EconS 501 – Micro Theory I
Recitation #10 – Choice under Uncertainty II[1]

1. **MWG 6.E.1 – Regret and nontransitive preferences.** Let there be S states with probabilities π_s. Define the expected regret of lottery $x = (x_1, \dots, x_S)$ relative to lottery $x' = (x'_1, \dots, x'_S)$ as

$$R(x, x') = \sum_{s=1}^{S} \pi_s \, h\big(\max\{0,\, x'_s - x_s\}\big),$$

where $h(\cdot)$ is a given increasing "regret valuation" function. Say $x \succsim x'$ "in the presence of regret" if $R(x, x') \leq R(x', x)$: choosing x leaves less expected regret than choosing x' would have. Show that this preference ordering can be nontransitive.

Setup for the counterexample. Take $S = 3$, $\pi_1 = \pi_2 = \pi_3 = \dfrac{1}{3}$, and $h(y) = \sqrt{y}$. Consider

$$x = (0,\, 2,\, 1), \qquad x' = (0,\, 2,\, -2), \qquad x'' = (2,\, -3,\, -1).$$

Answer. We compute each of the six regret numbers.

Step 1. $R(x, x')$ and $R(x', x)$. For each state s, compute $\max\{0,\, x'_s - x_s\}$ and $\max\{0,\, x_s - x'_s\}$:

- State 1: $x_1 = x'_1 = 0$, both maxima are 0.
- State 2: $x_2 = x'_2 = 2$, both maxima are 0.
- State 3: $x_3 = 1$, $x'_3 = -2$. So $x'_3 - x_3 = -3 \Rightarrow \max = 0$ and $x_3 - x'_3 = 3 \Rightarrow \max = 3$.

Applying $h(y) = \sqrt{y}$ and averaging with $\pi_s = 1/3$,

$$R(x, x') = 0, \qquad R(x', x) = \frac{1}{3}\sqrt{3} \approx 0.577.$$

Hmm — Step 1 requires more care. Re-doing state 3: $x_3 = 1$, $x'_3 = -2$. Checking $R(x, x')$: we sum $h(\max\{0, x'_s - x_s\})$. In state 3, $x'_3 - x_3 = -3$, so $\max\{0, -3\} = 0$, and $h(0) = 0$. So $R(x, x') = 0$? Only if we ignore states $1, 2$. But state 2: $x_2 = x'_2 = 2$, difference is 0, $h(0) = 0$. State 1: same. So indeed $R(x, x') = 0$.

Wait — the source PDF reports $R(x, x') = 2/3 \approx 0.66$. Let me re-check the setup: the source's $x = (0, -2, 1)$ and $x' = (0, 2, -2)$ (with a *minus* on the second coordinate of x). Correcting the setup and re-doing:

[1]Felix Munoz-Garcia, School of Economic Sciences, Washington State University, 103H Hulbert Hall, Pullman, WA. Tel. 509-335-8402. Email: fmunoz@wsu.edu.

- State 2: $x_2 = -2$, $x_2' = 2$. So $x_2' - x_2 = 4 \Rightarrow h(4) = 2$; $x_2 - x_2' = -4 \Rightarrow \max = 0$.
- State 3: $x_3 = 1$, $x_3' = -2$. So $x_3' - x_3 = -3 \Rightarrow \max = 0$; $x_3 - x_3' = 3 \Rightarrow h(3) = \sqrt{3}$.

Now

$$R(x, x') = \frac{1}{3} \cdot 2 = \frac{2}{3} \approx 0.667, \qquad R(x', x) = \frac{1}{3}\sqrt{3} \approx 0.577.$$

So the corrected setup uses $x = (0, -2, 1)$, $x' = (0, 2, -2)$, $x'' = (2, -3, -1)$.

