PE3101P: Decision and Social Choice

Lecture Notes 9 — Calibration

Instructor: Tomi Francis

15 October 2026

1 Introduction

1.1 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. There are 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.

1.2 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.

Figure 1 draws wellbeing against money. Each extra dollar adds less wellbeing than the one before. So some extra money for someone poorer gives a big rise in wellbeing, and the same money for someone richer gives a small rise. It does more for total wellbeing if the poorer person gets it.

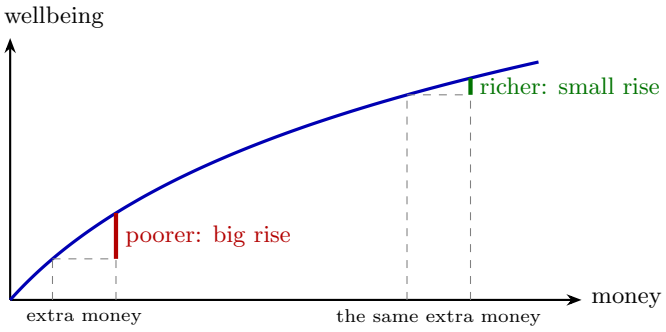


Figure 1: The same money at two wealth levels.

Many utilitarians think this explains our equality-favouring judgements about resources, without our needing to say that equality matters in itself.

1.3 The calibration problem

Roughly, the calibration problem is this: 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 (week 8, section 1.1).

So the calibration problem bites in three places:

(1) **Risk:** risk aversion, in preferences under risk.

(2) **Utilitarianism:** the diminishing value of money.

(3) **Egalitarianism, prioritarianism and views like them:** priority weighting.

It's an issue for everybody, not just for egalitarians and prioritarians.

2 Calibration and risk

2.1 Two gambles

Would you take these two gambles?

(1) A 50–50 gamble: gain \$1,100 or lose \$1,000.

(2) A 50–50 gamble: gain \$1,000,000,000 or lose \$7,000.

The rest of this section shows that for someone who maximises expected utility, with a concave utility function, turning down the first gamble at every wealth level means turning down the second.

2.2 Zooming in on a concave curve

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. Figure 2 shows this for utility = $\sqrt{\text{wealth}}$, looking more and more closely at the curve around \$10,000. So for small gambles, someone with a concave utility function is almost risk neutral.

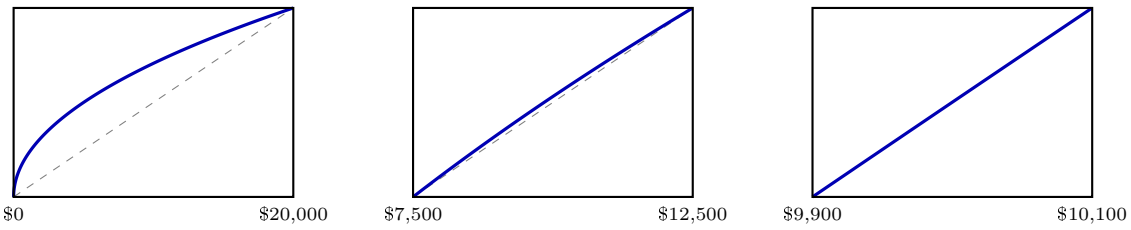


Figure 2: Utility = $\sqrt{\text{wealth}}$, from \$0 to \$20,000 (clearly curved), from \$7,500 to \$12,500, and from \$9,900 to \$10,100 (almost a straight line). In each picture, the dashed line is the straight line across it.

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.

2.3 Square root at \$10,000

Someone whose utility is $\sqrt{\text{wealth}}$ has \$10,000, which gives them 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. (Its expected utility is

half the utility of winning plus half the utility of losing, so it beats keeping $10,000 exactly when the gain in utility outweighs the loss.) Table 1 shows two gambles.

Table 1: Two 50–50 gambles, for someone with $10,000 whose utility is $\sqrt{\text{wealth}}$.

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

Losing $1,900 leaves $8,100, and $8{,}100 = 90^2$, so their utility would be 90. Gaining $2,000 gives $12,000, which is less than $12{,}100 = 110^2$, so their utility would be less than 110. So the loss outweighs the gain, and they turn down lose $1,900 / gain $2,000.

