Moral Decisions

Larry Blume

Cornell University & IHS Wien

2021-05-17

Who Should Get The Vaccine First

On Dec 3, 2020, the CDC Advisory Committee on Immunization Practices made a recommendation for COVID vaccine distribution.

https://www.cdc.gov/coronavirus/2019-ncov/vaccines/recommendations.html

Who Should Get The Vaccine First

On Dec 3, 2020, the CDC Advisory Committee on Immunization Practices made a recommendation for COVID vaccine distribution.

https://www.cdc.gov/coronavirus/2019-ncov/vaccines/recommendations.html

How would you decide this?

Who Should Get The Vaccine First

On Dec 3, 2020, the CDC Advisory Committee on Immunization Practices made a recommendation for COVID vaccine distribution.

https://www.cdc.gov/coronavirus/2019-ncov/vaccines/recommendations.html

They recommended:

Phase 1a) health care workers and long term care residents.

Phase 1b) Frontline essential workers and people over 75.

Who Should Get The Vaccine First

On Dec 3, 2020, the CDC Advisory Committee on Immunization Practices made a recommendation for COVID vaccine distribution.

https://www.cdc.gov/coronavirus/2019-ncov/vaccines/recommendations.html

What did states do?

Who Should Get The Vaccine First

On Dec 3, 2020, the CDC Advisory Committee on Immunization Practices made a recommendation for COVID vaccine distribution.

https://www.cdc.gov/coronavirus/2019-ncov/vaccines/recommendations.html

What did states do?

- 36 states followed the ACIP recommendations, 1 removed non-hospital health workers,
- 8 added fire and law enforcement personnel, 4 added correctional staff, 1 added long-term care staff and 1 added home health aides to the aged, 1 added frontline judicial staff, 1 added K-12 personnel, 1 state added "deployed and mission-critical personnel who play an essential role in national security."
- 3 added ages 65+, 1 added 75+, 1 added nursing home residents, 1 added 18+ non living independently

Discussion — Does the end justify the means?

- Consequentialists believe yes.
- Trolley car problem. The trolley car problem is not abstract:
  - If a jury can prevent riots that will cause many deaths only by convicting an innocent person of a crime and imposing a severe punishment on that person, should they convict and punish the innocent person?
  - If a doctor can save five people from death by killing one healthy person and using that person's organs for life-saving transplants, should she kill the one person to save five?

Welfarism

Consequentialism is the doctrine that states of affairs are judged by the consequences resulting from them.

Welfarism is the doctrine that "[t]he judgment of the relative goodness of alternative states of affairs must be based exclusively on, and taken as an increasing function of, the respective collections of individual utilities in these states." (Sen, 1979)

Utilitarianism is the doctrine that "[a]ny state of affairs x is at least as good as an alternative state of affairs y if and only if the sum total of individual utilities in x is at least as large as the sum total of individual utilities in y." (Sen, 1979)

A Critique of Welfarism

Peter and Paul must choose between baseline situation A and B which arises after redistributive taxation. The numbers are utilities.

	Peter	Paul	Total	Min
A	12	6	18	6
B	8	11	19	8

A simple utilitarian and a Rawlsian would both recommend the tax B.

A Critique of Welfarism

Peter and Paul must choose between baseline situation A and B which arises after redistributive taxation. The numbers are utilities.

	Peter	Paul	Total	Min
A	12	6	18	6
B	8	11	19	8

A simple utilitarian and a Rawlsian would both recommend the tax B.

Peter is rich and Paul is poor. Peter will not be able to buy his new speedboat, but Paul will be able to buy a house.

A Critique of Welfarism

Peter and Paul must choose between baseline situation A and B which arises after redistributive taxation. The numbers are utilities.

	Peter	Paul	Total	Min
A	12	6	18	6
B	8	11	19	8

A simple utilitarian and a Rawlsian would both recommend the tax B.

Peter is poor and Paul is rich. Peter will be able to buy a house, but Paul will be able to buy an ocean-going yacht.

