"Infinite Value and Permissibility"

Routledge Encyclopaedia of Philosophy, Electronic Supplement, Revised Feb. 19, 2019

Puzzles can arise in value theory and deontic (permissibility) theory when infinity is involved. These puzzles can arise for ethics, for prudence, or for any normative perspective. For simplicity, we focus on the ethical versions of these problems. We start by addressing problems that can arise in determining what is permissible, either in a given choice situation when there are an infinite number of options or in infinite sequences of choices situations each with only finitely many options. A common theme: standard and plausible decision rules, such as the dominance principle or maximizing principles, run into conflicts with other plausible principles. We then address addresses problems that can arise in determining whether one option is more valuable than a second, when the options have infinite, or undefined, value. This includes cases where a single bearer of value has infinite or undefined value, as well as cases where each bearer of value is finite but the total for the collection is infinite or undefined totals. Our focus is on introducing the puzzles, rather than canvassing potential solutions.

1. Puzzles about Permissibility

We shall first address problems in determining what options are permissible, when there are infinite number of options in a single choice situation. Following that, we shall address problems in determining what options are permissible, when the agent confronts an infinite sequence of choice situations, each with only finitely many options. (See Infinity for a general introduction to the concept of infinity.)

Suppose that an agent has infinitely many options (possible choices) in a given choice situation, each of which has a finite value. To say that there are infinitely many options is just to say that there are more options than any finite number. The presence of infinitely many options does not automatically generate problems. For example, if one option has a value of 1, and all the others have a lower value, a maximizing theory (one that judges only maximally good feasible options permissible) can judge just the first option permissible. In other others situations, maximizing theories can run into problems. Suppose, for example, that the options are numbered, that o_1 has a value of 1/2, that o_2 has a value of 2/3, and that in general o_n has a value of n/(n+1). In this case, there is no option with a maximal value. The values are all finite and less than one, but for any option, o_n, there is another option with greater value (e.g., o_{n+1}). No option is maximally good, and thus no option is permissible according to a value optimizing theory.

The result that nothing is permissible is puzzling, but it can be avoided by replacing the optimization requirement with a requirement that a chosen option be at least as good as "trivially less" (on some specified criterion) than the best one can do. For example, if one billionth of a unit of goodness is the cut-off for being trivial, then, in the above example, there are infinitely many options that satisfy this requirement (and they are all "almost" maximal).

In the above case, there are infinitely many options, each option has a finite value, and the values of the options are *bounded* (there is some finite value—1 in this case—such that no option has a greater value). Things are not so simple when the values are not bounded. Suppose, for example, that the value of o_1 is 1, o_2 is 2, and in general o_n is n. Given that there are infinitely many options, there is no finite limit on how high the values can be (even though each option has a finite value). In this case, optimizing and "almost optimizing" theories say that no option is permissible. *Absolute satisficing* theories—that is, theories that judge an option permissible just in case its value is "good enough" on some specified absolute sense—have no problem with this case. Whatever the criterion for being good enough, there are infinitely many options that are permissible. (Of course, even when there are finitely-many options, it is possible that no feasible option is good enough.) People who are inclined to defend an optimizing or almost optimizing theory in the finite case thus either have to accept that nothing is permissible in such infinite cases (a strange claim) or to explain why satisficing is acceptable in the infinite case but not in the finite case. (One possibility is to hold that one should maximize when possible, that when this is not possible one should almost-maximize, and that when this is not possible one should satisfice.)

The above puzzles involved infinitely many options for an agent in a given choice situation. Related puzzles can arise when there are only a finite number of options at a given time but there are infinitely many future choice situations, because time extends infinitely into the future. Here let us suppose for simplicity that the value of an option is the value (e.g., happiness) that it produces in the world, and that for each time there is well defined "next" time. Furthermore, suppose for simplicity that there is a first time, and that at each time there is exactly one agent (either one agent who exists forever, or a new agent that comes into being when the previous agent dies). At each time, the agent has two options. One option is to produce a certain amount of value immediately, in which case no further value will be later produced in the world. The other option is to produce no value immediately, in which case, at the next time, the agent will have a choice between (1) producing significantly more value immediately and nothing thereafter and (2) producing no value immediately, but having a similarly structured choice situation at the next time. For example, the sequences of possible choices might look like this: <1 unit immediately vs. postpone>, <2 units immediately vs. postpone>, …. <n units immediately vs. postpone>, etc. (Assume here that any relevant discounting of value for temporal delays is already reflected in the numbers.) An optimizing, or almost optimizing, theory says that, at each time, the choice should be to postpone, but this will have the result that no value is ever produced!