Over $1,000, though, the curve is almost straight. Gaining $1,100 beats losing $1,000, so they take lose $1,000 / gain $1,100. Square root isn't risk averse enough to turn down small gambles like this.

2.4 A utility function made of straight pieces

Suppose someone turns down lose $1,000 / gain $1,100 at every wealth level. Figure 3 shows one utility function that does that, made of straight $1,000 pieces. Below the curve is the slope of each piece: how much utility each extra dollar adds at that point. Every $1,000, the slope is multiplied by 0.9.

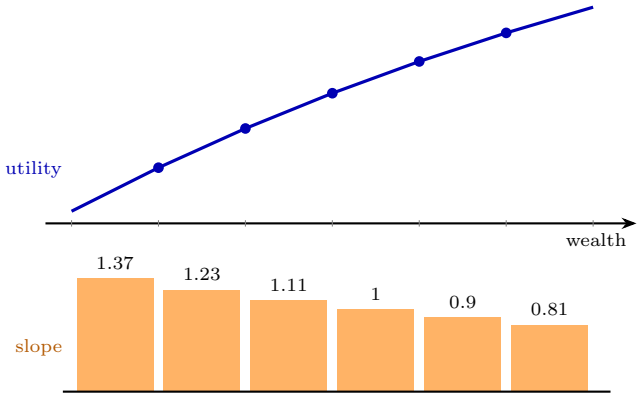


Figure 3: A utility function made of straight $1,000 pieces (above), and the slope of each piece (below).

2.5 Diminishing marginal utility

Figure 4 shows the slope again, with real wealth levels. Say you are the person with this utility function, and you have $50,500, in the middle of a piece. What the $1,000 below you is worth to you is the area under the slope over that $1,000: the red area. What the $1,000 above you is worth is the green area. It's the red area moved up $1,000, where every slope is 0.9 as big. So it's worth 0.9 of the red.

What about the last $100 of the $1,100 gain? The slope never rises as wealth goes up, so each $100 of the red area is worth at least as much as the $100 just above you. There are ten of them, so the $100 just above you is worth at most a tenth of the red. The $100 that takes the gain from $1,000 to $1,100 is that $100 moved up $1,000, so it is worth 0.9 of it: at most 0.09 of the red.

So gaining $1,100 is worth at most $0.9 + 0.09 = 0.99$ of losing $1,000, and you turn the gamble down. Nothing here depended on where you start. 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.

Zoomed out, as in Figure 5, the slope keeps climbing as wealth falls: the less wealth, the higher and higher the slope.

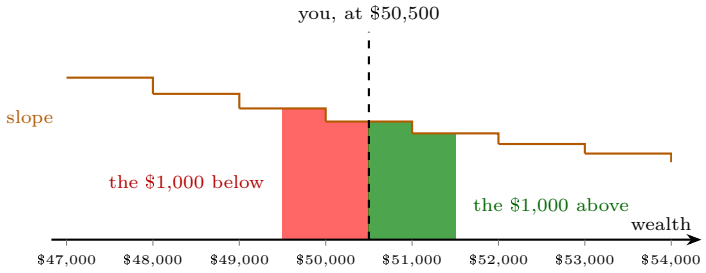


Figure 4: The slope, with you at $50,500. Red: what the $1,000 below you is worth. Green: what the $1,000 above you is worth.

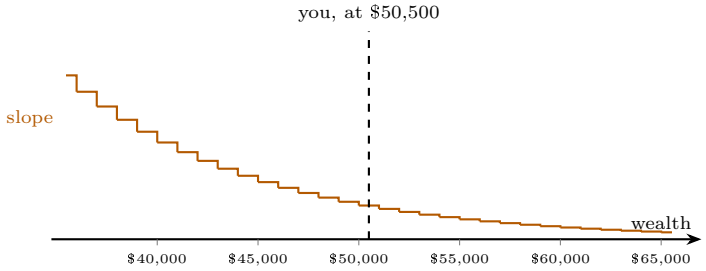


Figure 5: The same slope, zoomed out, from about $35,000 to $65,000.

2.6 Seven steps halve the slope

Start anywhere, and call the slope there 1. With $1,000 more, it is multiplied by 0.9. With $2,000 more, it is multiplied by 0.9 again: $0.9^2 = 0.81$. And so on, with every $1,000 more (Figure 6). After seven steps, with $7,000 more, the slope is $0.9^7 \approx 0.48$, less than half. So with every $7,000 more, the slope at least halves.

