Decision and Social Choice

PE3101P — Lecture 9: Calibration

Tomi Francis

National University of Singapore

15th October, 2026

Today

1. Introduction
2. Calibration and Risk
3. Calibration and Distribution
4. Gini Egalitarianism
5. Calibration Without Numbers
6. Ways Out

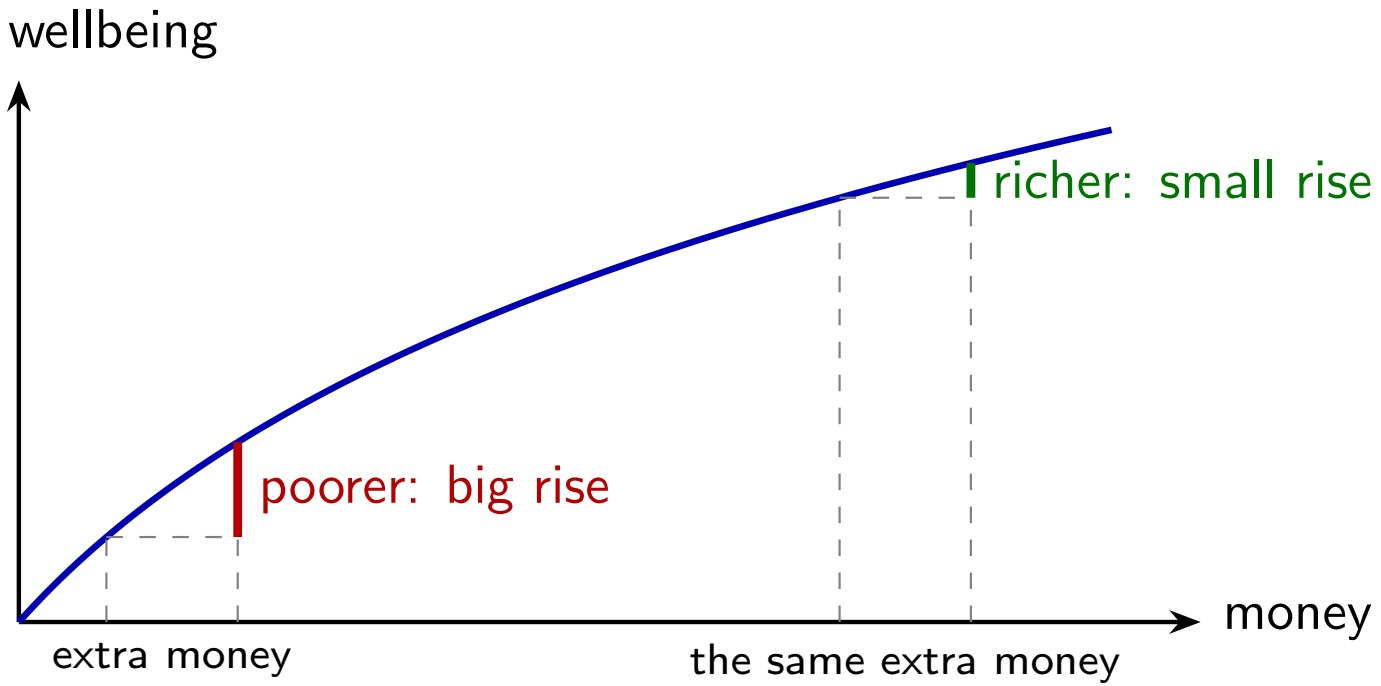
An Extra Dollar

- Almost everyone believes that, other things being equal, matters go better if an extra dollar goes to somebody who is worse off rather than somebody who is better off.
- Three possible reasons:
  1. benefiting the worse off matters more;
  2. a poorer person getting the dollar makes the situation more equal;
  3. a poorer person gets more use out of an additional dollar than a richer person.
- Prioritarians care about the first. Egalitarians care about the second. Virtually everybody cares about the third.

Diminishing Marginal Value

- If I have no money at all, a dollar might buy me a bag of rice, which would matter a lot if I have no other food.
- If I have a million dollars, I can already buy all the rice I can eat.
- Many utilitarians think this explains our equality-favouring judgements about resources, without our needing to say that equality matters in itself.

The Same Money at Two Wealth Levels



The same money for someone richer: a small rise.
So it does more for total wellbeing if the poorer person gets it.

The Calibration Problem

- Roughly: if we favour equality moderately in low-stakes cases, where not much money or wellbeing is at stake, then we have to favour equality to an implausibly extreme extent in high-stakes cases.
- The calibration problem undermines the utilitarian explanation, and also many views on which equality or priority to the worse off *is* significant in itself.
- The problem arises for risk as well as for distribution, because states of nature play the role that people play.

Where the Calibration Problem Bites

1. **Risk.** Risk aversion, in preferences under risk.
2. **Utilitarianism.** The diminishing value of money.
3. **Egalitarianism, prioritarianism and views like them.** Priority weighting.

- So it's an issue for everybody, not just for egalitarians and prioritarians.

Would You Take This Gamble? (1)

A 50–50 gamble:

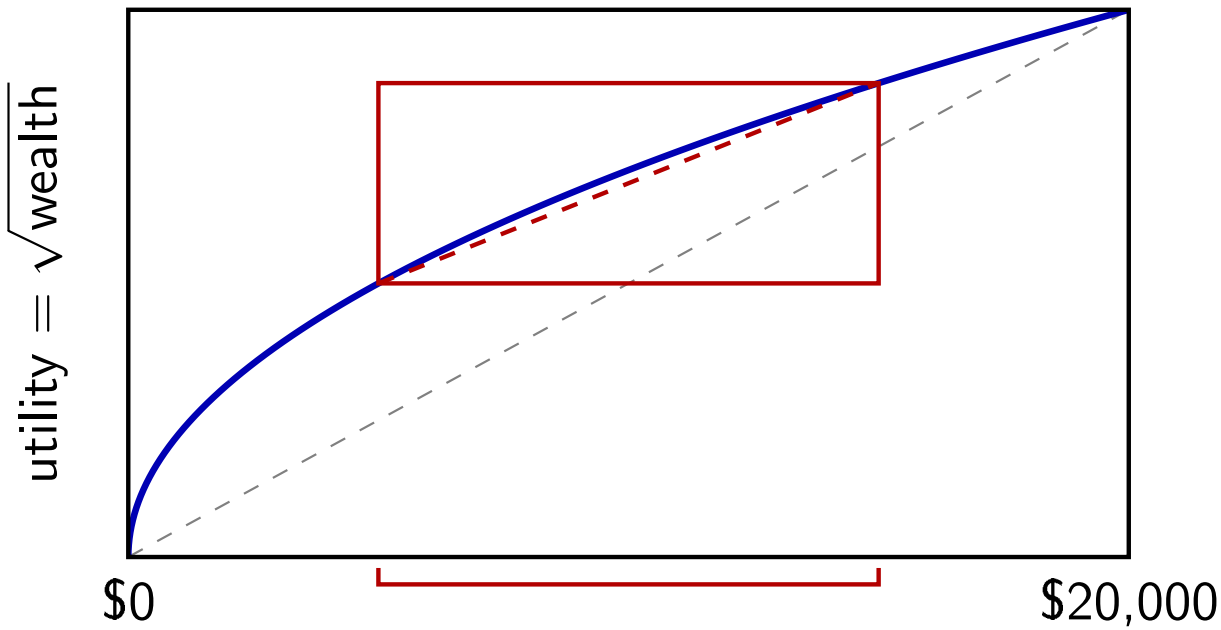
gain \$1,100 or lose \$1,000.

Would You Take This Gamble? (2)

A 50–50 gamble:

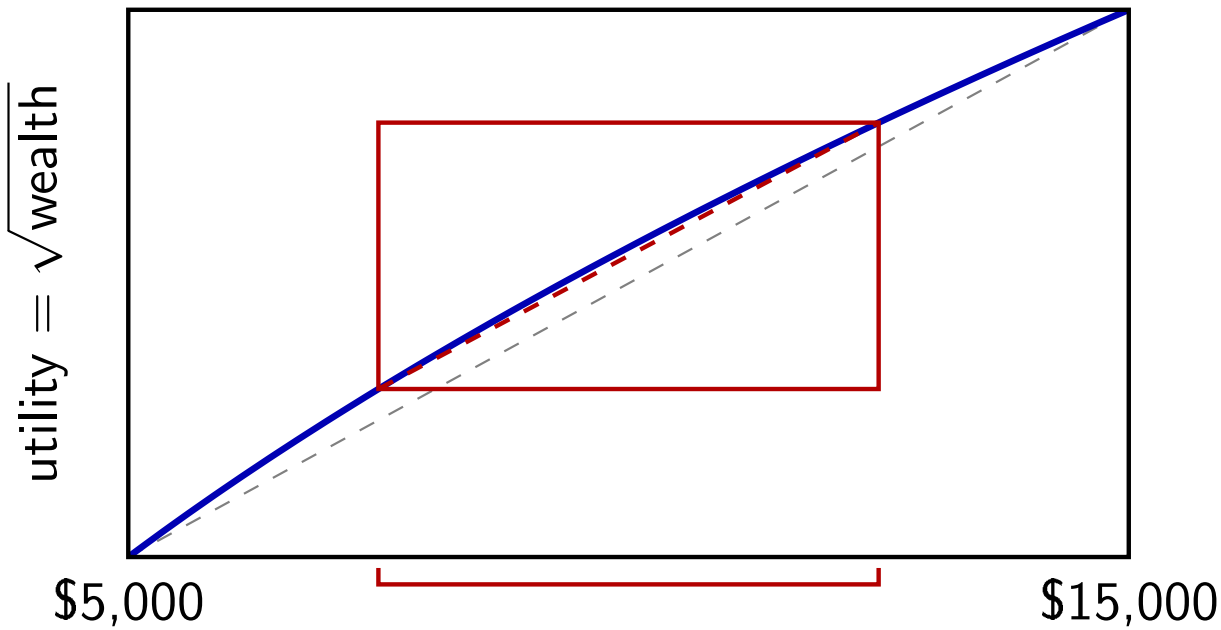
gain $1,000,000,000 or lose $7,000.