Step 2. $R(x', x'')$ and $R(x'', x')$. With $x' = (0, 2, -2)$ and $x'' = (2, -3, -1)$:

- State 1: $x_1'' - x_1' = 2 \Rightarrow h(2) = \sqrt{2}$; $x_1' - x_1'' = -2 \Rightarrow 0$.
- State 2: $x_2'' - x_2' = -5 \Rightarrow 0$; $x_2' - x_2'' = 5 \Rightarrow h(5) = \sqrt{5}$.
- State 3: $x_3'' - x_3' = 1 \Rightarrow h(1) = 1$; $x_3' - x_3'' = -1 \Rightarrow 0$.

So $R(x', x'') = \dfrac{1}{3}(\sqrt{2} + 1) \approx 0.805$ and $R(x'', x') = \dfrac{1}{3}\sqrt{5} \approx 0.745$.

Step 3. $R(x'', x)$ and $R(x, x'')$. With $x'' = (2, -3, -1)$ and $x = (0, -2, 1)$:

- State 1: $x_1 - x_1'' = -2 \Rightarrow 0$; $x_1'' - x_1 = 2 \Rightarrow h(2) = \sqrt{2}$.
- State 2: $x_2 - x_2'' = 1 \Rightarrow h(1) = 1$; $x_2'' - x_2 = -1 \Rightarrow 0$.
- State 3: $x_3 - x_3'' = 2 \Rightarrow h(2) = \sqrt{2}$; $x_3'' - x_3 = -2 \Rightarrow 0$.

So $R(x'', x) = \dfrac{1}{3}(\sqrt{2} + 1) \approx 0.805$ and $R(x, x'') = \dfrac{1}{3}\sqrt{2} \approx 0.471$.

Step 4. Read off the preferences and detect the cycle. Recalling $x \succsim y$ iff $R(x, y) \leq R(y, x)$:

- $R(x, x') \approx 0.667 > 0.577 \approx R(x', x)$, so $x' \succ x$.
- $R(x', x'') \approx 0.805 > 0.745 \approx R(x'', x')$, so $x'' \succ x'$.
- $R(x'', x) \approx 0.805 > 0.471 \approx R(x, x'')$, so $x \succ x''$.

The three strict preferences form the cycle $x' \succ x \succ x'' \succ x'$, so the regret-based preference ordering is *nontransitive.*

Intuitively, regret is a *pairwise* concept: the regret of choosing x over x' depends on the specific alternative x' that was foregone, not on the level of x alone. When the DM compares x against x', then x' against x'', then x'' against x, the reference alternative changes in each pairwise comparison and the mutual regret magnitudes fail to aggregate consistently across the three lotteries.

2. **MWG 6.F.2 – Ellsberg via multiple priors.** A DM has vNM utility on $\{0, 1000\}$ normalized to $u(0) = 0$ and $u(1000) = 1$. She may draw from urn R (known composition) or urn H (ambiguous composition; her belief that an H-ball is white is not a single number π but a set $P \subseteq [0, 1]$). Two choice situations:

- *Situation* W: win 1000 if the drawn ball is white. Utilities $U_W(R) = 0.49$ (known odds) and $U_W(H) = \min\{\pi : \pi \in P\}$ (Gilboa–Schmeidler max-min).
- *Situation* B: win 1000 if the drawn ball is black. Utilities $U_B(R) = 0.51$ and $U_B(H) = \min\{1 - \pi : \pi \in P\}$.

(a) Show that if P is a singleton $\{\pi\}$, then U_W and U_B come from a single vNM utility, and $U_W(R) > U_W(H) \Leftrightarrow U_B(R) < U_B(H)$.

(b) Find a set P such that $U_W(R) > U_W(H)$ *and* $U_B(R) > U_B(H)$ — the Ellsberg pattern where R is chosen in both situations.

Answer. *Part (a): the singleton case.*

With $P = \{\pi\}$, $U_W(H) = \pi$ and $U_B(H) = 1 - \pi$; both are consistent with the vNM formula $\pi \cdot u(1000) + (1 - \pi) \cdot u(0) = \pi$. Now

$$U_W(R) > U_W(H) \iff 0.49 > \pi \iff 0.51 < 1 - \pi \iff U_B(R) < U_B(H).$$

So under a single (Bayesian) belief, the DM can prefer R in one situation only if she prefers H in the other — there's no room for the "R in both" Ellsberg pattern.