A Critique of Welfarism

Peter and Paul must choose between baseline situation A and B which arises after redistributive taxation. The numbers are utilities.

	Peter	Paul	Total	Min
A	12	6	18	6
B	8	11	19	8

A simple utilitarian and a Rawlsian would both recommend the tax B.

The information contained in utilities alone is not enough to determine how we feel about the two states.

What is Utility?

- self-satisfaction?
  The preference relation that describes you behavior in Wegmans: What food and drink you like, the tradeoff between another frozen pizza and more beer, etc.
- deliberative preferences?
  I might say, I think extreme wealth inequality is a bad thing. At the same time I might spend money on something insignificant rather than give it away. So deliberative preferences and self-satisfaction preferences can be in conflict.
- the good?
  $\succ$ might rank acts not by preference but by goodness.

Act utilitarianism vs rule utilitarianism. RU rules out self-satisfaction.

Quality Adjusted Life Year

QALYs are estimated by assigning every life-year a weight between 0 and 1, where a weight of 0 reflects a health status that is valued as equal to being dead and a weight of 1 represents full health. For example, consider a patient with a colon cancer who has a health state of 0.9. Without surgery, he will die in 2 years. With surgery, his health state deteriorates slightly to 0.7, but he lives 5 more years (in total 7 years). Therefore, the QALY gained with surgery is

$$\Delta \text{QUALY} = [(0.7 \times 7) - (0.9 \times 2)] = 4.9 - 1.8 = 3.1.$$

Why It's Hard to be Welfarist

Society has N citizens choosing over least three alternatives. A social welfare function (SWF) assigns to each profile of utility functions $u = (u_1, \ldots, u_N)$ an aggregate or social preference relation.

- The domain of the SWF F is all possible utility profiles.
- A SWF respects unanimity if it ranks A above B whenever every citizen ranks A above B.
- It respects independence of irrelevant alternatives if the ranking of A and B depends only on the citizens' rankings of A and B.
- A SWF is a dictatorship by citizen n if $u_n(a) > u_n(b)$ implies society prefers a to b.

Arrow's Theorem. Any SWF that respects universal domain, independence of irrelevant alternatives, and unanimity is a dictatorship.

Proof of Arrow's Theorem

Lemma. Choose an arbitrary alternative b. In every profile where each citizen puts b either at the very top or the very bottom of her ranking, society puts b either at the top or at the bottom of its ranking.

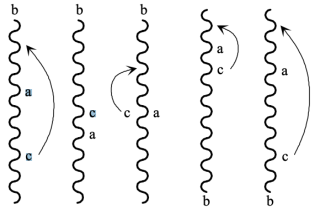

Proof of the lemma. Suppose for some profile u and alternatives a and c we have $a \succeq b \succeq c$. Now take a new utility profile u' which moves c up in everybody's rankings so that c beats a and such that if $u_n(b) > u_n(c)$, this individual ranking is preserved. The b, c rankings in both profiles are the same, and the a, b rankings in both profiles are the same. IIA implies $a \succeq_{u'} b$ and $b \succeq_{u'} c$, and unanimity implies that $c \succ_{u'} a$, so $\succ_{u'}$ is not a preference relation. ■

Proof of Arrow's Theorem

Start with a profile where b is at the very bottom of everyone's ranking. Starting with citizen 1, move b to the very top of each citizen's ranking. At some point, b will move from the very bottom to the very top of the social ranking. Label the citizen who caused the switch is $n(b)$. Call the profile just before $n(b)$ moved b up u, and call by v the profile just after he moved b up.

Citizen $n(b)$ must exist, because by unanimity b must be at the bottom of the social ranking before anyone moved b up, and at the top after everyone moved b up.

Proof of Arrow's Theorem

Claim: Citizen $n(b)$ is a dictator over all pairs of alternatives $a, c \neq b$.