Zooming In on a Concave Curve



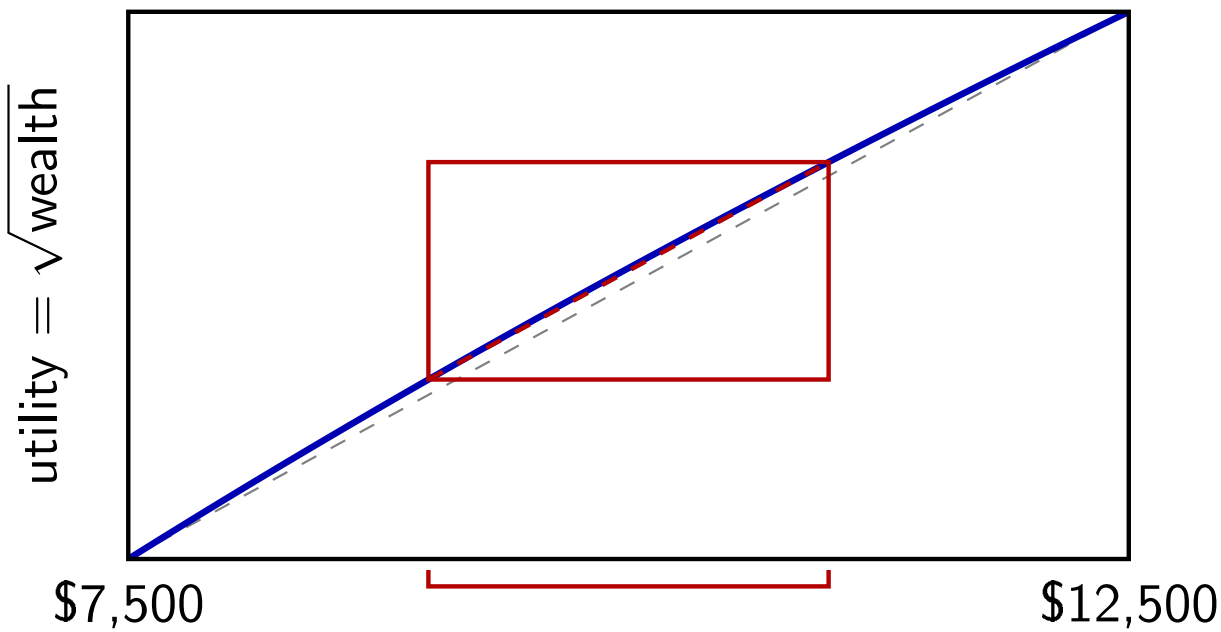
The straight line across the box:
the dashed line in the next picture.

Zooming In on a Concave Curve



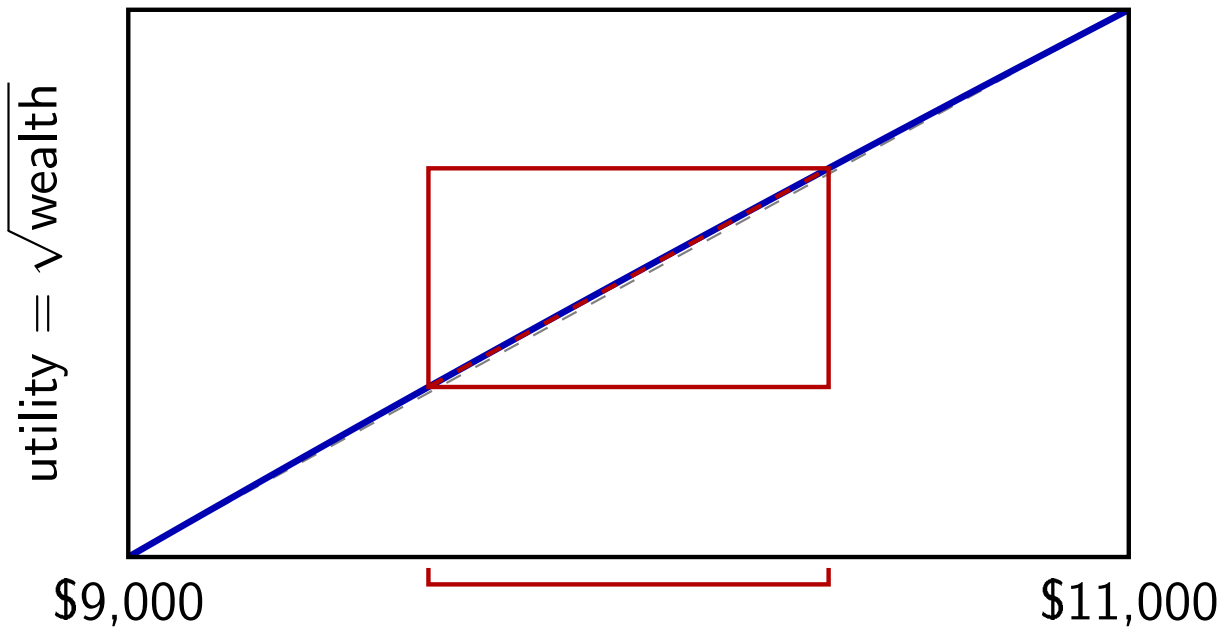
The straight line across the box:
the dashed line in the next picture.

Zooming In on a Concave Curve



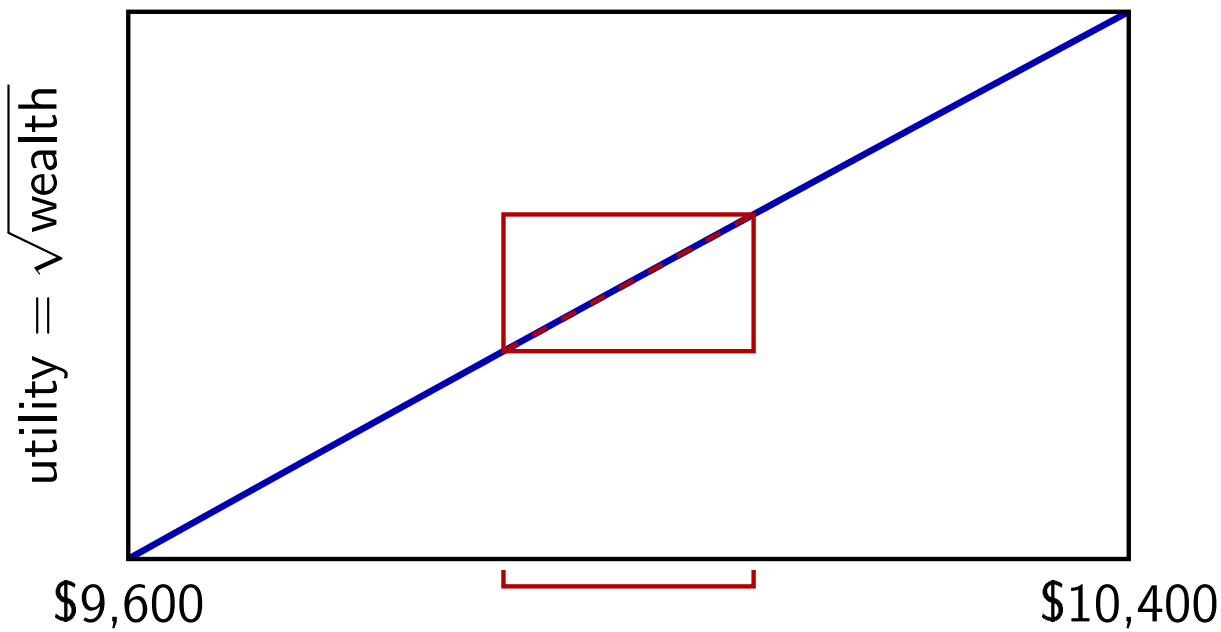
The straight line across the box:
the dashed line in the next picture.

Zooming In on a Concave Curve



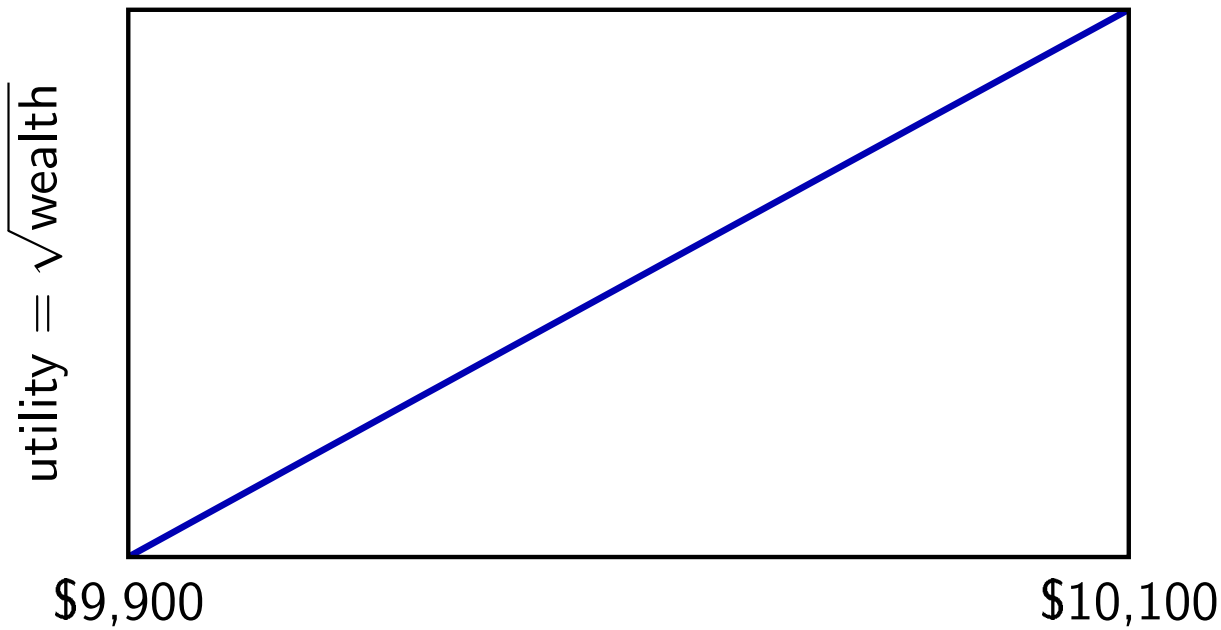
The straight line across the box:
the dashed line in the next picture.

Zooming In on a Concave Curve



The straight line across the box:
the dashed line in the next picture.

Zooming In on a Concave Curve



From $9,900 to $10,100.
Almost a straight line.

Less Risk Averse, the Closer You Look

- A concave utility function is a risk-averse one: the curve bends.
- Zoom in on a concave curve and it gets flatter and flatter, closer and closer to a straight line. So for small gambles, it is almost risk neutral.
- So if someone is still risk averse when we zoom in a lot, their curve must still bend at that small scale.
- Then, zoomed out, it bends a lot: they are *extremely* risk averse about large gambles.