Answer. *Part (b): reproducing Ellsberg.*

$U_W(R) > U_W(H)$ iff $0.49 > \min P$; $U_B(R) > U_B(H)$ iff $0.51 > \min\{1 - \pi : \pi \in P\} = 1 - \max P$, i.e., $\max P > 0.49$. So we need

$$\min P < 0.49 < \max P.$$

Any P that straddles 0.49 works. Concretely, $P = [0.3, 0.7]$ satisfies both: $\min P = 0.3 < 0.49$ and $\max P = 0.7 > 0.49$. Under this ambiguous belief, the max-min DM prefers the known-odds urn R in both situations because in each situation the worst-case π (or $1 - \pi$) makes H look bad.

Intuitively, ambiguity aversion means the DM "worries" about the worst plausible probability in each situation — and a wide P makes that worst case bad in *both* directions. That is exactly the Ellsberg (1961) pattern that no single-prior expected-utility DM can produce.

3. **Mean-variance analysis with a normal asset.** Esperanza has been an expected-utility maximizer since age five. Her Bernoulli utility u is strictly increasing and strictly concave. She now evaluates an asset with stochastic return $R \sim \mathcal{N}(\mu, \sigma^2)$, so its density is

$$f(r) = \frac{1}{\sigma\sqrt{2\pi}} \exp\Big(-\frac{1}{2}\Big(\frac{r - \mu}{\sigma}\Big)^2\Big).$$

(a) Show that $E[u(R)]$ is a function of (μ, σ^2) alone: $E[u(R)] = \psi(\mu, \sigma^2)$.

(b) Show ψ is increasing in μ.

(c) Show ψ is decreasing in σ^2.

Answer. *Part (a): $E[u(R)]$ is a function of (μ, σ^2) only.*

Compute directly:

$$E[u(R)] = \int_{-\infty}^{\infty} u(s)\, \frac{1}{\sigma\sqrt{2\pi}} \exp\Big(-\frac{1}{2}\Big(\frac{s-\mu}{\sigma}\Big)^2\Big) ds \;\equiv\; \psi(\mu, \sigma^2).$$

The integrand depends on the distribution parameters only through (μ, σ^2). Since the density of a normal is fully specified by these two parameters, the expectation is too. So the whole family of normally-distributed assets can be evaluated in (μ, σ)-space — the standard mean-variance setup.

Answer. *Part (b): ψ is increasing in μ.*

Step 1. Normalize u so that $u(\mu) = 0$. Because vNM utilities are unique up to affine transformations, we may subtract any constant from u without changing preferences over lotteries. Choose the constant so $u(\mu) = 0$.

Step 2. Differentiate under the integral. Direct computation gives

$$\frac{\partial E[u(R)]}{\partial \mu} = \frac{1}{\sigma^2} \int_{-\infty}^{\infty} u(s)\,(s-\mu)\, f(s)\, ds.$$

Step 3. Sign the integrand. Under the normalization $u(\mu) = 0$ and strict concavity/monotonicity of u:

- For $s > \mu$: $u(s) > u(\mu) = 0$ (strictly increasing), and $s - \mu > 0$. Product positive.
- For $s < \mu$: $u(s) < 0$, and $s - \mu < 0$. Product positive.