These puzzles are called *discontinuous in the limit.* There is a series of choices, each of which has the potential to produce more value than the previous, but the limit strategy (always taking the choice that gives the option to produce more value) ends up producing less value than some of the other strategies for making choices.

In this case, a satisficing theory using an absolute criterion of being good enough will judge it permissible, at some point (and all later points), to produce immediate value. Such a theory, however, will also judge it permissible *at every point in time* to postpone the production of value, which seems puzzling. For in that case, no value is produced. One way of avoiding this problem (probably the only way) is to appeal to a rule-based, rather than an act-based, approach to deontic assessment. The idea is that, at a given time, the agent faces infinitely many rules or strategies that she could adopt and then comply with in the future. In the problem situation, the possible strategies are: never produce immediate value (i.e., postpone at each step), produce immediate value at time 1, produce immediate value at time 2, etc. If n units are the minimal satisfactory level, then all strategies that produce at least n units of value are permissible. The strategy of never producing immediate value is clearly not satisfactory and thus not permissible. This solves the puzzle, but is also raises questions about whether the permissibility of options is indeed based on the consequences of compliance with rules rather than directly on their consequences.

The previous puzzle arose in part because time was unbounded (infinitely long). Puzzles can also arise when time is bounded (finitely long) but infinitely divisible (or dense). Suppose that time is dense, which is to say that between any two points in time there is at least one other point in time, and consider the points in time between any two arbitrarily chosen points in time—between 0 and 1, say, as measured in complete days. Suppose further that the time has a metric (which measures the length of temporal intervals). For example, halfway between 0 and 1, there is .5 (half a day), and halfway between .5 and 1 there is .75, and so on. Suppose that at time 0 there are infinitely many one dollar bills that have been numbered using the natural numbers (1, 2, 3, etc.), and that no dollar bills later come into existence. Suppose that, at time 0, God has possession of all the dollar bills, but offers to transfer them to you by the following scheme: At each of the times, 1/2, 3/4, 7/8, … $(2^n-1)/2^n$, …, she will give you two arbitrarily chosen bills from those not yet given to you, and she will then destroy the lowest numbered bill in your possession since the last transaction. All bills in your possession between time 0 and time 1 remain in your possession, unless destroyed by God by the above process at a time prior to time 1. You are not permitted to use these bills until time 1, and nothing else of value is affected by this scheme. Is God's offer worth accepting? Does it increase the amount of money that you will have at time 1?

It may seem that God's offer will increase the amount of money that you have at time 1, but it will not. For when time 1 arrives, God will have destroyed every single dollar bill. This is so because (1) every dollar bill has a number, and (2) for any given dollar bill, no matter what its number, at some point prior to time 1, that number will be the lowest of the numbers of the bills in your possession. Hence, the bill will be destroyed prior to time 1. The specification of the problem entails that all the dollar bills are destroyed by time 1.

Still, this result is puzzling. Each transaction prior to time 1 increases the number of bills you have by one. Nonetheless, at time 1 you have no bills. Furthermore, very different results can be generated by seemingly trivial differences in the specification. Suppose, for example, that, at each of the specified times, God arbitrarily chooses two bills in her possession, destroys the bill with the smaller number of the two bills, and then gives you the remaining one. In this case, you will have infinitely many bills at time 1. It seems quite strange, however, that merely changing which bill is destroyed at each of the given times should have any effect on how much money you end up with. Depending on the background assumption, however, it can.

2. Puzzles about Betterness

The above puzzles arise when trying to determine what is permissible. Puzzles can also arise when trying to determine what is better than what (e.g., see Axiology). These can arise (1) when some of one good (e.g., being loved) is better than any finite amount of a different good (e.g., being amused), and (2) when there are an infinite number of locations of finite value (e.g., times)

For simplicity, we shall assume that the value of an option (or world) is, at least in the finite case, *additive* in the sense that the value of an option (or world) is the sum of values in the associated locations (e.g., see Utilitarianism is). Of course, many theories of value are not additive in this sense, but, for at least some of them, there may be related problems.

Can the value of some events or states (e.g., being loved) be infinitely greater in relative terms than the value of some other event or state (e.g., being amused)? Is it possible, that is, that there is no finite number of the latter events such that the value of all those events together is at least as great as the value of the former event? This is impossible, if all states and events have some standard finite value. One can, however, coherently reject this assumption.