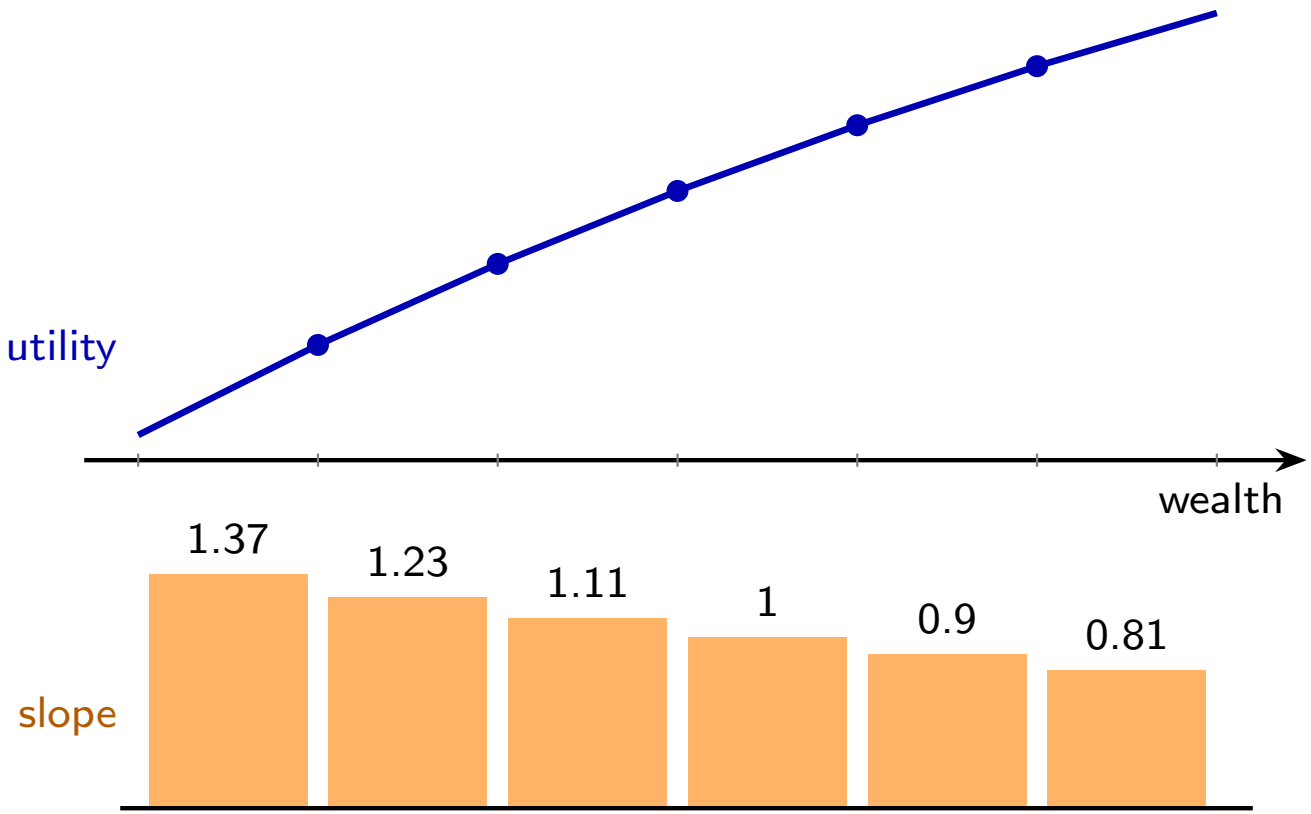
Square Root at $10,000

- Someone whose utility is $\sqrt{\text{wealth}}$ has \$10,000: utility 100.
- A 50–50 gamble is worth taking if the utility they'd gain if they win is more than the utility they'd lose if they lose.

50–50 gamble	Utility lost	Utility gained
Lose \$1,900 / gain \$2,000	$100 - 90 = 10$	less than $110 - 100 = 10$
Lose \$1,000 / gain \$1,100	about 5.1	about 5.4

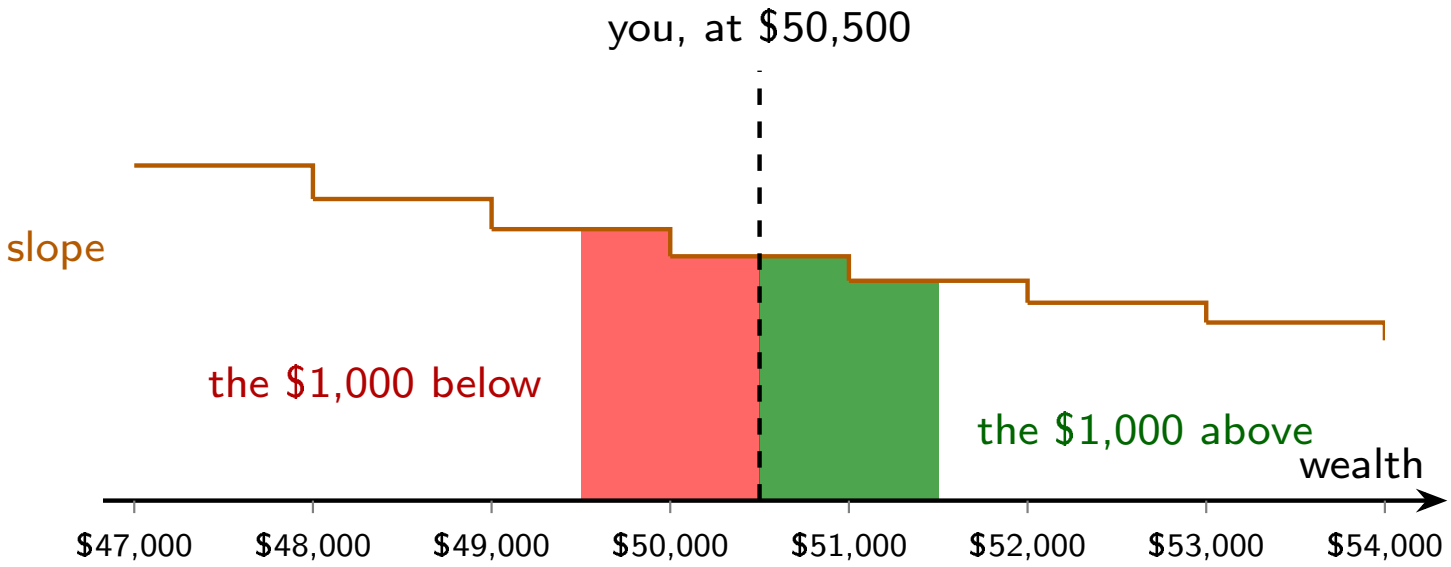
Over \$1,000 the curve is almost straight:
gaining \$1,100 beats losing \$1,000, so they take it.
Square root isn't risk averse enough to turn down small gambles like this.

A Utility Function Made of Straight Pieces



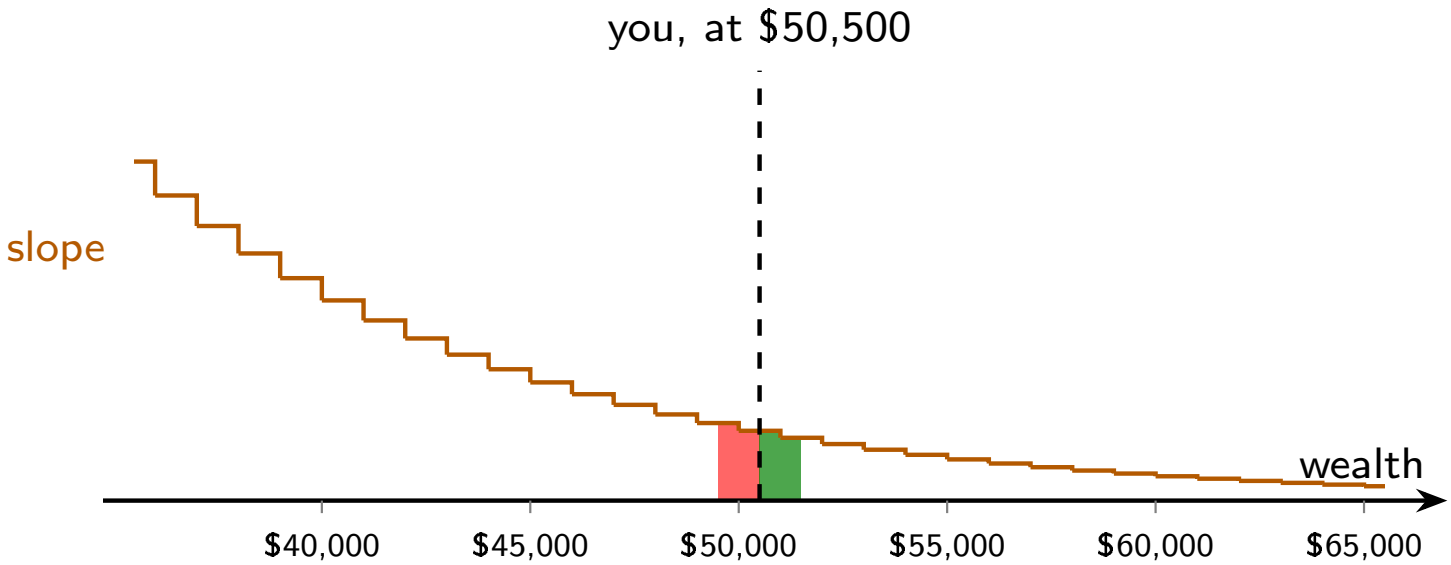
Every \$1,000, the slope is multiplied by 0.9.

Diminishing Marginal Utility



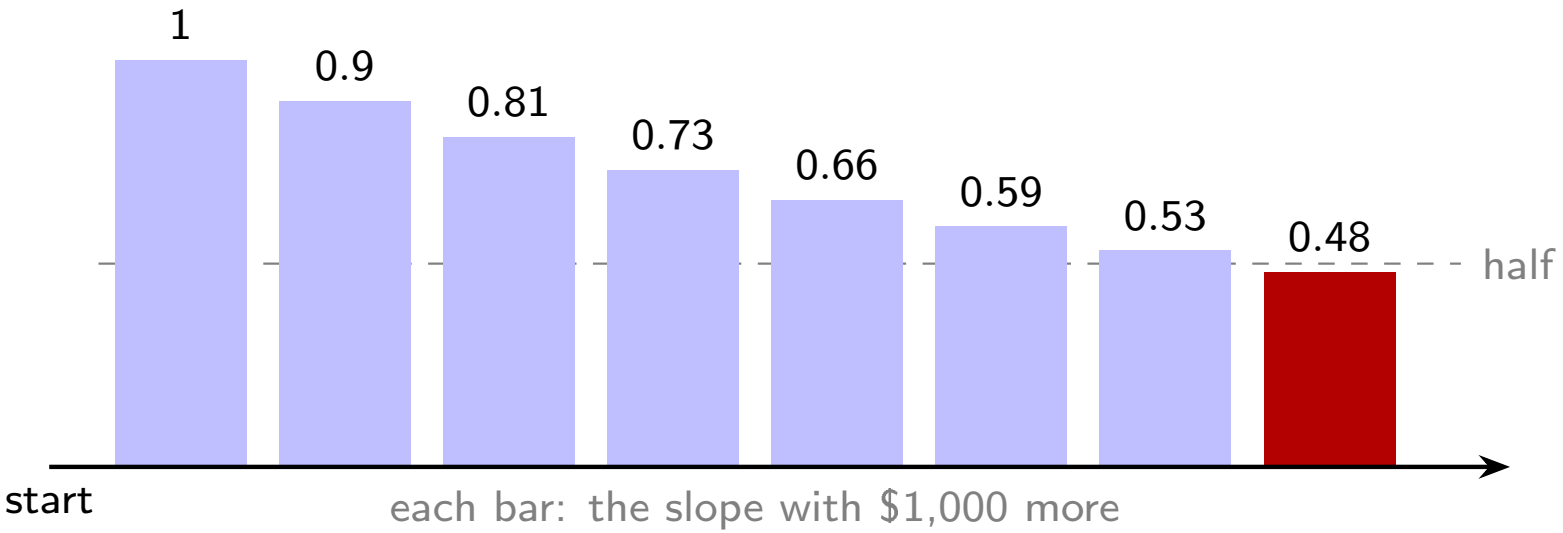
So gaining \$1,100 is worth at most $0.9 + 0.09 = 0.99$ of losing \$1,000.
That's why we chose 0.9: it's the kind of drop-off needed
to turn down 50–50 lose \$1,000 / gain \$1,100, wherever you start.