Choose a from the pair a, c. Construct profile w from profile v by having $n(b)$ put a above b, so $u_{n(b)}(a) > u_{n(b)}(b) > u_{n(b)}(c)$, and letting everyone else arbitrarily arrange their rankings, leaving b in its extreme position. All individual a, b rankings are as they were in profile u, so IIA implies $a \succ_w b$. All b, c rankings are as they were in profile v, so IIA implies $b \succ_w c$. Transitivity of the social ranking implies that $a \succ_w c$. Now, by IIA, the social preference must agree with $n(b)$ whenever $u_{n(b)}(a) > u_{n(b)}(c)$.

Proof of Arrow's Theorem

Claim: Citizen $n(b)$ is a dictator over all pairs of alternatives a, b.

Perform the same exercise with alternative c to uncover citizen $n(c)$. She is a dictator over a, b. But $n(b)$ affected the ranking of a, b by changing her preferences at profile u to create profile v. Thus, $n(c) = n(b)$. ∎

Diagrammatic Intuition

This two-person example illustrated the force of the axioms. $X = R^2_+$. Choose A, B and C as in the figure. Suppose the social preference is $A \sim B$. Ordinality implies that $A \sim C$. Transitivity implies that $B \sim C$, which violates Pareto.

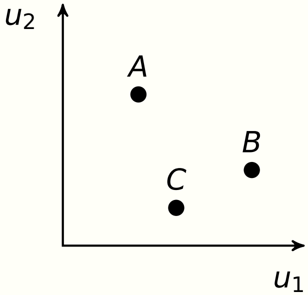

If the social preference is $A \succ B$, then whenever $u_2(x) > u_2(y)$, $x \succ y$, either because we are as in the picture, or if $u_1(x) > u_1(y)$, because of the Pareto axiom. Therefore 2 is a dictator. If $B \succ A$ then, for the same reason, 1 is a dictator.

Levels and Differences

- To talk about inequality in the distribution of welfare, one needs to compare individuals’ utility levels.
- To compare the aggregate welfare of different social states, one needs to compare only individuals’ utility differences.

Cardinal Utility

Cardinal utility requires that individuals be able to say things like, "I am made much better off by moving from a to b than by moving from c to d." We suppose individuals have a binary order on $X \times X$ where X is some set of choices. $(w, x) \rhd (y, z)$ means,

w is preferred to x more than y is preferred to z.

Cardinal Utility

The more than relationship has the following properties.

- $\rhd$ is asymmetric and negatively transitive on $X \times X$.
- If $(w, x) \rhd (y, z)$, then $(z, y) \rhd (x, w)$.
- If $(a, b) \rhd (a', b')$ and $(b, c) \rhd (b', c')$, then $(a, c) \rhd (a', c')$.
- An Archimedean axiom.
- If $(w, x) \unrhd (y, z) \unrhd (w, w)$, then there are a, b such that $(a, z) \equiv (y, z)$ and $(b, x) \equiv (y, z)$.

Theorem. If these five axioms are satisfied, then there is a function $u : X \to \mathbb{R}$ such that $(w, x) \rhd (y, z)$ iff $u(w) - u(x) > u(y) - u(z)$. If v also represents $\rhd$, then $v(x) = a + bu(x)$ with $b > 0$.

Quiz I

1. How would you derive a preference order $\succ$ on X from $\rhd$ on $X \times X$.

2. Show that your method works, that it delivers a preference relation.

Additive Social Welfare

Suppose that $u(x)$ and $v(x)$ measure social welfare in state of affairs x. Φ_u and Φ_v are the increasing transformations of u and v. Suppose we want an additive social welfare funtion: $W(u, v)(x) = u(x) + v(x)$. Can utility be ordinal? Cardinal? Something else?

What is the set Φ of transformations (ϕ_u, ϕ_v) of the utility functions u, v such that

$$u(x) + v(x) > u(y) + v(y)$$

$$\text{iff}$$

$$\phi_u \circ u(x) + \phi_v \circ v(x) > \phi_u \circ u(y) + \phi_v \circ v(y)$$

Additive Social Welfare

- Clearly Φ cannot be the set of increasing transformations — utility cannot be ordinal.