First, value need not be representable by numbers. It may simply be ordinally representable by a ranking relation (i.e., as more, less, or equally valuable; but no assignment of specific numbers for value). All else being equal, more of the infinitely less valuable sources of value makes the world more valuable, but no finite number of such sources can ever compensate for the loss of one of the former sources of value. The infinitely more valuable sources of value are simply *lexicographically prior* in the generation of overall value to the infinitely less valuable sources of value. Thus, we can make perfect sense of the idea, if we do not require that value be numerically representable.

Second, even if one requires that numbers be assigned to the value of states of the world, this can be done using the infinitesimal numbers of non-standard arithmetic. In standard mathematics, there are no numbers that are infinitesimally small. In the 1960s, however, Abraham Robinson, a mathematician, proved that one can make perfect mathematical sense of such infinitesimals, and thus that it is legitimate to posit them. The addition of a positive infinitesimal to a given number produces a larger number, but the sum of finitely many infinitesimals still is still infinitesimally small, and hence smaller than any finite number (although greater than each of the original infinitesimals). If infinitesimals are recognized, then some sources of value may generate only infinitesimal value relative to other sources of value.

A second problem for the assessment of betterness arises when both options have infinite value. Suppose again that time extends infinitely into the future and that an agent has a choice between producing two units of value at each time or one unit of value at each time. Intuitively, it would seem that the former outcome is better than the latter outcome. The total value produced, however, is the same infinity in each case. Thus, if overall value is simply the sum of the values at each time (as we here assume for illustration), then it would seem that neither is better than the other is. This, however, seems strange, given that for every single time one option produces more value than the other does. One proposal for escaping this paralysis involves using non-standard, rather than standard, arithmetic. Here we will confine our discussion to hyperreal mathematics. Unlike in standard mathematics, in hyperreal mathematics, we can add and subtract finite numbers from infinite numbers and return different infinite numbers. For example, for any given hyperreal infinity, H, H+1 is also infinite and greater than H. Unfortunately, this attractive solution comes with some problems. Most notably: properties of the hyperreal field differ depending on the details of its construction. For example, as a result of one arbitrary choice to be made in the construction of the hyperreals, one can, for example, make there be more odd numbers than even numbers (more precisely: the set of odd numbers can be larger than the set of even numbers). Thus, in a world with infinitely many people numbered by the natural numbers and such that even numbered people are happy while odd-numbered people are sad, if we want to use hyperreals to model the aggregate value in that world, whether we have more happy or sad people will depend on an arbitrary choice we make in the construction of the hyperreals. This is an unfortunate result; whether it can be overcome depends on still-open questions in non-standard analysis.

An alternative approach is to appeal to a generalized dominance principle. This approach preserves the commitment to additivity in the finite case, but rejects it in the infinite case. In particular, it rejects the claim that two infinitely value options must be equally valuable. Instead, it holds that more at each time is better than less at each time (in both finite and infinite cases). The core idea (ignoring some variations) is that, if there is a time in the future such that at each later time the cumulative total value in one state of affairs is at least as great as that of the other, then the first state of affairs is at least as valuable as the second is. In our example, the cumulative total of the two-at-each-time state of affairs is always greater than that of the one-at-each-time, and hence it is judged as more valuable (even though both have infinite totals). The principle also allows that a more valuable state of affairs may be worse in the short run (e.g., have a lower cumulative total initially) as long as it eventually prevails.

This principle (and other related ones) gives the seemingly correct judgement about the puzzle case. It also satisfies a weak kind of dominance condition that seems central to finitely additive theory of values: if, at each location, one state of affairs is at least as valuable as a second state of affairs is, and more valuable at some locations, then it is more valuable. On the other hand, it violates a kind of neutrality condition that seems also to be central: if one state of affairs is identical to a second except that the values at locations have been placed in a different order (permuted), then the two states of affairs are equally valuable. The above principle satisfies the condition, when only finitely many locations have been reordered. It can, however, violate this condition, when an infinite number of locations are affected. For example, the principle judges <1,1,1,0,1,0,1,0,1,…> (one unit of value for the first three times and alternating zeros and ones thereafter) as better than <1,0,1,0,1,0,…> (alternating zeros and ones starting immediately). For, starting with the second time, the cumulative total of the first thereafter remains greater than that of the second. Hence, it is judged more valuable. The first sequence, however, is simply an infinite reordering of the second (obtained by moving all 0s two positions to the right, the first 1 one position to the left, and all remaining 1s two positions to the left). Hence, there is a genuine puzzle here as to what finitely additive theories of value should do in the infinite case.