Diminishing Marginal Utility



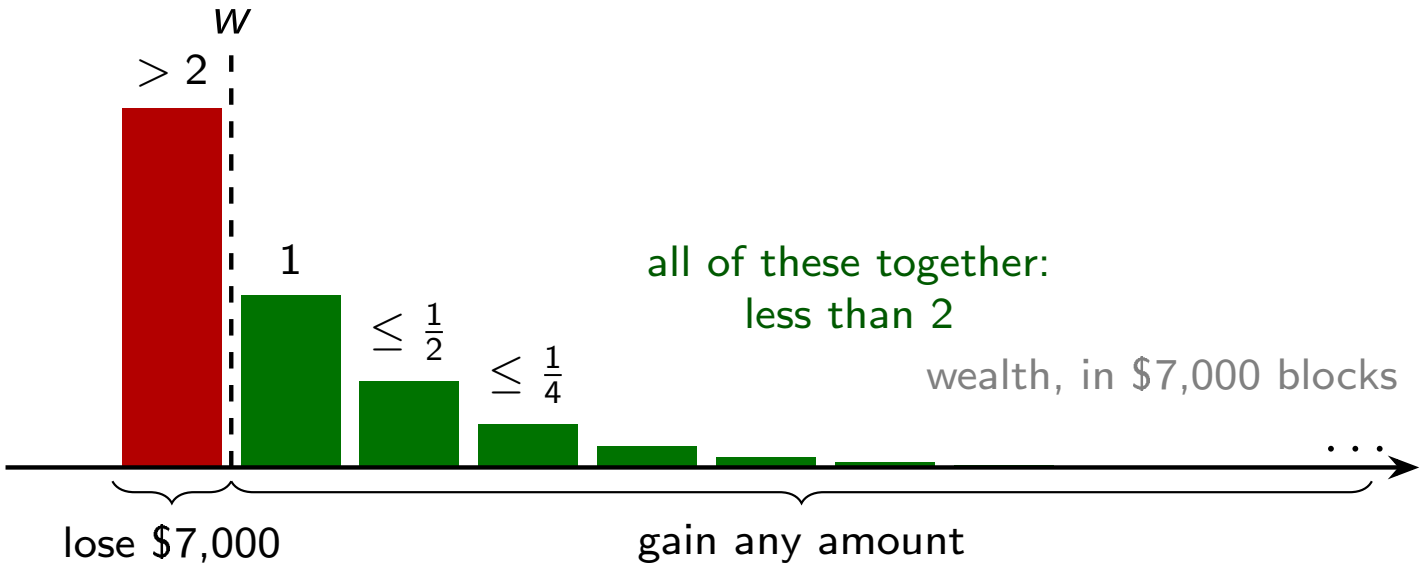
Zoomed out, from about $35,000 to $65,000:
the less wealth, the higher and higher the slope.

Seven Steps Halve the Slope



After seven steps, with \$7,000 more: $0.9^7 \approx 0.48$, less than half.
So with every \$7,000 more, the slope at least halves.

Halving, Again and Again



So they turn down 50–50 lose \$7,000 / gain G, however large G is: even $G = \$1{,}000{,}000{,}000$, an expected gain of \$499,996,500.

Calibration, More Generally

- Our step function is just one example.
- Rabin (2000) showed that it works for every concave utility function that turns down 50–50 lose \$1,000 / gain \$1,100 at every wealth level.

Risk and Distribution

- Last week: states of nature play the role that people play.

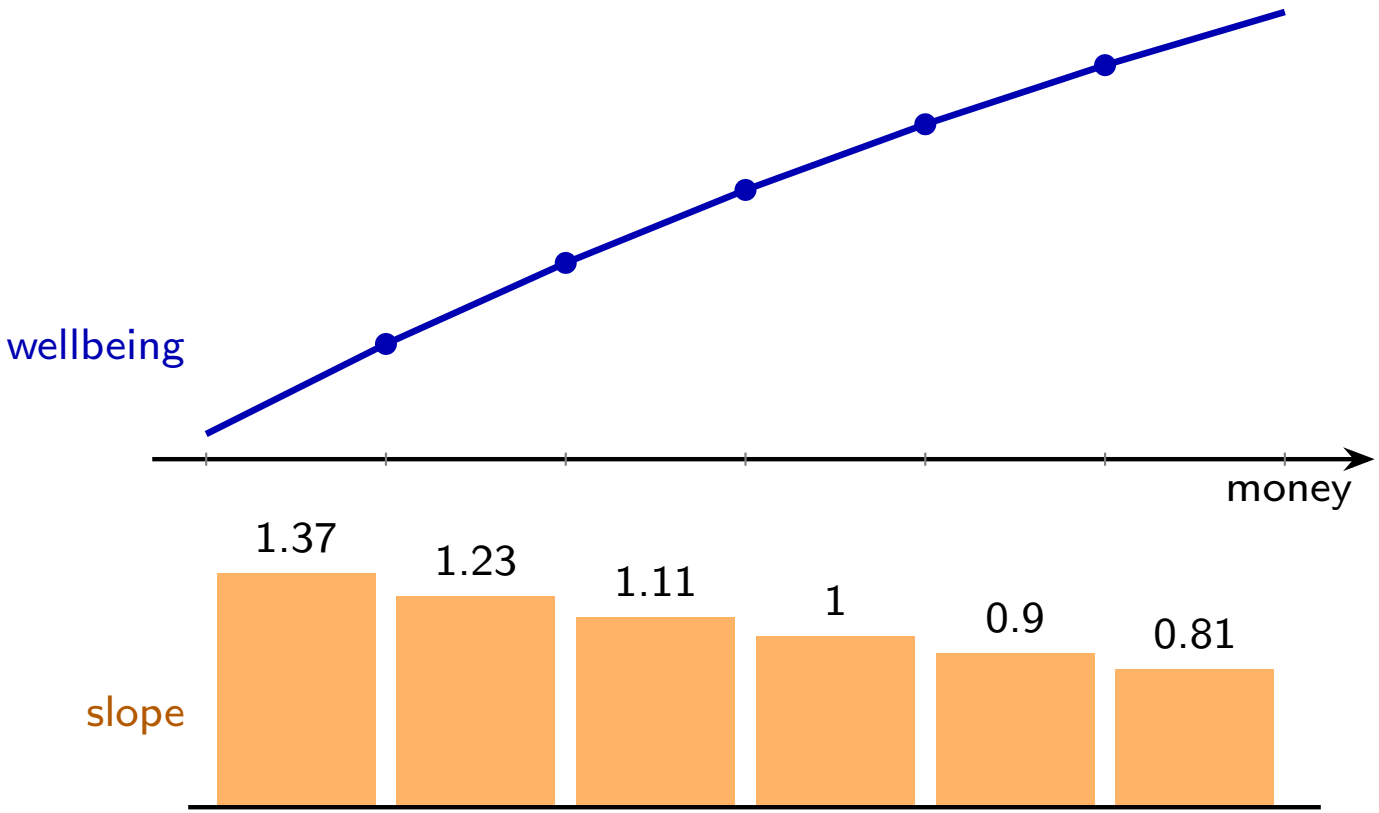
Risk	Distribution
Equally likely states of nature	People
Utility of money	Wellbeing from money, or priority weighting
Turn down lose \$1,000 / gain \$1,100	Prefer (w, w) to $(w - 1{,}000,\ w + 1{,}100)$

So we can run the same argument twice more:
for a utilitarian about money, and for a prioritarian.

The Same Problem for Money (1)

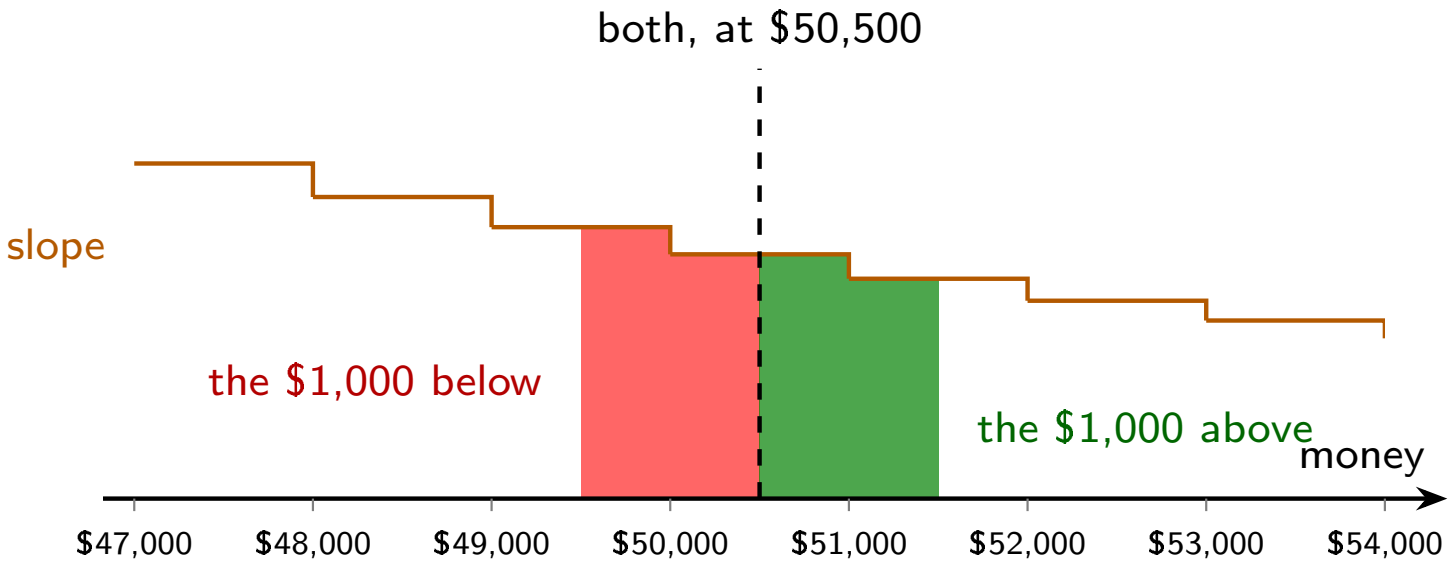
- Two people, equally well off.
- Is it better to take $1,000 from one and give $1,100 to the other?
- Is it better to take $7,000 from one and give $1,000,000,000 to the other?