  Suppose $u(x) - u(y) > 0$ and $v(x) - v(y) < 0$. By choosing appropriate ϕ_u and ϕ_v, one can make either the Δu or the Δv dominate the sum.

- Clearly Φ cannot be the set of positive affine transformations — utility cannot be cardinal.

  The same thing can be done with positive affine transformations by appropriate choices of b_u and b_v.

Cardinal Utility & Social Welfare

One can make interpersonal utility comparisons with just cardinal preferences. Suppose utility differences are measured relative to a standard difference. Choose social states $p_n \not\sim_n q_n$, and define the welfare function

$$W(x) = \sum_n \frac{u_n(x)}{|u_n(q_n) - u_n(p_n)|}$$

If each $\phi_n \circ u = a_n + b_n u$, then

$$\sum_n \frac{\phi_n \circ u_n(x)}{|\phi_n \circ u_n(q_n) - \phi_n \circ u_n(p_n)|} = \sum_n \frac{a_n + b_n \circ u_n(x)}{|a_n + b_n u_n(q_n) - (a_n + b_n u_n(p_n))|}$$

$$= K + \sum_n \frac{u_n(x)}{|u_n(q_n) - u_n(p_n)|}$$

If the $b_n > 0$.

Arrow's Theorem and Cardinal Utility

Arrow's Theorem requires that if two utility profiles u and v are related by strictly increasing transformations, $v = \phi \circ u$, then $F(v) = F(u)$. What happens if u and v cardinal representations of the same "more preferred than" relation.

Cardinality. If two utility profiles u and v are related by positive affine transformations, $v = \phi \circ u$, then $F(v) = F(u)$.

IIA'. If u and v are utility profiles such that for all citizens n, $u_n(p) = v_n(p)$ and $u_n(q) = v_n(q)$, then the social ranking of a and b is the same for both profiles.

Theorem. Any SWF that satisfies universal domain, unanimity, IIA', and cardinality is a dictatorship.

Arrow's Theorem and Cardinal Utility

The proof shows that the IIA' and cardinality imply IIA. The conclusion follows from Arrow's Theorem.

IIA states the following: Suppose u and v are any two utility profiles, and suppose that p and q are two choices such that for each citizen n, $u_n(p) > u_n(q)$ iff $v_n(p) > v_n(q)$. Then the social rankings of p, q for profile u and for profile v are identical.

Suppose that for any pair p, q of alternatives one can find positive affine transformations ϕ_n such that $\phi_n \circ u_n(p) = v_n(p)$ and $\phi_n \circ u_n(q) = v_n(q)$. The social ranking of p, q for profile $\phi \circ u$ has to be identical to that of v by IIA'. Cardinality implies that the social ranking of p, q for profile u has to be identical to that of $\phi \circ u$. This would give us the Theorem.

Arrow's Theorem and Cardinal Utility

Here are the positive affine transformations:

If $u_n(p) = u_n(q)$, then $v_n(p) = v_n(q)$. Taking $a_n = v_n(p) - u_n(p)$ and $b_n = 1$ gives $\phi_n \circ u_n(p) = v_n(p)$, and $\phi_n \circ u_n(q) = v_n(q)$.

Suppose $u_n(p) > u_n(q)$ and (therefore) $v_n(p) > v_n(q)$. The equations

$$a + bu_n(p) = v_n(p)$$
$$a + bu_n(q) = v_n(q)$$

have the solution

$$a = \frac{u_n(p)v_n(q) - u_n(q)v_n(p)}{u_n(p) - u_n(q)}, \qquad b = \frac{v_n(p) - v_n(q)}{u_n(p) - u_n(q)} > 0.$$

■

Quiz II

Which of the cardinal Arrow assumptions does $W(x)$ from slide 19 violate?

Relaxing Other Assumptions

If the universal domain assumption is dropped, SWFs satisfying Arrow's other axioms can be found.