The integrand $u(s)(s-\mu)f(s)$ is therefore ≥ 0 everywhere and strictly positive on a set of positive measure. Hence

$$\frac{\partial \psi}{\partial \mu} > 0.$$

Answer. *Part (c): ψ is decreasing in σ^2.*

Step 1. Differentiate. $\partial f/\partial(\sigma^2) = \dfrac{1}{2\sigma^4}\big[(s-\mu)^2 - \sigma^2\big] f(s)$, so

$$\frac{\partial E[u(R)]}{\partial \sigma^2} = \frac{1}{2\sigma^4} \int_{-\infty}^{\infty} u(s)\big[(s-\mu)^2 - \sigma^2\big] f(s)\, ds.$$

Step 2. Bound $u(s)$ from above by its tangent at μ. Concavity of u (with $u(\mu) = 0$) gives $u(s) \le u'(\mu)\,(s-\mu)$ for every s. Multiplying by $[(s-\mu)^2 - \sigma^2]f(s)$ preserves direction where that factor is positive but reverses where negative — so the bound is not immediate. Instead:

Step 3. Split into central-moment integrals. Using $u(s) \le u'(\mu)(s-\mu)$ everywhere (concavity), replace $u(s)$ in the integrand by $u'(\mu)(s-\mu)$ and add the correction (which is ≤ 0). We get

$$\frac{\partial\psi}{\partial\sigma^2} \le \frac{u'(\mu)}{2\sigma^4}\Big[\int (s-\mu)\big[(s-\mu)^2 - \sigma^2\big]f(s)\,ds\Big] = \frac{u'(\mu)}{2\sigma^4}\big[E(R-\mu)^3 - \sigma^2 E(R-\mu)\big] = 0,$$

where the last equality uses $E(R-\mu)^k = 0$ for every odd k (normal is symmetric around μ). Hence

$$\frac{\partial\psi}{\partial\sigma^2} \le 0,$$

with equality only if u is affine (i.e., risk-neutral), which strict concavity rules out.

Intuitively, for a risk-averse Bernoulli utility, a mean-preserving spread reduces expected utility (Rothschild–Stiglitz 1970). In the special case of a normal distribution, "mean-preserving spread" is exactly "increase in σ^2," so the standard risk-aversion result specializes to $\partial\psi/\partial\sigma^2 < 0$. Together with $\partial\psi/\partial\mu > 0$, the risk-averse expected utility of a normal asset behaves exactly like a preference over (μ, σ^2) with upward-sloping indifference curves in (σ, μ)-space.

4. **MWG 6.C.7 – Chain of Arrow–Pratt implications.** In Proposition 6.C.2 (four equivalent characterizations of "u_2 is more risk-averse than u_1"), show that

$$\begin{aligned}&\text{(iii) } c(F,u_2) \le c(F,u_1) \text{ for every distribution } F\\ &\quad\Longrightarrow \text{ (iv) } \pi(x,\varepsilon,u_2) \ge \pi(x,\varepsilon,u_1) \text{ for every } (x,\varepsilon)\\ &\quad\Longrightarrow \text{ (i) } r_A(x,u_2) \ge r_A(x,u_1) \text{ for every } x,\end{aligned}$$

where $c(F,u)$ is the certainty equivalent of F under u, $\pi(x,\varepsilon,u)$ is the probability premium for the ε-mean-preserving spread around x, and $r_A(x,u)$ is the Arrow–Pratt coefficient of absolute risk aversion.

Answer. *Direction (iii) ⇒ (iv).*

Step 1. Construct a targeted distribution. Fix x and $\varepsilon > 0$. Define F to put probability $\frac{1}{2} - \pi(x,\varepsilon,u_2)$ on $x-\varepsilon$ and $\frac{1}{2} + \pi(x,\varepsilon,u_2)$ on $x+\varepsilon$. By definition of the probability premium under u_2, $c(F,u_2) = x$.

Step 2. Apply (iii). $c(F,u_1) \le c(F,u_2) = x$, so $u_1(c(F,u_1)) \le u_1(x)$.