Wellbeing From Money in Straight Pieces



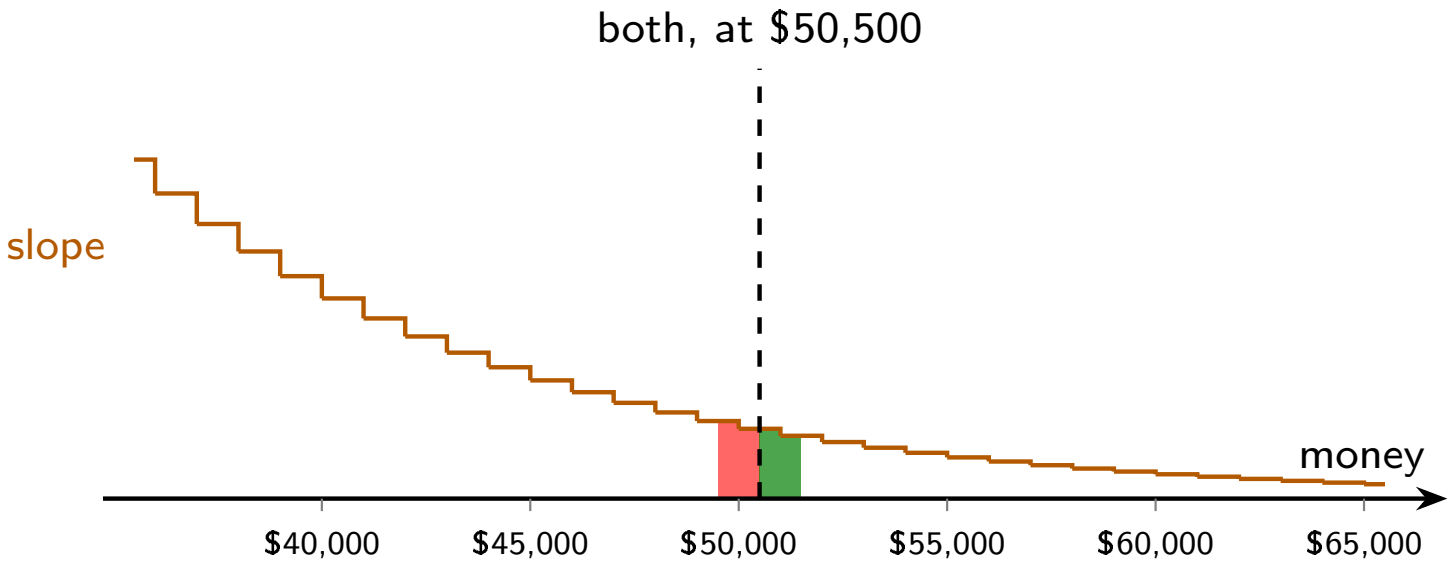
Every \$1,000, the slope is multiplied by 0.9.

Diminishing Marginal Wellbeing



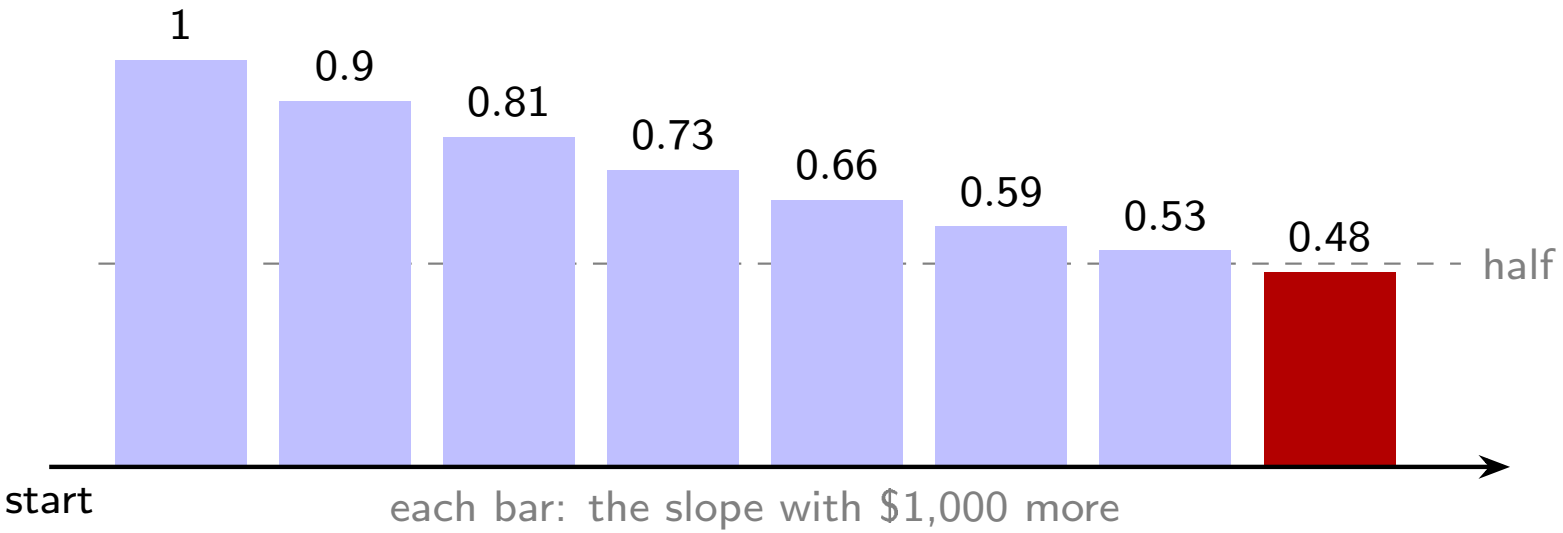
So \$1,100 more for one adds at most $0.9 + 0.09 = 0.99$ of what \$1,000 less costs the other: the drop-off needed for the utilitarian never to take \$1,000 from one to give \$1,100 to the other.

Diminishing Marginal Wellbeing



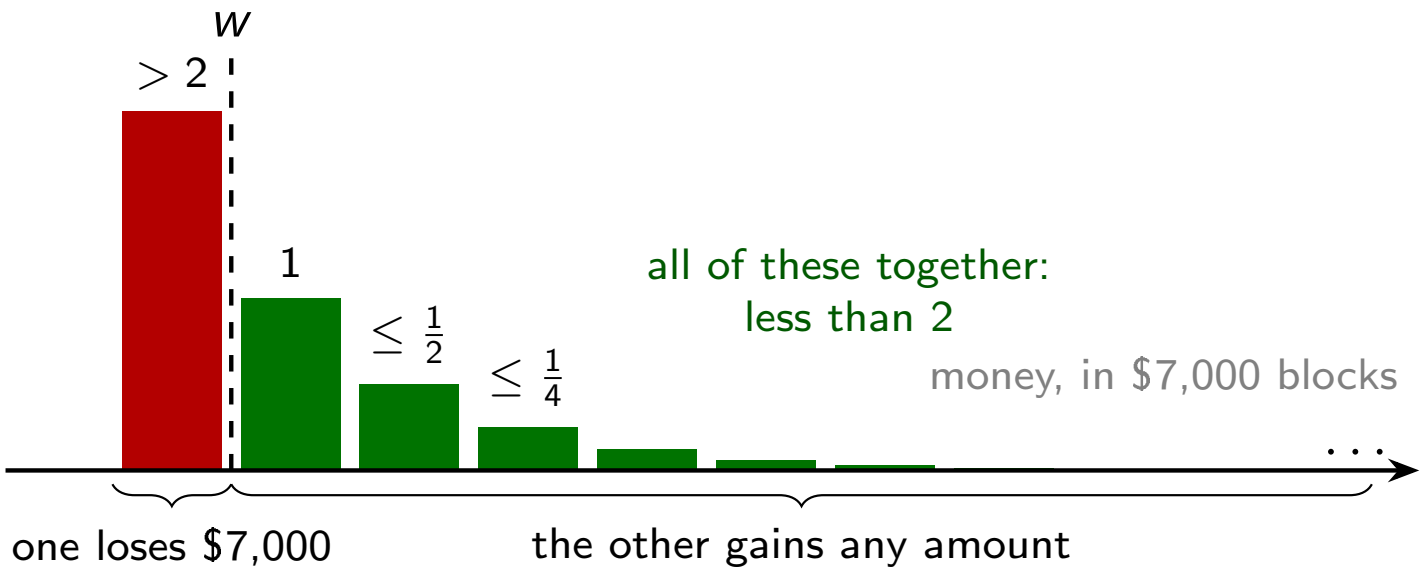
Zoomed out, from about $35,000 to $65,000:
the less money, the higher and higher the slope.

Seven Steps Halve the Wellbeing per Dollar



After seven steps, with \$7,000 more: $0.9^7 \approx 0.48$, less than half.
So with every \$7,000 more, the wellbeing per dollar at least halves.

Halving for Utilitarians



So the utilitarian won't take $7,000 from one to give G to the other, however large G is: not even $1,000,000,000.

What This Means for Utilitarians

Risk	Money, for a utilitarian
One person, two equally likely states	Two people, equally well off
Utility per dollar	Wellbeing per dollar
Turn down lose \$1,000 / gain \$1,100	Don't take \$1,000 from one to give \$1,100 to the other
Turn down lose \$7,000 / gain G	Don't take \$7,000 from one to give G to the other

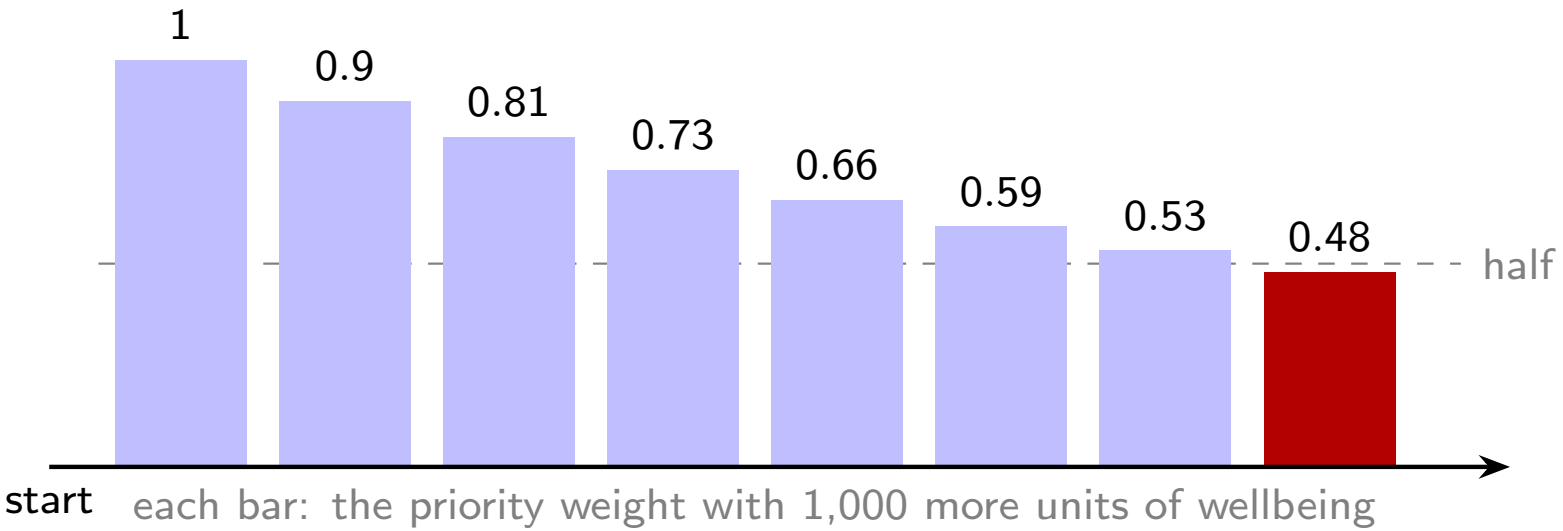
However large G is.