A third problem with assessment of betterness arises when, for one or both options, the sum of the positive values and the sum of the negative values are both infinite. Consider a world with infinitely people with one unit of happiness and infinitely many people with one unit of unhappiness. Is this better than a world in which all the same people get zero units of happiness? The latter value is well defined and is (assuming an absolute zero) zero. The first world, however, has no well-defined total utility. This is because one gets different sums depending on the order in which the terms are added. For example, 1+1-1+1+1-1+1+1-1 … produces a positively infinite total, but -1-1+1-1-1+1-1-1+1… produces a negatively infinite total, even though the same terms are merely added together in a different order. (A second example of an undefined total is 1-1/2+1/3-1/4+… .) This case does not involve infinite value, but it does involve some similar issues.

One might suppose that, where two options (or worlds) each have undefined values, it is impossible to say whether one option is better than the other is. As with the case of infinite value, however, there are dominance reasons for supposing that, at least sometimes, one option has greater value than the other is. In the above example, suppose that each happy person has 2 units of happiness, rather than 1. Is this world better than a world in which each happy person has only 1 unit of happiness (and all the other people have -1)? Again, the total happiness is undefined (because the sum depends on the order of addition). Still, the 2-unit world dominates the 1-unit world in the sense that all the happy people are better off in the 2-unit world, and the unhappy people are unaffected. As with the infinite case, this dominance reasoning can be strengthened in certain ways to cover cases where there is no dominance (e.g., in the 1-unit world, increase one happy person's happiness by 2 units and decrease another person's happiness by 1 unit).

As we saw above, it is plausible that (1) some infinitely valuable options are more valuable than some other infinitely valuable options, and (2) some options with no defined value are more valuable than other options with no defined value. The core case is where one of them *strongly dominates* (is strongly Pareto superior) to the other in the sense that, at *every* location of value (e.g., person), the first option has more value than the second option. We now note, however, that this is incompatible with a standard form of anonymity (impartiality).

Consider, then, the following two conditions:

Weak Dominance (Weak Pareto): If, at every location, of value, one option has more value than another does, then the first option is more valuable than the second is.

Full Anonymity (Full Impartiality): If two options have the same pattern of distribution of values over locations (even if the values at specific locations are different), then they are equally valuable.

Consider now the following two options, O_1 and O_2, with their values listed at the various locations (l_1, l_2, etc.):

	…	l_{-n}	…	l_{-2}	l_{-1}	l_0	l_1	l_2	…	l_n	…
O_1	…	-n	…	-2	-1	0	1	2	…	n	…
O_2	…	-(n+1)	…	-3	-2	-1	0	1	…	n-1	…

Here, O_1 strongly dominates O_2 (has more value at every location). So, Weak Dominance requires that O_1 be more valuable than O_2. O_2, however, has exactly the same pattern of distribution as O_1 (for each n, there is one location with value n and one location with value –n). So, Full Anonymity requires that O_1 and O_2 be equally valuable. Hence, the two conditions are incompatible.

It is worth noting, without proof, that Dominance is fully compatible with a weakened version of anonymity:

Finite Anonymity (Finite Impartiality): If two options have the same pattern of distribution of values over locations, and there are only finitely-many locations at which they have different values, then they are equally valuable.

The two options above do not satisfy this condition, since there are infinitely-many locations at which they have different values.

Thus, those who endorse Weak Dominance must reject Full Anonymity, but they can accept Finite Anonymity. By contrast, those who accept Full Anonymity must reject Weak Dominance. Finding a good way to navigate this dilemma is one of the central challenges of infinite value theory.

3. Probability and Infinity

The puzzles addressed above do not involve appeals to probability. Related puzzles (e.g., Pascal's Wager, St. Petersburg Paradox, The Pasadena Game, and Two Envelopes Problem) can arise when probabilities are involved. Discussion of such puzzles is, however, beyond the scope of this article.

DANIEL RUBIO

PETER VALLENTYNE

See also: Infinity, Utilitarianism, Axiology

References and further reading

Arntzenius, F. (2014). "Utilitarianism, Decision Theory, and Eternity," *Philosophical Perspectives* 28: 31-58. (This presents a challenge to the use of hyperreal utilities stemming from unforced choices in the hyperreal construction.)