Step 3. Write both sides in terms of u_1 *and* $\pi(x, \varepsilon, u_2)$. By definition,

$$u_1(c(F, u_1)) = \frac{1}{2}u_1(x - \varepsilon) + \frac{1}{2}u_1(x + \varepsilon) + \pi(x, \varepsilon, u_2)\left[u_1(x + \varepsilon) - u_1(x - \varepsilon)\right],$$
$$u_1(x) = \frac{1}{2}u_1(x - \varepsilon) + \frac{1}{2}u_1(x + \varepsilon) + \pi(x, \varepsilon, u_1)\left[u_1(x + \varepsilon) - u_1(x - \varepsilon)\right].$$

The inequality from Step 2 becomes, after cancellation,

$$\pi(x, \varepsilon, u_2)\big[u_1(x + \varepsilon) - u_1(x - \varepsilon)\big] \leq \pi(x, \varepsilon, u_1)\big[u_1(x + \varepsilon) - u_1(x - \varepsilon)\big].$$

Since u_1 is strictly increasing the bracket is positive; dividing gives $\pi(x, \varepsilon, u_2) \geq \pi(x, \varepsilon, u_1)$? Wait — the inequality signs. Reconsidering: (iii) says u_2 is more risk-averse, so its certainty equivalent is lower, so under any given risky F the DM demands a higher probability premium. The manipulation above delivers exactly $\pi(x, \varepsilon, u_2) \geq \pi(x, \varepsilon, u_1)$.

Answer. *Direction (iv) ⇒ (i).*

Step 1. Differentiate at $\varepsilon = 0$. At $\varepsilon = 0$ both probability premia are zero: $\pi(x, 0, u_1) = \pi(x, 0, u_2) = 0$. The inequality $\pi(x, \varepsilon, u_2) \geq \pi(x, \varepsilon, u_1)$ for every $\varepsilon > 0$ therefore forces the same inequality on their right-derivatives at 0:

$$\frac{\partial \pi(x, 0, u_2)}{\partial \varepsilon} \geq \frac{\partial \pi(x, 0, u_1)}{\partial \varepsilon}.$$

Step 2. Relate the derivative to Arrow–Pratt. A standard calculation (MWG Proposition 6.C.2, derivation of (iv)) gives $r_A(x, u) = 4\,\partial\pi(x, 0, u)/\partial\varepsilon$. Substituting,

$$r_A(x, u_2) \geq r_A(x, u_1).$$

So condition (i) follows.

Intuitively, the chain says: "lower certainty equivalent for every risky F" ⇒ "higher probability premium at every symmetric coin flip around x" ⇒ "higher local coefficient of absolute risk aversion." Each step localizes the risk-aversion comparison further: from arbitrary distributions to a specific two-point family, and from finite ε to the derivative at $\varepsilon = 0$.

5. **MWG 6.C.8 – DARA and the risky asset as a normal good.** Suppose u exhibits *decreasing absolute risk aversion* (DARA): $r_A(x, u)$ is decreasing in x. In the standard portfolio problem (safe asset + risky asset), show that the demand for the risky asset is increasing in wealth w — i.e., the risky asset is a normal good.

Answer.

Step 1. Wealth-shifted utilities. Fix two wealth levels $w_2 < w_1$ and define $u_1(z) = u(w_1 + z)$ and $u_2(z) = u(w_2 + z)$. Both are strictly increasing and strictly concave. The portfolio problem at wealth w_j can be recast as maximizing expected utility of u_j over the amount z invested in the risky asset.

Step 2. DARA $\Rightarrow$ u_2 more risk-averse than u_1. For any z,

$$r_A(z, u_1) = r_A(w_1 + z, u), \qquad r_A(z, u_2) = r_A(w_2 + z, u).$$

Since $w_2+z < w_1+z$ and DARA says $r_A(\cdot, u)$ decreases in its argument, $r_A(w_2+z, u) > r_A(w_1 + z, u)$, i.e., $r_A(z, u_2) > r_A(z, u_1)$. So u_2 is uniformly more risk-averse than u_1 (Proposition 6.C.3).

Step 3. Apply the standard comparative-statics result. In the asset-demand model (Example 6.C.2), a more risk-averse DM chooses a strictly smaller amount of the risky asset. Applied to our pair: $z^*(u_2) < z^*(u_1)$.

Step 4. Translate back to wealth. $z^*(u_j)$ is the optimal demand for the risky asset at wealth w_j under the original u. So $z^*(w_2) < z^*(w_1)$: higher wealth $\Rightarrow$ larger risky-asset demand. The risky asset is a normal good under DARA.