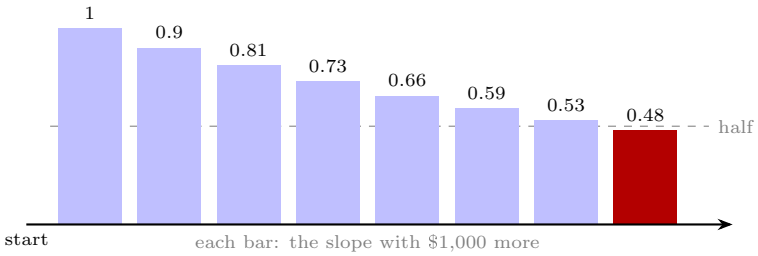


Figure 6: Seven steps halve the slope.

2.7 Halving, again and again

Now divide wealth into blocks of $7,000, starting at your wealth w. In Figure 7, each bar's height is what that $7,000 block is worth to you, compared with the $7,000 just above w, which we call 1. Every dollar in the next $7,000 up is worth less than half the dollar $7,000 below it, so that block is worth at most $\frac{1}{2}$. By the same reasoning, the next block up is worth at most $\frac{1}{4}$, and so on, all the way up. The other way, every dollar in the $7,000 just below w is worth more than double the dollar $7,000 above it, so that block is worth more than 2.

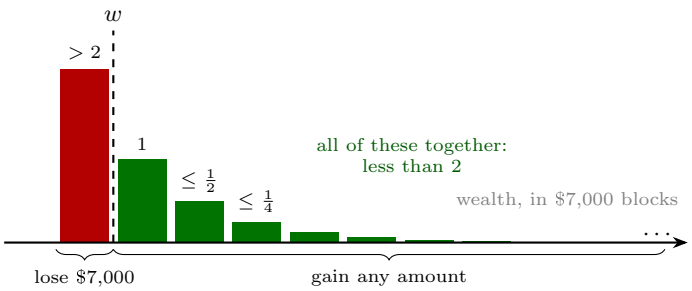


Figure 7: What each $7,000 block is worth, compared with the $7,000 just above w.

So losing $7,000 costs more than 2. Gaining any amount, however large, is worth at most $1+\frac{1}{2}+\frac{1}{4}+\ldots$, which stays below 2, however many blocks we add. So the loss outweighs any possible gain, and, as in section 2.3, you turn down 50–50 lose $7,000 / gain G, however large G is. That includes $G = \$1{,}000{,}000{,}000$, a gamble with an expected gain of $\frac{1}{2} \cdot \$1{,}000{,}000{,}000 - \frac{1}{2} \cdot \$7{,}000 = \$499{,}996{,}500$.

2.8 Calibration, more generally

Our step function is just one example. Rabin (2000) showed that the same kind of conclusion holds for every concave utility function that turns down 50–50 lose $1,000 / gain $1,100 at every wealth level.[1]

3 Calibration and distribution

3.1 Risk and distribution

Last week we saw that states of nature play the role that people play (week 8, section 1.1). Table 2 lines up the two sides for this week's argument.

Table 2: Risk and distribution.

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.

3.2 The same problem for money

Take 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?

Take a utilitarian who never takes $1,000 from one person to give $1,100 to another who is equally well off. One way wellbeing could depend on money is in straight $1,000 pieces: Figure 3 again, with wellbeing in place of utility and money in place of wealth. Now the slope is how much wellbeing each extra dollar adds at that point, and every $1,000 it is multiplied by 0.9.

[1] Rabin, Matthew (2000). "Risk Aversion and Expected-Utility Theory: A Calibration Theorem". In: *Econometrica* 68.5, pp. 1281–1292. In his words, someone who turns down lose $100 / gain $110 from any initial wealth level "will turn down 50-50 bets of losing $1,000 or gaining any sum of money" (p. 1282).