Suppose that X is ordered; e.g. X is a subset of the real line. Suppose that for any n there is an x_n such that if $y > x > x_n$ or if $y < x < x_n$ then $u_n(y) < u_n(x) < u_n(x_n)$. Such preferences are single-peaked.

Theorem. If a preference profile is single-peaked, the majority voting profile, $x \succ y$ iff $\#\{n : u_n(x) > u_n(y)\}$ is a preference relation, and majority voting satisfies all of Arrow's other axioms on the set of single-peaked profiles.

Quiz III

Prove the claims about majority voting.

Relaxing Other Assumptions

If the number of alternatives is 2 rather than 3 or more, majority voting (or any variant) satisfies all of Arrows other axioms.

Relaxing Other Assumptions

Suppose that the output of the social welfare function need only be transitive, but not negatively transitive.

Theorem. Arrow’s axioms imply that if negative transitivity is relaxed to transitivity of strict preference, then the SWF must map each preference profile onto its Pareto order.

Utility Differences

Let X be the set of social choices, and suppose utility functions u and v map X onto an interval of $\mathbb{R}$.

Theorem. The welfare comparisons of $W(x; u, v)$ and $W(x; \phi_u, \phi_v)$ are identical if $\phi_u \circ u = a_u + bu$ and $\phi_v \circ v = a_v + bv$.

This is interpersonal comparisons of utility differences. It entails:

- cardinal individual utility — admissible transforms are positive affine transformations;
- a common utility unit — $b_u = b_v$.

Utility Levels

Other social welfare functions require even more cardinality. The Rawlsian SWF (maximin SWF) is

$$W(x) = \min\{u(x), v(x)\}.$$

Theorem. Maximin welfare comparisons are preserved by the set of transformations Φ iff $\Phi = \{(\phi_u, \phi_v)\}$ such that $\phi_u \circ u = a + bu$ and $\phi_v \circ v = a + bv$.

Rawls requires interpersonal comparisons of utility levels. This entails cardinality, with a common utility unit and a common utility level: $b_u = b_v$ and $a_u = a_v$. This is like measuring length.

Summary

- No interpersonal comparisons — ordinal utility

$$\Phi_o = \{(\phi_u, \phi_v) : \text{ both are increasing.}\}$$

- Interpersonal comparisons with cardinal utility

$$\Phi_c = \{(\phi_u, \phi_v) : \phi_u \circ u = a_u + b_u u, \phi_v = a_v + b_v v\}$$

- Interpersonal comparisons of utility differences

$$\Phi_d = \{(\phi_u, \phi_v) : \phi_u \circ u = a_u + bu, \phi_v = a_v + bv\}$$

- Interpersonal comparisons of utility levels

$$\Phi_l = \{(\phi_u, \phi_v) : \phi_u \circ u = a + bu, \phi_v = a + bv\}$$

Is there a moral right to privacy?

The Problem with Pareto

The Pareto axiom appears to be consistent with liberal values. In fact, in some situations it can be deeply illiberal. The following is due to Sen (1970). Consider an Arrowvian SWF, and a new axiom, liberalism. A SWF respects liberalism if there is some group of at least two people, and one pair of alternatives for each person in that group, such that their preferences are decisive for social choice. That is, if n is in the group, then there are two alternatives x and y such that x is socially preferred to y iff $x \succ_n y$.

The Problem with Pareto

Theorem. There is no SWF satisfying universal domain, unanimity, and liberalism.

Proof. Suppose 1 is determinative for x, y and 2 is determinate for w, x. Suppose that $x \succ_1 y$ and $w \succ_2 x$, and that everyone holds $y \succ_n w$. Then the social order is $x \succ y$, $y \succ w$ (unanimity), and $w \succ x$ so preferences are cyclic.

Now suppose 1 is determinative for x, y and 2 is determinate for w, z (no overlap). Suppose $x \succ_1 y$, $w \succ_2 x$, and all hold $y \succ_n w$. Again preferences are cyclic.