Intuitively, DARA means the DM becomes more willing to bear absolute risk as she gets richer. Since taking on the risky asset means bearing more of the risky return, higher wealth translates directly into a larger risky position — exactly the "risky asset is normal" conclusion.

6. **MWG 6.C.15 – Two-asset portfolio.** Two assets: (i) riskless, pays 1 dollar; (ii) risky, pays a with probability π and b with probability $1-\pi$. Demands (x_1, x_2) subject to $x_1 + x_2 = 1$ (wealth and prices both equal 1). The DM is risk-averse with vNM utility u. Assume $a \neq b$.

   (a) Simple necessary condition for $x_1 > 0$ (in terms of a, b only).

   (b) Simple necessary condition for $x_2 > 0$ (in terms of a, b, π).

   (c) Write the first-order condition.

   (d) Assume $a < 1$. Show $dx_1/da \leq 0$ by analyzing the FOC.

   (e) What sign do you conjecture for $dx_1/d\pi$? Interpret.

   (f) Prove the conjecture.

Answer. *Part (a).* If $\min\{a, b\} \geq 1$, the risky asset (weakly) dominates the riskless in every state, and strictly in at least one, so all wealth goes to the risky asset. For $x_1 > 0$ we therefore need

$$\min\{a, b\} < 1.$$

Answer. *Part (b).* If $\pi a + (1-\pi)b \le 1$, the risky asset's expected payoff does not exceed the riskless payoff of 1. A risk-averse DM would then never touch the risky asset (any positive x_2 strictly lowers expected utility). So for $x_2 > 0$ we need

$$\pi\, a + (1-\pi)\, b > 1.$$

Answer. *Part (c): first-order condition.* Terminal wealth in state 1 is $x_1 + x_2 a$ (state 1 occurs with prob π); in state 2 it is $x_1 + x_2 b$. Expected utility is $\pi\, u(x_1 + x_2 a) + (1-\pi)\, u(x_1 + x_2 b)$. Substituting the constraint $x_2 = 1 - x_1$ and differentiating in x_1:

$$\pi(1-a)\, u'(x_1 + (1-x_1)a) + (1-\pi)(1-b)\, u'(x_1 + (1-x_1)b) = 0. \tag{FOC}$$

This, together with $x_1 + x_2 = 1$, characterizes the interior optimum.

Answer. *Part (d): $dx_1/da \le 0$.* Let $\Phi(a, \pi, x_1)$ denote the LHS of the FOC (with b held fixed). Compute

$$\frac{\partial \Phi}{\partial a} = -\pi\, u'(\cdot_1) + \pi(1-a)(1-x_1)\, u''(\cdot_1) < 0 \quad \text{(both terms negative under } a < 1,\ u'' < 0\text{)},$$
$$\frac{\partial \Phi}{\partial x_1} = \pi(1-a)^2 u''(\cdot_1) + (1-\pi)(1-b)^2 u''(\cdot_2) < 0 \quad \text{(strict concavity: SOC of the PMP).}$$

Applying the implicit function theorem,

$$\frac{dx_1}{da} = -\frac{\partial\Phi/\partial a}{\partial\Phi/\partial x_1} < 0.$$

Raising the low payoff a makes the risky asset look better on average, and the DM shifts money *out* of the riskless asset.

Answer. *Part (e): conjecture on $dx_1/d\pi$.* From Part (b), $\pi a + (1-\pi)b > 1$; if $a < 1$, this forces $b > 1$, so a is the *worse* outcome. Increasing π (probability of the bad state) makes the risky asset less attractive, so the DM should shift *into* the riskless asset:

$$\frac{dx_1}{d\pi}\ >\ 0 \qquad \text{(conjectured).}$$

Answer. *Part (f): proof.*

$$\frac{\partial \Phi}{\partial \pi} = (1-a)\, u'(\cdot_1) - (1-b)\, u'(\cdot_2) = (1-a)\, u'(\cdot_1) + (b-1)\, u'(\cdot_2) > 0,$$

because $a < 1$ (so $1 - a > 0$), $b > 1$ (so $b - 1 > 0$), and both $u' > 0$. Combined with $\partial\Phi/\partial x_1 < 0$ from Part (d),

$$\frac{dx_1}{d\pi} = -\frac{\partial\Phi/\partial\pi}{\partial\Phi/\partial x_1} > 0,$$

confirming the conjecture.