Say the two people each have \$50,500, in the middle of a piece (Figure 8). The red area is the wellbeing one person loses if we take \$1,000 from them. The green area is the wellbeing the other gains from \$1,000 more. It's the red area moved up \$1,000, where every slope is 0.9 as big, so it's worth 0.9 of the red. So, by the same reasoning as in section 2.5, \$1,100 more for one adds at most $0.9 + 0.09 = 0.99$ of what \$1,000 less costs the other. That's the drop-off needed for the utilitarian never to take \$1,000 from one to give \$1,100 to the other. Zoomed out, as in Figure 5, the less money, the higher and higher the slope.

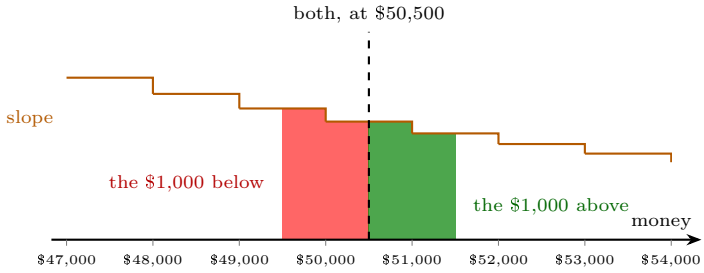


Figure 8: Wellbeing per dollar, with both people at \$50,500. Red: the wellbeing one loses from \$1,000 less. Green: the wellbeing the other gains from \$1,000 more.

As in section 2.6, after seven steps of \$1,000 the wellbeing per dollar is multiplied by $0.9^7 \approx 0.48$, so with every \$7,000 more, the wellbeing per dollar at least halves.

Now let w be the two people's common wealth, and divide money into blocks of \$7,000 from w (Figure 9). Compared with the \$7,000 just above w, which we call 1, the next \$7,000 up adds at most $\frac{1}{2}$, the next at most $\frac{1}{4}$, and so on; the \$7,000 just below w adds more than 2. So one person losing \$7,000 costs more than 2, and the other gaining any amount, however large, adds at most $1 + \frac{1}{2} + \frac{1}{4} + \ldots$, which stays below 2, however many blocks we add.

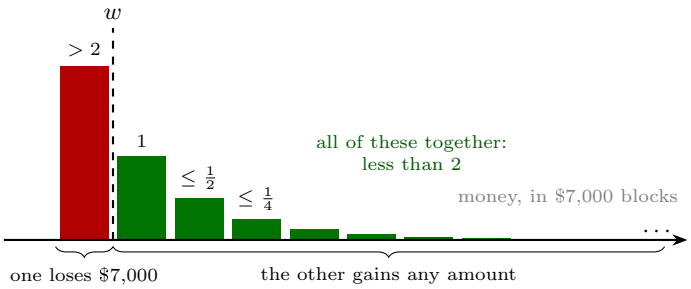


Figure 9: The wellbeing each \$7,000 block adds, compared with the \$7,000 just above w.

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. Table 3 sets the two arguments side by side.

Table 3: 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

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

3.3 Priority weights in straight pieces

The *priority weight* at a level of wellbeing is 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, by the same reasoning as in section 2.5, this prioritarian prefers (w, w) to $(w - 1{,}000,\ w + 1{,}100)$, at every level w.

Again, seven steps halve the priority weight. Start at any level, and call the priority weight there 1. With 1,000 more units it is 0.9; with 2,000 more units, $0.9^2 = 0.81$; and so on. After seven steps, with 7,000 more units, it is $0.9^7 \approx 0.48$, less than half. So with every 7,000 more units, the priority weight at least halves (Figure 10).

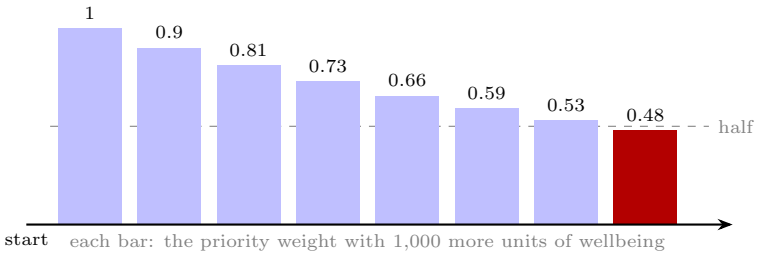


Figure 10: Seven steps halve the priority weight.