Arntzenius, F., A. Elga and J. Hawthorne (2004). "Bayesianism, Infinite Decisions, and Binding," *Mind* 113: 251–283. (This explores decision problems that are discontinuous in the limit and defends the 'binding' solution to them.)

Bartha, P., J. Barker and A. Hajek (2014). "Satan, St. Peter, and St. Petersburg," *Synthese* 191: 629-660. (This explores dynamical solutions to problems that are discontinuous in the limit and gives them their general characterization.)

Bostrom, N. (2011). "Infinite Ethics," *Analysis and Metaphysics* 10: 9–59. (This explores problems resulting from the presence of positive and negative infinities, and defends the use of hyperreals as a solution to them.)

Barrett J. and F. Arntzenius (1999). "An infinite decision puzzle." *Theory and Decision* 46: 101-103. (This introduces the final puzzle about dollar bills and infinitely divisible time.)

Chen, E. and D. Rubio (Forthcoming). "Surreal Decisions." *Philosophy and Phenomenological Research.* (This defends a non-standard utility theory using surreal numbers, which, unlike hyperreal numbers, do not suffer from arbitrariness in their construction, and shows how it can solve many of the puzzles resulting from infinite utilities in finite state spaces.)

Fine, T. (2008). "Evaluating the Pasadena, Altadena, and St. Petersburg Gambles," *Mind* 117:613-632. (Explores/explains the conflict between dominance, order-invariance, and the classical theory of infinite sums).

Hajek, A. and H. Nover (2004). "Vexing Expectations," *Mind* 113: 237-249. (This introduces the Pasadena Game and the problems of undefined value.)

Hamkins J. and B. Montero (2000). "With Infinite Utility, More Needn't Be Better," *Australasian Journal of Philosophy* 78(2): 231-240. (This argues against the principle identified in the text for assessing outcomes when the future is infinitely long. It also contains many references to recent work on this problem.)

Hurka, T. (2004). “Satisficing and Substantive Values,” in *Satisficing and Maximizing: Moral Theorists on Practical Reason*, ed. Michael Byron (ed.). Cambridge: Cambridge University Press. (This gives a good introduction to and critical discussion of satisficing.)

Keisler, H.J. (1976) *Elementary Calculus*, Boston: Prindle, Weber & Schmidt. (Ch. 1 of this mathematics book provides the best intuitive introduction to non-standard infinite and infinitesimal numbers.)

Landesman, C. (1995) "When to Terminate a Charitable Trust?" *Analysis* 55: 12-13. (This lays out the puzzle case where postponing gains makes greater gains possible and thus no choice is maximally valuable.)

L. Lauwers and P. Vallentyne (2004) “Infinite Utilitarianism: More Is Always Better”, *Economics and Philosophy* 20: 307-330. (This generalizes and simplifies the approach of Vallentyne and Kagan 1997.)

L. Lauwers and P. Vallentyne (2016) “Decision Theory without Finite Standard Expected Value”, *Economics and Philosophy* 32: 383-407. (This addresses problems of undefined value in the context of the Pasadena problem, but the core ideas are applicable even where there are probabilities involved.)

Nelson, M. (1991) “Utilitarian Eschatology,” *American Philosophical Quarterly* 28: 339-47. (This lays out the problem for utilitarianism and related ethical theories of assessing options when the future is infinitely long.)

Robinson, A. (1966) *Non-Standard Analysis*, Amsterdam: North Holland. (This is a technical mathematics book that lays out the author’s groundbreaking proofs that one can recognize the existence of infinitesimals.)

Slote, M. (1989) *Beyond Optimizing*, Cambridge, MA: Harvard University Press. (This defends the view that morality only requires satisficing.)

Sorensen, R. (1993) "Infinite Decision Theory," in J. Jordan (ed.) *Gambling with God*, Totowa, NJ: Rowman & Littlefield. (This highly accessible survey addresses some of the issues that arise in decision theory when payoffs are infinite or when they are infinitely many finite payoffs.)

Vallentyne, P. (1993) "Utilitarianism and Infinite Utility," *Australasian Journal of Philosophy* 71: 212-17. (This defends the principle identified in the text for assessing outcomes when the future is infinitely long. It also contains many references to earlier work by economists and philosophers on this problem.)

Vallentyne, P. and S. Kagan (1997) "Infinite Utility and Finitely Additive Value Theory," *Journal of Philosophy* 94: 5-26. (This is an abstract generalization of the core idea of the previous article. It contains references to criticisms and replies to that article.)