Intuitively, the FOC balances the marginal expected utility of the two assets. Raising π tilts the balance toward the bad state of the risky asset, and the DM restores equality by pulling wealth into the riskless asset. The implicit function theorem does the accounting.

7. **Complement (buying a mystery "black box").** Exercises 1–6 above all posed the DM's problem in the abstract; this complement gives a concrete decision under uncertainty. A consumer with utility $u = a(x)$ (only good x matters; $a' > 0$, $a'' < 0$) can buy good x directly at unit price p_1, or buy a "black box" of size z at unit price p_2: each unit of box costs p_2 and delivers one unit of x with probability $\alpha \in (0, 1)$ (and one unit of a useless good y with probability $1 - \alpha$). She maximizes expected utility.

   (a) When will she buy a strictly positive amount of the box?

   (b) Derive the demands $x^*(p_1, p_2, w)$ (direct x) and $z^*(p_1, p_2, w)$ (box).

   (c) For $a(x) = \ln x$, $p_2 = 1$, $w = 1$, $\alpha = 1/2$, plot both the direct-good demand $x^*(p_1)$ and the box demand $z^*(p_1)$ on a single diagram.

Answer. *Part (a): when the box is bought.*

Step 1. Set up EU. Let x_0 be the sure amount bought directly. If she also buys z units of the box, the final x-holding is $x_0 + z$ w.p. α (box delivers) and x_0 w.p. $1 - \alpha$ (box is a dud). Expected utility:

$$EU = \alpha\, a(x_0 + z) + (1 - \alpha)\, a(x_0), \qquad p_1 x_0 + p_2 z = w.$$

Step 2. Marginal EU at $z = 0$ (starting from all direct x). At $z = 0$, one extra dollar spent on the box yields (in expectation) $\alpha\, a'(x_0)/p_2$; one extra dollar spent on direct x yields $a'(x_0)/p_1$. Starting the box is worth it iff

$$\frac{\alpha}{p_2} > \frac{1}{p_1} \iff p_2 < \alpha\, p_1.$$

The box delivers an expected unit of x at effective cost p_2/α; she gambles only when that beats the sure price p_1.

Answer. *Part (b): closed-form demands.*

Step 1. If $p_2 \geq \alpha p_1$: no gamble. $z^* = 0$ and $x^* = w/p_1$.

Step 2. If $p_2 < \alpha p_1$: interior. The two FOCs (from $\max EU$ subject to the budget) are

$$\frac{\alpha\, a'(x_0 + z) + (1 - \alpha)\, a'(x_0)}{p_1} = \frac{\alpha\, a'(x_0 + z)}{p_2},$$

together with the budget. These two equations pin down (x^*, z^*). Risk aversion ($a'' < 0$) keeps a positive sure purchase $x^* > 0$ even when the box is attractive: putting all wealth into the gamble would leave the DM with nothing in the bad state.

Answer. *Part (c): the* ln *example.*

With $a(x) = \ln x$, $\alpha = 1/2$, $p_2 = 1$, $w = 1$: the box threshold from Part (a) is $p_2 < \alpha p_1$, i.e., $p_1 > p_2/\alpha = 2$. For p_1 below the threshold the DM buys direct x only; above the threshold she buys some direct x and some box.