Figure 11 shows what each block of 7,000 units counts for, compared with the 7,000 units just above w, which we call 1. The next 7,000 units up count for at most $\frac{1}{2}$, the next for at most $\frac{1}{4}$, and so on, all the way up. The other way, the 7,000 units just below w count for more than 2.

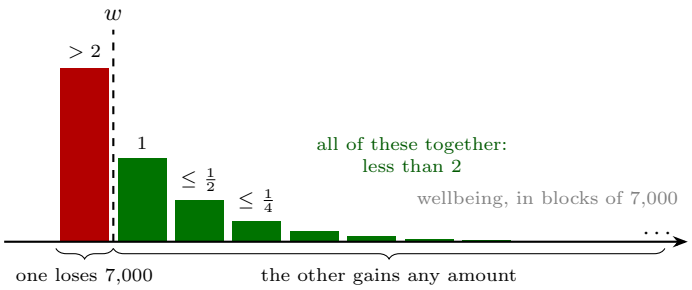


Figure 11: What each block of 7,000 units of wellbeing counts for, compared with the 7,000 units just above w.

So one person losing 7,000 units costs more than 2, and the other gaining any amount, however large, adds at most $1 + \frac{1}{2} + \frac{1}{4} + \ldots$, which stays below 2, however many blocks we add. So (w, w) beats $(w - 7{,}000,\ w + G)$, however large G is.

3.4 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. The priority-weighted sum is the sum of what each pair contributes, and each pair at (w, w) contributes more than a pair at $(w - 7{,}000,\ w + G)$. So $2n$ people at w beat n people at

$w - 7{,}000$ and n at $w + G$.

As with risk, the same conclusion holds for any concave priority weighting. Nebel and Stefánsson (2023) make this point about prioritarianism.[2]

4 Gini egalitarianism

4.1 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 (week 8, section 4.5). 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.

4.2 The Gini coefficient

Definition 1. Let S be the sum of the gaps $|w_i - w_j|$ over all i and j, where w_i is person i's wellbeing. The *Gini coefficient* of an outcome is

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

where n is the number of people and w_A is the average wellbeing.

First add up the gaps between every pair of people, each pair counted both ways: that sum is S. Each person is also paired with themselves, with a gap of 0, so there are n^2 pairs, and S/n^2 is the average gap. Then divide by twice the average wellbeing, $2w_A$, so that the scale doesn't matter. The Gini coefficient is 0 when everyone is equal, and close to 1 when one person has nearly everything.

For example, take the outcome $(2, 10, 20)$. The total is 32, and the average is $w_A = \frac{32}{3}$. Table 4 gives the gap between each pair, counted both ways.

Table 4: The gaps in $(2, 10, 20)$.

Gap	Person 1	Person 2	Person 3
Person 1	0	8	18
Person 2	8	0	10
Person 3	18	10	0

So $S = 72$. Divided by $n^2 = 9$, the average gap is 8. Divided by twice the average wellbeing, $2w_A = 2 \times \frac{32}{3} = \frac{64}{3} \approx 21.3$, that gives a Gini coefficient of $8 \div \frac{64}{3} = \frac{24}{64} = 0.375$.

4.3 Gini egalitarianism

Definition 2. *Gini egalitarianism* ranks outcomes by total wellbeing $\times$ (1 $-$ Gini coefficient), that is, by

$$\text{total wellbeing} \times \Big(1 - \frac{S}{2n^2 w_A}\Big).$$

(It is for positive wellbeing levels.) Call this number the outcome's *Gini value*.

Table 5 compares two outcomes with the same total. In UNEQUAL, the only gap is 12, counted twice, so $S = 24$; with $n = 2$ and $w_A = 10$, the Gini coefficient is $\frac{S}{2n^2 w_A} = \frac{24}{2 \times 4 \times 10} = 0.3$.

[2] Nebel, Jacob M. and H. Orri Stefánsson (2023). "Calibration Dilemmas in the Ethics of Distribution". In: *Economics and Philosophy* 39.1, pp. 67–98.

Table 5: Gini egalitarianism at work.

	Outcome	Total	Gini coefficient	Gini value
UNEQUAL	$(4, 16)$	20	0.3	$20 \times 0.7 = 14$
EQUAL	$(10, 10)$	20	0	$20 \times 1 = 20$

Utilitarianism says the two are equally good. Gini egalitarianism says EQUAL is better.