So the utilitarian explanation of our equality-favouring judgements about money runs into the calibration problem too.

Priority Weights in Straight Pieces

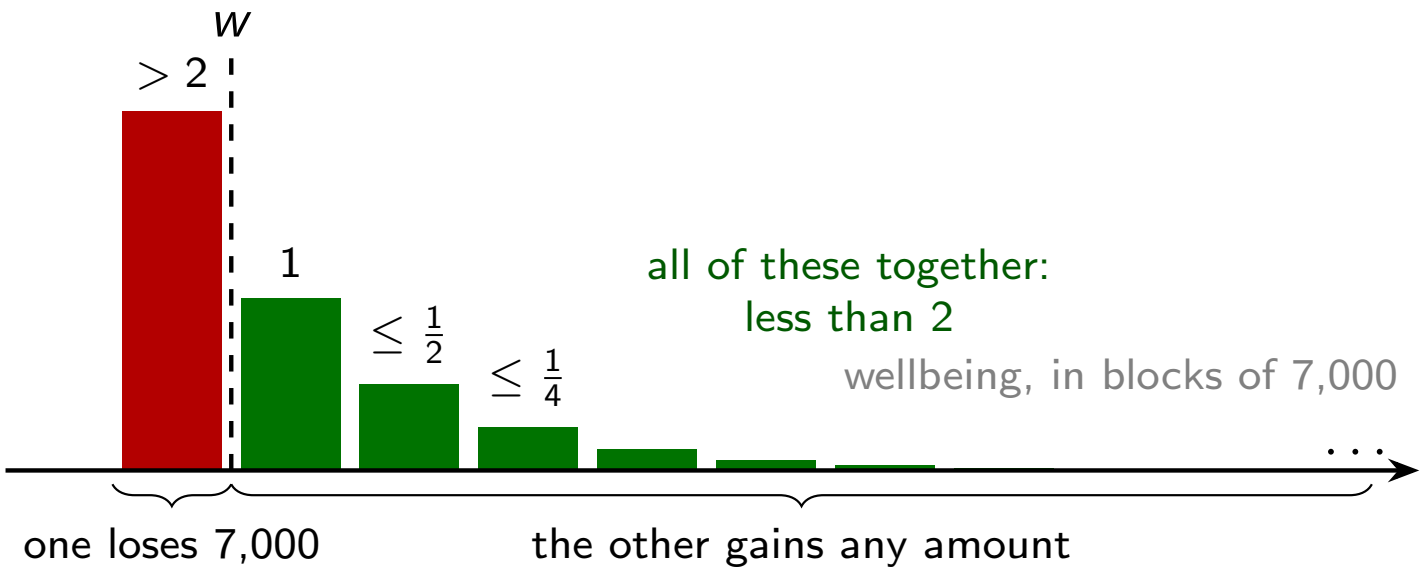
- The priority weight: how much each extra unit of wellbeing adds to the priority-weighted sum, at that level.
- Suppose that every 1,000 units of wellbeing, the priority weight is multiplied by 0.9.
- Then the 1,000 units above someone count for 0.9 of the 1,000 units below them, wherever they are.
- So this prioritarian prefers (w, w) to $(w - 1{,}000, w + 1{,}100)$, at every level w.

Seven Steps Halve the Priority Weight



After seven steps, with 7,000 more units: $0.9^7 \approx 0.48$, less than half.
So with every 7,000 more units, the priority weight at least halves.

Halving for Prioritarians



So (w, w) beats $(w - 7{,}000, w + G)$, however large G is.

What This Means for Prioritarians

- A prioritarian who prefers (w, w) to $(w - 1{,}000, w + 1{,}100)$ at every level must prefer (w, w) to $(w - 7{,}000, w + G)$, however large G is.
- With $2n$ people, pair them off: $2n$ people at w beat n at $w - 7{,}000$ and n at $w + G$.
- As with risk, this works for any concave priority weighting. Nebel and Stefánsson (2023) make this point about prioritarianism.

Egalitarianism Without Adding Up

- The argument with numbers added up a priority-weighted number for each person. So it works against prioritarianism.
- But an egalitarian needn't rank outcomes by adding up a number for each person.
- Adding up like that is separable. An egalitarian who denies Separability over Persons can't add up like that.
- So let's look at one egalitarian view that doesn't: Gini egalitarianism.

The Gini Coefficient

The Gini coefficient

$$\frac{S}{2n^2 w_A}$$

where S is the sum of the gaps $|w_i - w_j|$ over all i and j, w_i is person i's wellbeing, n is the number of people and w_A is the average wellbeing.

- Add up the gaps between every pair of people, each pair counted both ways: that is S. Dividing by n^2 gives the average gap.
- Then divide by twice the average wellbeing, $2w_A$, so that the scale doesn't matter.
- It's 0 when everyone is equal, and close to 1 when one person has nearly everything.

The Gini Coefficient at Work

Alex at 2, Darius at 10, Philip at 20. Total 32, average $w_A = \frac{32}{3}$.

Gap	Alex	Darius	Philip
Alex	0	8	18
Darius	8	0	10
Philip	18	10	0

Divided by twice the average wellbeing, $2w_A = 2 \times \frac{32}{3} \approx 21.3$:
Gini coefficient $8 \div 21.3 = 0.375$.

Gini Egalitarianism

Gini egalitarianism

Rank outcomes by total wellbeing $\times$ (1 $-$ Gini coefficient). (For positive wellbeing levels.)

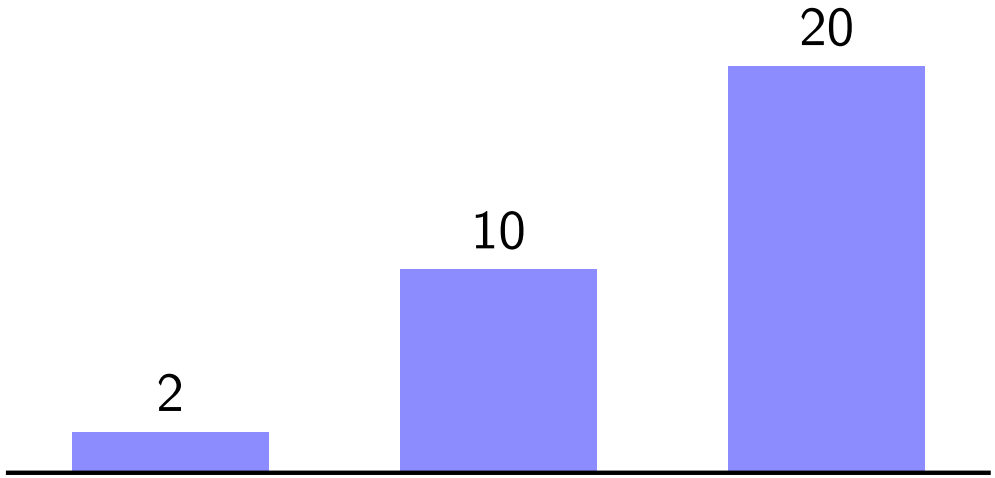
	Alex	Darius	Total	Gini coefficient	Value
UNEQUAL	4	16	20	0.3	14
EQUAL	10	10	20	0	20

- Utilitarianism: equally good. Gini egalitarianism: EQUAL is better.
- Gini egalitarianism satisfies Pigou–Dalton and Continuity, but not Separability over Persons.

Gini as Rank Weighting

- A fact we won't prove here: total wellbeing $\times$ (1 $-$ Gini coefficient) is a weighted sum.
- Line people up from best off to worst off. Weight the best off by $\frac{1}{n}$, the next by $\frac{3}{n}$, then $\frac{5}{n}$, and so on, up to just below 2 for the worst off:
$$\frac{1}{n}, \frac{3}{n}, \frac{5}{n}, \ldots, \frac{2n-1}{n}.$$
- Three people: $\frac{1}{3}$, $\frac{3}{3}$, $\frac{5}{3}$. Four people: $\frac{1}{4}$, $\frac{3}{4}$, $\frac{5}{4}$, $\frac{7}{4}$.
- The weight goes with your *rank*, not with you: the fewer people worse off than you, the bigger your weight.