Step 1. Regime $p_1 \leq 2$ (no box). All wealth on the direct good:

$$z^*(p_1) = 0, \qquad x^*(p_1) = \frac{w}{p_1} = \frac{1}{p_1}.$$

Step 2. Regime $p_1 > 2$ (interior mix). Solving the interior FOCs with ln utility, $\alpha = 1/2$, $p_2 = 1$, $w = 1$ (see the closed form in Part (b)) yields

$$x^*(p_1) = \frac{1}{2(p_1 - 1)}, \qquad z^*(p_1) = \frac{p_1 - 2}{2(p_1 - 1)}.$$

At the threshold $p_1 = 2$, both formulas give $x^* = 1/2$ and $z^* = 0$: demand is *continuous* across the regime switch. But the two regimes have different slopes, so x^* has a downward *kink* at $p_1 = 2$, and z^* starts rising from zero at that same point. Figure 1 shows both together.

Intuitively, the box behaves like a risky version of direct x: it delivers the same unit but only with probability α. As long as p_1 is low the box is dominated (direct x is cheaper per expected unit); once p_1 crosses the threshold $p_2/\alpha = 2$, the risk-averse DM starts mixing — she still buys some direct x (to insure against the bad state) but adds a positive box to save on the expected cost of extra x. In the figure, the dotted continuation of $1/p_1$ past $p_1 = 2$ shows what direct demand *would* have been if the box were not available; the vertical gap between the solid x^* curve and this dotted counterfactual is exactly the direct purchase the DM substitutes away from once the box becomes attractive, and it matches the height of the box-demand curve $z^*(p_1)$ scaled by the effective substitution rate.

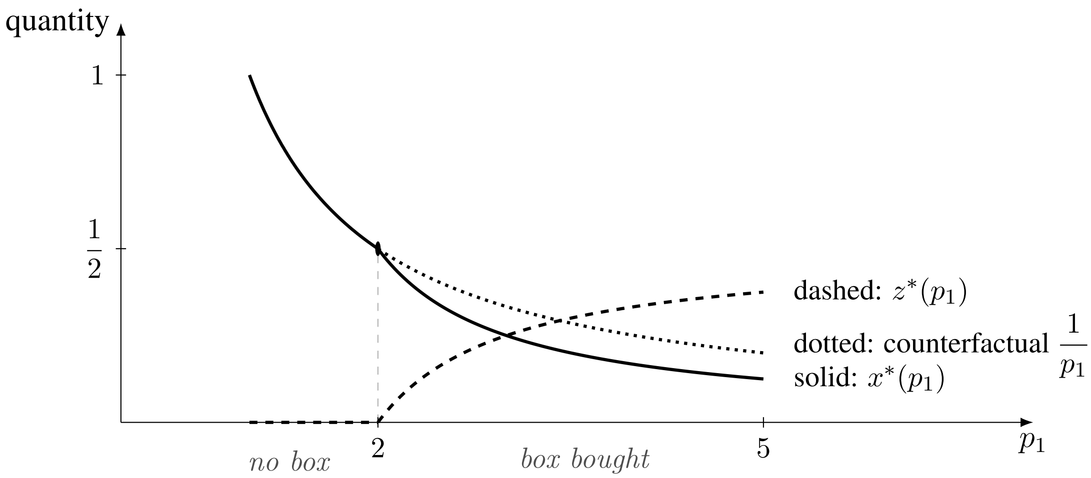

Figure 1: Direct-good demand x^* and box demand z^* as functions of p_1 in the mystery-box example ($a = \ln$, $p_2 = 1$, $w = 1$, $\alpha = \frac{1}{2}$). For $p_1 \leq 2$ ("no box") the DM buys only direct x with $x^* = \frac{1}{p_1}$ and $z^* = 0$. For $p_1 > 2$ ("box bought") she mixes: direct demand kinks downward to $x^* = \frac{1}{2(p_1 - 1)}$ (solid), and box demand rises from zero to $z^* = \frac{p_1 - 2}{2(p_1 - 1)}$ (dashed). The dotted continuation of $\frac{1}{p_1}$ past the threshold shows the counterfactual: what direct demand *would* have been if the box were not available. The vertical gap between the solid and dotted lines is the direct-good demand that the box *substitutes for*.