Gini egalitarianism satisfies Pigou–Dalton and Continuity (week 8, section 4.7), but not Separability over Persons.

4.4 Gini as rank weighting

Here is a fact we won't prove here: an outcome's Gini value is a weighted sum of the people's wellbeing levels. You will prove it yourself in Assignment 2, Question 4. Line people up from best off to worst off. Weight the best-off person by $\frac{1}{n}$, the next one by $\frac{3}{n}$, the next by $\frac{5}{n}$, and keep going until the worst-off person is weighted just below 2, at $\frac{2n-1}{n}$. So the weights are just the odd numbers $1, 3, 5, \ldots$, each divided by the number of people:

$$\frac{1}{n}, \frac{3}{n}, \frac{5}{n}, \ldots, \frac{2n-1}{n}.$$

With three people they are $\frac{1}{3}$, $\frac{3}{3}$ and $\frac{5}{3}$; with four, $\frac{1}{4}$, $\frac{3}{4}$, $\frac{5}{4}$ and $\frac{7}{4}$.

The weight goes with your *rank*, not with you: the fewer people worse off than you, the bigger your weight.

Figure 12 works this out for $(2, 10, 20)$. To keep it simple, it uses the weights 5, 3 and 1, and divides by 3 at the end. The weighted sum is $60 \div 3 = 20$. That is the same as total wellbeing $\times$ (1 − Gini coefficient): $32 \times (1 - 0.375) = 20$.

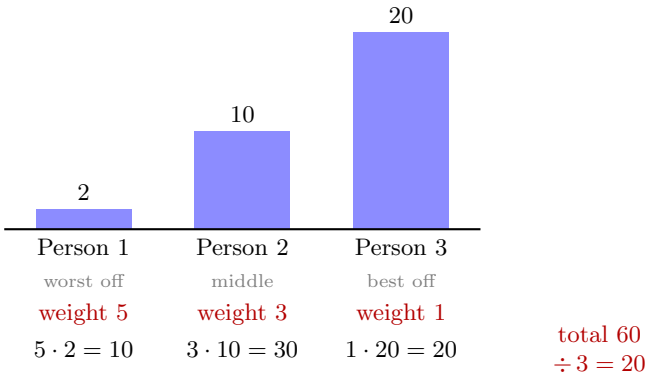


Figure 12: Weighting by rank.

4.5 Picturing the Gini coefficient

Line people up, worst off first. In Figure 13, the horizontal axis gives a share of the people, and the vertical axis gives their share of total wellbeing. Perfect equality is the diagonal. In $(2, 10, 20)$, out of 32, person 1 has $\frac{2}{32}$ of the total, and persons 1 and 2 together have $\frac{12}{32}$. The curve through these points is the *Lorenz curve*.

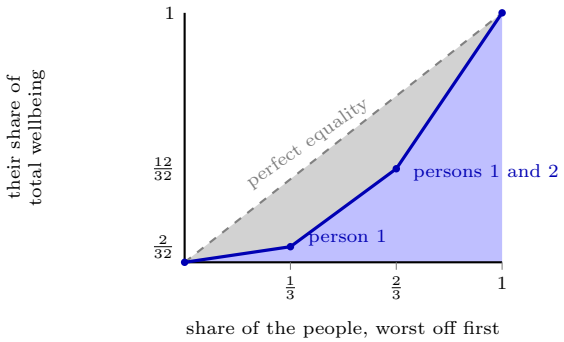


Figure 13: The Lorenz curve for $(2, 10, 20)$. Grey: the gap between the diagonal and the Lorenz curve. Blue: the area under the Lorenz curve.

The Gini coefficient is the grey gap between the diagonal and the Lorenz curve, as a fraction of the whole triangle under the diagonal. The area under the Lorenz curve is 0.3125, so the grey gap is $0.5 - 0.3125 = 0.1875$, and $0.1875 \div 0.5 = 0.375$.

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.

4.6 Gini and separability

Table 6 gives four outcomes. To compare them, it is enough to use the weights 5, 3 and 1, from worst off to best off: that gives 3 times the Gini value, so the same ranking.

Table 6: Gini and separability. Score: weights 5, 3, 1 from worst off to best off.