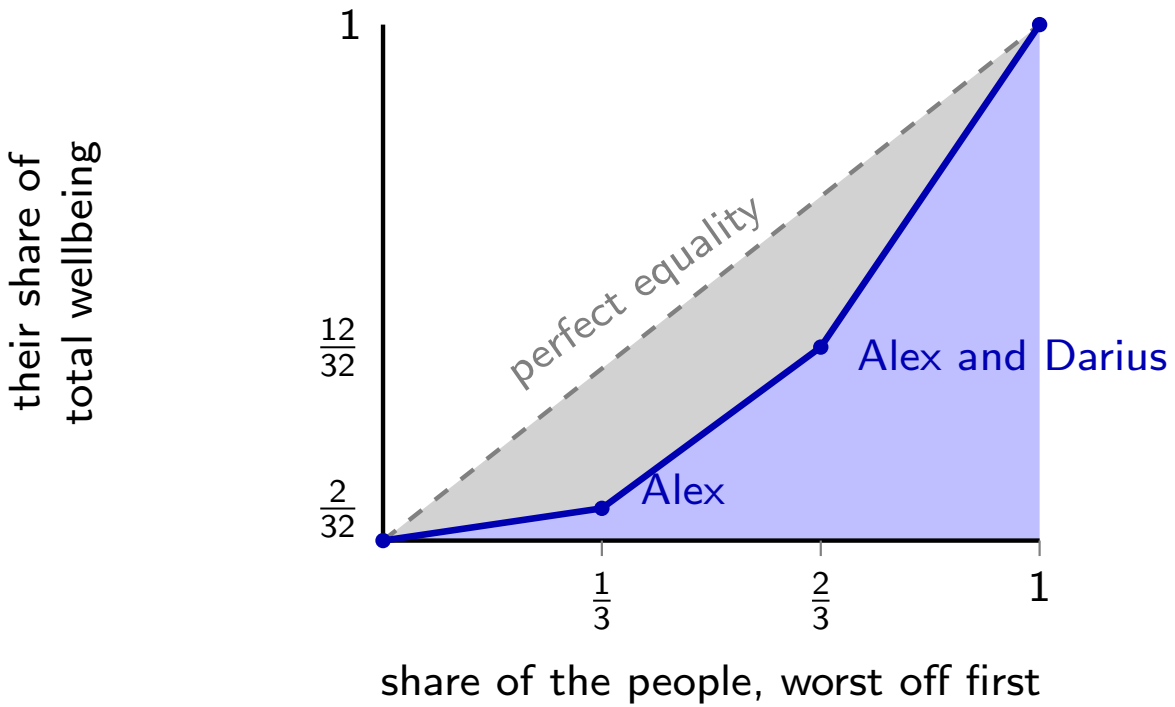
Weighting by Rank



total 60
$\div 3 = 20$

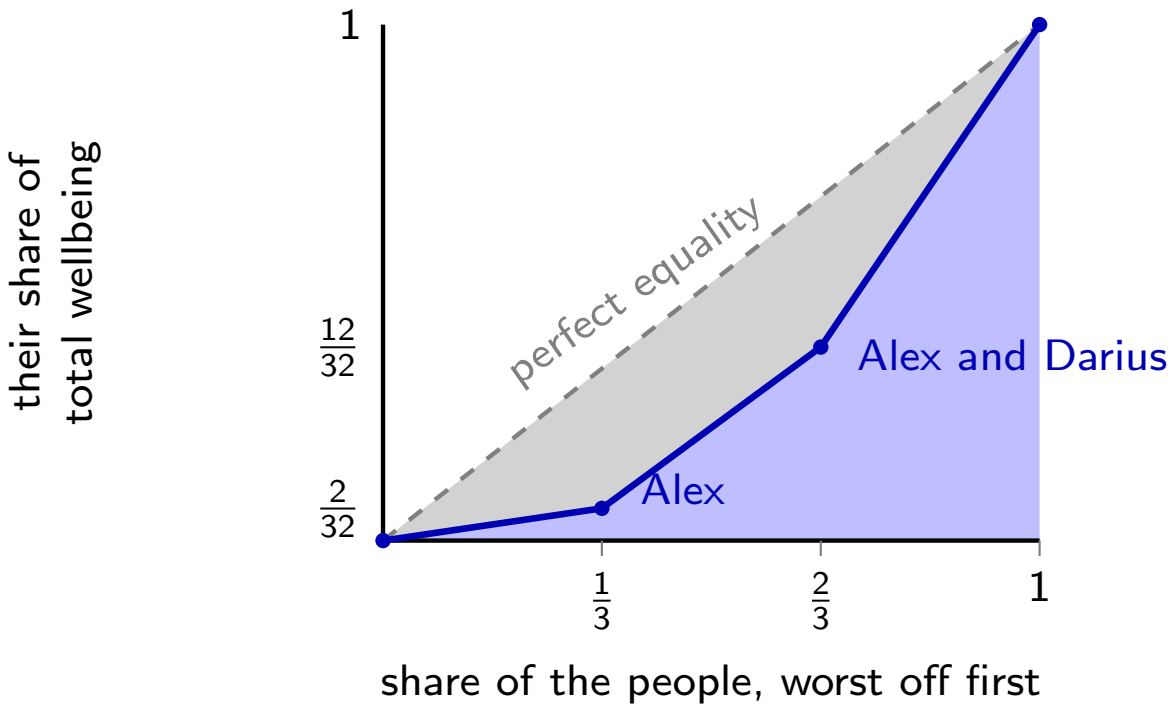
The weighted sum is $60 \div 3 = 20$.
The same as total $\times$ (1 $-$ Gini coefficient): $32 \times (1 - 0.375) = 20$.

Picturing the Gini Coefficient



Gini coefficient: the grey gap as a fraction of the whole triangle under the diagonal: $0.1875 \div 0.5 = 0.375$.

Picturing the Gini Coefficient



The grey gap is the share of total wellbeing "lost to inequality": $32 \times (1 - 0.375) = 20$, so 12 of the 32 are lost. The blue part still counts.

Gini and Separability

Score: weights 5, 3, 1 from worst off to best off (3 $\times$ the Gini value, so the same ranking).

	Alex	Darius	Philip	Score
X	6	6	1	$5 \cdot 1 + 3 \cdot 6 + 1 \cdot 6 = 29$
Y	2	14	1	$5 \cdot 1 + 3 \cdot 2 + 1 \cdot 14 = 25$
X'	6	6	20	$5 \cdot 6 + 3 \cdot 6 + 1 \cdot 20 = 68$
Y'	2	14	20	$5 \cdot 2 + 3 \cdot 14 + 1 \cdot 20 = 72$

- With Philip at 1, X beats Y.
- With Philip at 20, Y' beats X'. Alex and Darius are just as before: only Philip, who is unaffected, differs.
- Separability over Persons would need the same verdict both times.

Trade-offs Between Two People

- When a change moves wellbeing between just two people, and nobody changes place in the line, only the ratio of their two weights matters.
- The worse off has weight a, the better off weight b: a loss L to the worse off is outweighed by a gain G to the better off exactly when $G > \frac{a}{b} \cdot L$.
- From X to Y, Alex loses 4 and Darius gains 8.
- With Philip at 1, their weights are 3 and 1, a ratio of 3: the gain would have to beat $3 \times 4 = 12$. X is better.
- With Philip at 20, their weights are 5 and 3, a ratio of $\frac{5}{3}$: the gain only has to beat $\frac{5}{3} \times 4 \approx 6.7$. Y' is better.

Does Gini Egalitarianism Escape?

Two people: weights $\frac{3}{2}$ and $\frac{1}{2}$.
Value $= (3 \times$ worse off $+ 1 \times$ better off$) \div 2$.

Distribution	Value
(20,000, 20,000)	40,000
(19,000, 22,000)	39,500
(13,000, 20,000 + G)	29,500 + $G/2$

- Equality beats one losing 1,000 and the other gaining twice as much. So it beats losing 1,000 to give 1,100, at every level: the premise of the argument with numbers.
- But one losing 7,000 and the other gaining G beats equality as soon as G is more than 21,000.
- So Gini egalitarianism seems to escape the calibration problem.

Gini With Other People Around

Now add ten more people, each at 50,000. Twelve people: the two worst off get weights $\frac{23}{12}$ and $\frac{21}{12}$.

The two worst off	Their part of the value, $\times 12$
(20,000, 20,000)	$23 \cdot 20{,}000 + 21 \cdot 20{,}000 = 880{,}000$
(19,000, 21,100)	$23 \cdot 19{,}000 + 21 \cdot 21{,}100 = 880{,}100$

- The ten others add the same either way. So now Gini egalitarianism prefers one losing 1,000 and the other gaining 1,100.
- It seemed to escape when we pictured those two as the only people in existence. With other people around, it may not.
- So the escape is not robust.

Small Inequalities Aversion (1)

- One person has $30,000; another has $31,000.
- Adding $1,000 to the first, equalising their wealth, seems better than giving the second another $2,000.

	Poorer	Richer	Total
Now	$30,000	$31,000	$61,000
$1,000 to the poorer	$31,000	$31,000	$62,000
$2,000 to the richer	$30,000	$33,000	$63,000

Giving $2,000 to the richer person adds more money in total, but leaves the two more unequal.

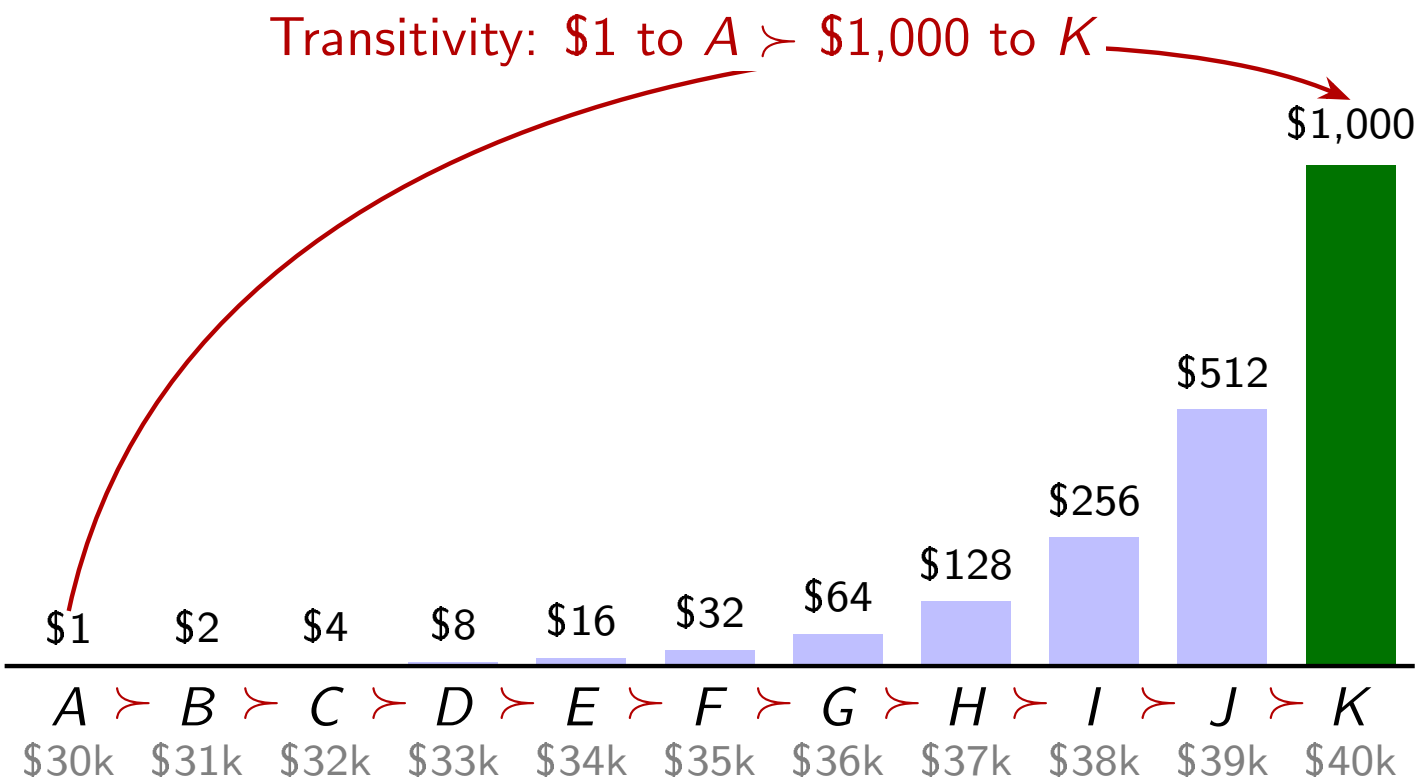
Small Inequalities Aversion (2)

Small inequalities aversion

For any two people, one of whom is worse off than the other by at least \$1,000: it is better to give the poorer person up to \$1,000 than to give the richer person up to twice as much.

