insurance premium for the same gamble is exactly the same.

Ingersoll, Jonathan E. 1987. *Theory of Financial Decision Making*. Vol. 3. Rowman & Littlefield.