	Outcome	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 person 3 at 1, X beats Y. With person 3 at 20, Y' beats X'. Persons 1 and 2 are just as before: only person 3, who is unaffected, differs. What changes is their ranks: with person 3 at 1, person 3 is the worst off, so persons 1 and 2 get the weights 3 and 1; with person 3 at 20, they get 5 and 3. Separability over Persons would need the same verdict both times.

4.7 Trade-offs between two people

When a change moves wellbeing between just two people, and nobody changes place in the line, there is a quick way to see what Gini egalitarianism says. Everyone else's term in the weighted sum stays the same, so all that matters is the ratio of the two people's weights. If the worse off of the two has weight a and the better off has weight b, a loss of L to the worse off is outweighed by a gain of G to the better off exactly when $b \cdot G > a \cdot L$, that is, when $G > \frac{a}{b} \cdot L$.

Table 6 shows this at work. Going from X to Y, person 1 loses 4 and person 2 gains 8. With person 3 at 1, persons 1 and 2 have the weights 3 and 1, a ratio of 3, so the gain would have to be more than $3 \times 4 = 12$: X is better. With person 3 at 20, they have the weights 5 and 3, a ratio of $\frac{5}{3}$, so the gain only has to be more than $\frac{5}{3} \times 4 \approx 6.7$: Y' is better.

4.8 Does Gini egalitarianism escape the calibration dilemma?

With two people, the weights are $\frac{3}{2}$ and $\frac{1}{2}$, so an outcome's Gini value is

$$\frac{3 \times \text{the worse off's wellbeing} + 1 \times \text{the better off's wellbeing}}{2}.$$

Table 7 gives three outcomes.

Table 7: Does Gini egalitarianism escape?

Distribution	Gini value
(20,000, 20,000)	$\frac{3\times 20{,}000+20{,}000}{2} = 40{,}000$
(19,000, 22,000)	$\frac{3\times 19{,}000+22{,}000}{2} = 39{,}500$
(13,000, $20{,}000+G$)	$\frac{3\times 13{,}000+20{,}000+G}{2} = 29{,}500+\frac{G}{2}$

Equality beats one person losing 1,000 and the other gaining twice as much. So it beats one losing 1,000 to give the other 1,100. The loss counts three times and the gain once, whatever the level, so this holds at every level: it is the premise of the argument with numbers.

But one person losing 7,000 and the other gaining G beats equality as soon as G is more than 21,000, since $29{,}500 + \frac{G}{2} > 40{,}000$ exactly when $G > 21{,}000$. So Gini egalitarianism seems to escape the calibration problem.

4.9 Gini with other people around

Now add ten more people, each at 50,000. With twelve people, the two worst off get the weights $\frac{23}{12}$ and $\frac{21}{12}$. Table 8 compares the two worst off's part of the Gini value, multiplied by 12.

Table 8: Gini with ten more people, each at 50,000.

The two worst off	Their part of the Gini 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: here it doesn't even satisfy the premise, since its aversion to small inequalities depends on who else is around. So the escape is not robust.

5 Calibration without numbers

5.1 Small inequalities aversion

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 (Table 9). Giving \$2,000 to the richer person adds more money in total, but leaves the two more unequal.

Table 9: $1,000 to the poorer, or $2,000 to the richer.

	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

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: everyone else, who gets nothing either way.

5.2 A chain

Take eleven people, A to K, $1,000 apart in wealth: A has $30,000, B has $31,000, and so on up to K with $40,000. Each option gives *one* person a gift, and each gift is double the last: $1 to A, $2 to B, $4 to C, and so on, up to $512 to J; K's gift is $1,000 (Figure 14).

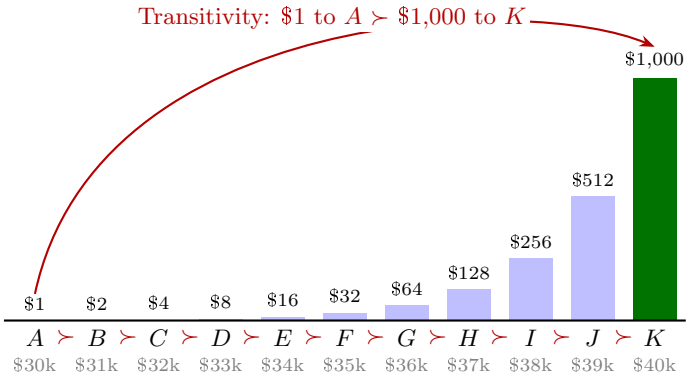


Figure 14: A chain. Each bar is that person's gift, to scale; below each letter, their wealth.