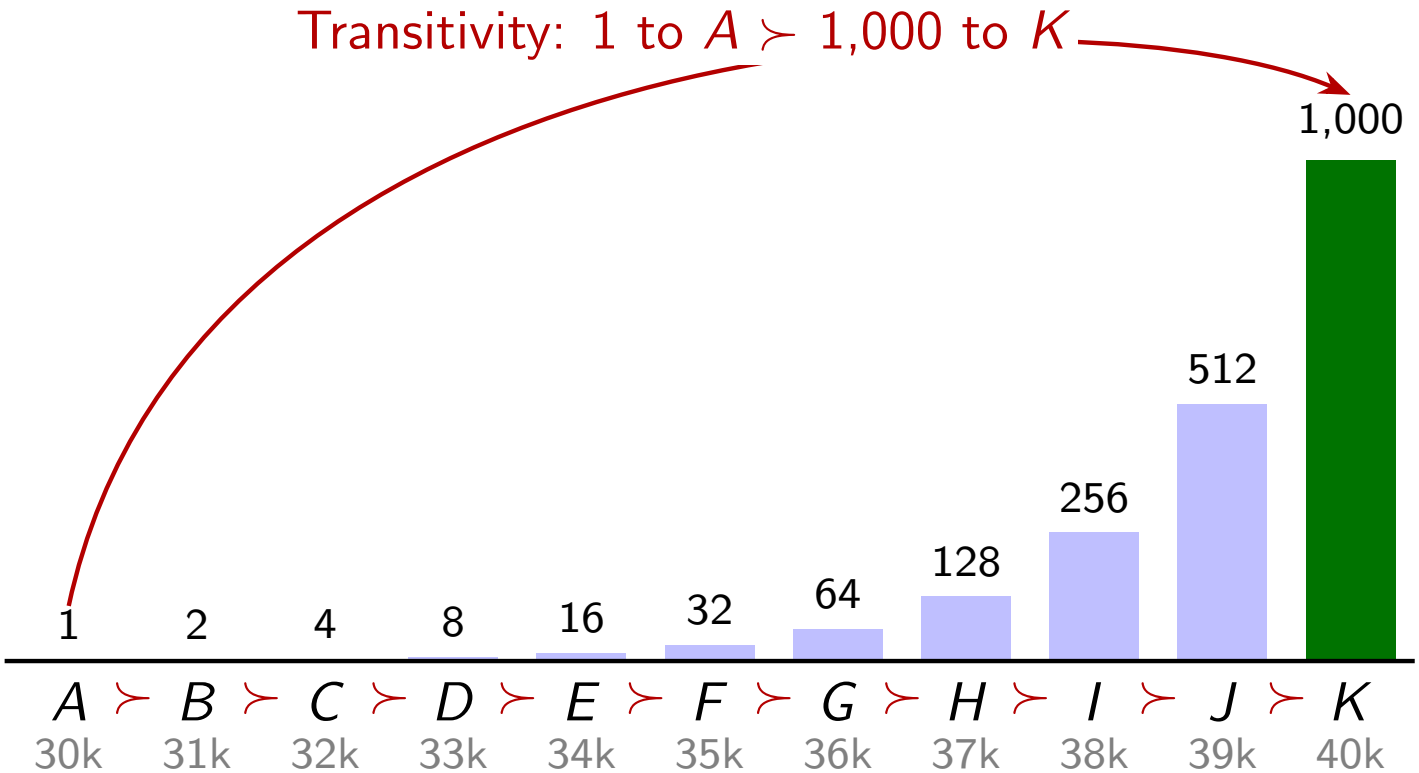
- Small inequalities aversion is meant to hold regardless of the unaffected background population.
- The background population: everyone else, who gets nothing either way.

A Chain



But it seems that the opposite is true.

The Same Chain in Wellbeing



So the same problem arises for prioritarians and egalitarians.

What the Chain Uses

- The chain uses only small inequalities aversion and Transitivity: no sum, no concave function.
- It doesn't assume any numbers at all: no utility function, no priority weighting. Only these inequality-aversion conditions.
- So the problem isn't just for prioritarians. It's a problem for all transitive rankings of distributions: egalitarian and utilitarian ones too.

Four Claims That Can't All Hold

Claim	In the chain
Small inequalities aversion	\$1 to $A \succ$ \$2 to B, and so on
... regardless of background	whoever else is in the line
Transitivity	so \$1 to $A \succ$ \$1,000 to K
Non-extreme larger inequalities aversion	\$1,000 to $K \succ$ \$1 to A

The chain shows that the four can't all hold.
Each way out gives up one of them.

Rankings Relative to Option Sets

- One way out: make the ranking relative to the option set. For example, build in an aversion to "losses" for each person, weighted more heavily than gains.
- Losses and gains are measured from a reference level, such as where the person would be in the most equal feasible distribution.
- Reference levels depend on which options are available, so the ranking of two options can differ from one option set to another.
- Another: count only inequalities among those affected by the choice. Perhaps inequalities are only bad in themselves when they are avoidable.

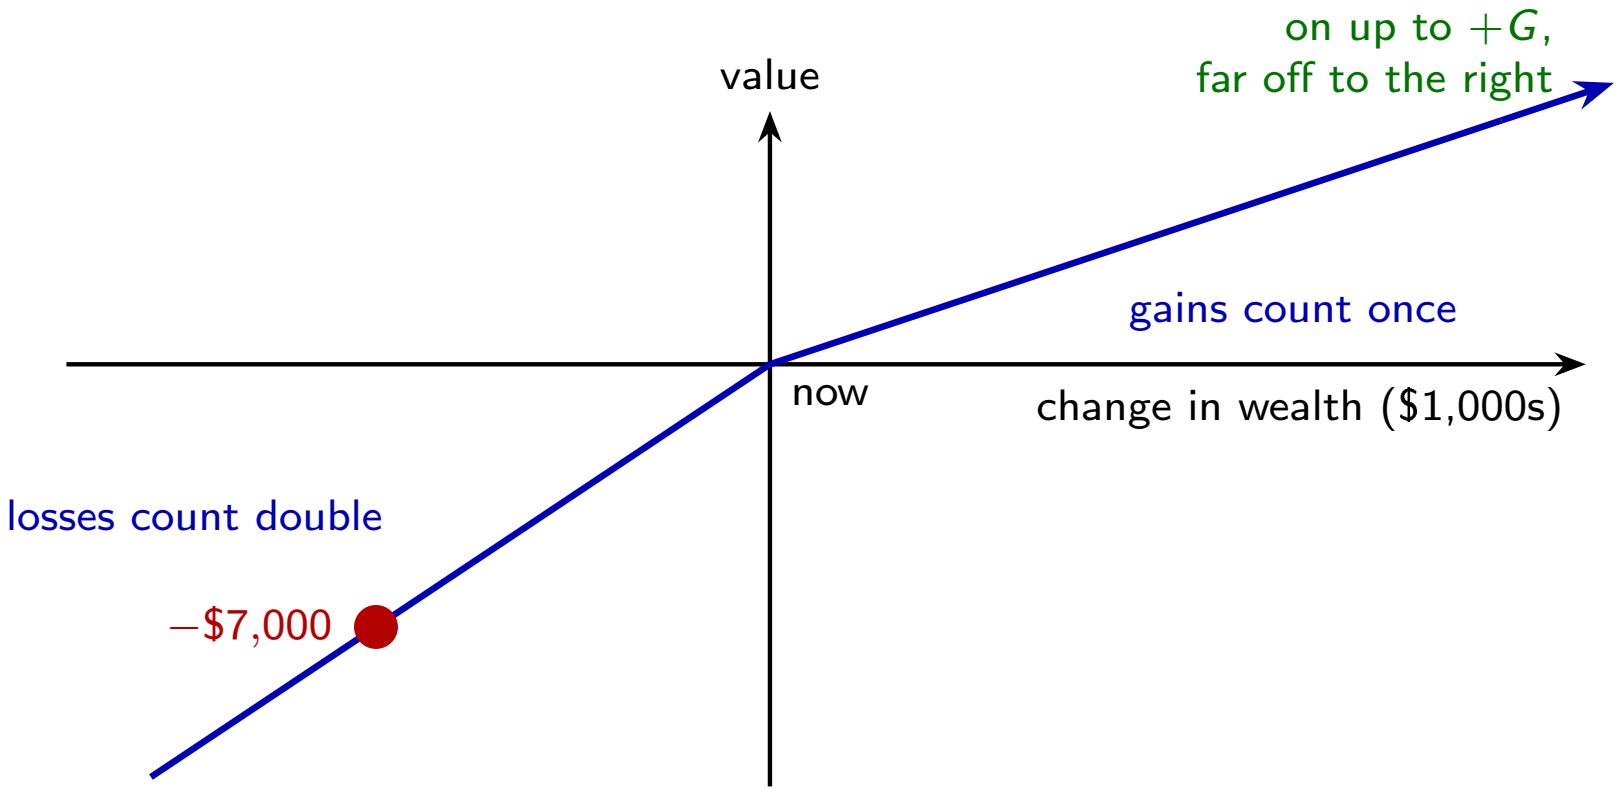
What Rabin Concludes

Rabin (2000: 1282)

"From such observations we should conclude that aversion to modest-stakes risk has nothing to do with the diminishing marginal utility of wealth."

- Rabin's explanation is *loss aversion*: people are much more averse to losses relative to the status quo than they are attracted by gains.
- On this view, utilities depend on changes in wealth, not on absolute levels.

Loss Aversion at Work



Lose \$7,000 / gain G: $\frac{1}{2}(-14{,}000) + \frac{1}{2}G$.

That's positive whenever G is more than \$14,000, so the agent takes it.

References

Adler, Matthew D. and Nils Holtug (2019). "Prioritarianism: A Response to Critics". In: *Politics, Philosophy & Economics* 18.2, pp. 101–144.

Nebel, Jacob M. and H. Orri Stefánsson (2023). "Calibration Dilemmas in the Ethics of Distribution". In: *Economics and Philosophy* 39, pp. 67–98.

Rabin, Matthew (2000). "Risk Aversion and Expected-Utility Theory: A Calibration Theorem". In: *Econometrica* 68.5, pp. 1281–1292.