B has $1,000 more than A, so by small inequalities aversion, $1 to $A \succ$ $2 to B. The same holds between B and C: $2 to $B \succ$ $4 to C. And between C and D: $4 to $C \succ$ $8 to D. And so on up the line, to $512 to $J \succ$ $1,000 to K. In each of these comparisons, everyone else in the line gets nothing either way: they are the background population.

By Transitivity, giving $1 to A is better than giving $1,000 to K. But it seems that the opposite is true.

5.3 The same chain in wellbeing

Now let the amounts be units of wellbeing. The eleven people are 1,000 units apart, and each gift is double the last. Small inequalities aversion, now about wellbeing, gives 1 to $A \succ$ 2 to B, then 2 to $B \succ$ 4 to C, then 4 to $C \succ$ 8 to D, and so on, to 512 to $J \succ$ 1,000 to K. By Transitivity, 1 unit to A is better than 1,000 units to K. So the same problem arises for prioritarians and egalitarians.

5.4 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.

6 Ways out

6.1 Four claims that can't all hold

Table 10 lists four claims and where each one comes into the chain. The last of them is *non-extreme larger inequalities aversion*: a big enough benefit to someone better off can outweigh a small benefit to someone worse off.

Table 10: 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.

6.2 Rankings relative to option sets

One way out is to make the ranking relative to the option set. For example, we could 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 is to count only inequalities among those affected by the choice. Perhaps inequalities are only bad in themselves when they are avoidable.

6.3 Loss Aversion

Having worked through examples like those in section 2, Rabin draws this conclusion about risk:

> From such observations we should conclude that aversion to modest-stakes risk has nothing to do with the diminishing marginal utility of wealth.[3]

His 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.[4]

6.4 Loss aversion at work

A loss-averse agent values changes from where they are now. Say losses count twice as much as gains, as in Figure 15.

[3] Rabin (2000: 1282).
[4] Rabin (2000: 1288).

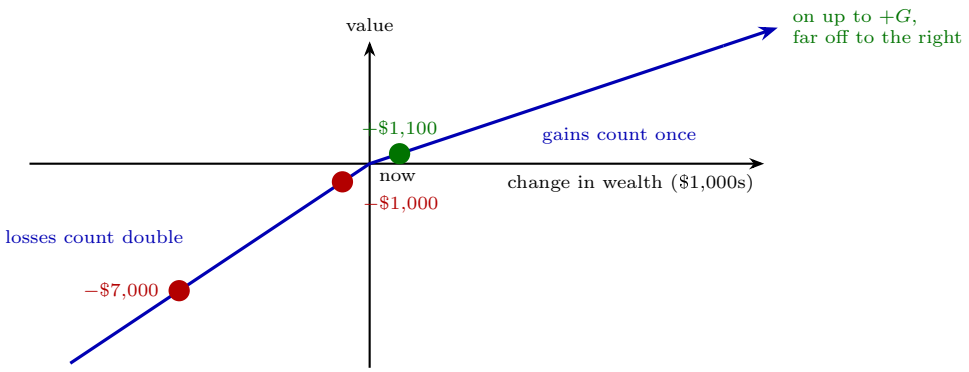


Figure 15: Loss aversion: losses count twice as much as gains.

For lose \$1,000 / gain \$1,100, the agent's expected value is

$$\tfrac{1}{2} \cdot (-2{,}000) + \tfrac{1}{2} \cdot 1{,}100 = -450,$$

so the agent turns it down. The kink in the line, where losses start counting double, is always at the current wealth level, so the agent turns this gamble down at every wealth level.

For lose \$7,000 / gain G, the gamble from section 2.7, it is

$$\tfrac{1}{2} \cdot (-14{,}000) + \tfrac{1}{2} \cdot G.$$

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

References

Nebel, Jacob M. and H. Orri Stefánsson (2023). "Calibration Dilemmas in the Ethics of Distribution". In: *Economics and Philosophy* 39.1, pp. 67–98.

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