THE QUESTION-CENTERED ACCOUNT OF HARM AND BENEFIT

Forthcoming in *Noûs*

Aaron Thieme

University of North Carolina at Chapel Hill

Abstract: The counterfactual comparative account of harm and benefit (CCA) has faced a barrage of objections from cases involving preemption, overdetermination, and choice. In this paper I provide a unified diagnosis of CCA's vulnerability to these objections: CCA is susceptible to them because it evaluates each act by the same criterion. This is a mistake because, in a sense I make precise, situations raise prudential questions, and only some acts—the *relevant alternatives*—are directly relevant to these questions. To answer the objections, we must revise CCA so that its evaluations foreground the relevant alternatives. The result is a question-centered account of harm and benefit.

0. Introduction

Suppose that a young student gets a vaccine. Had the student not been vaccinated, she would have gotten sick, and she would have been worse off. The vaccine strikes us as benefitting the student. It certainly benefits her in a respect since it prevents her from getting sick. But benefits and harms in a respect come cheap: getting the vaccine also harms the student in a respect since (suppose) it involved a painful shot in the arm. When it strikes us that getting the vaccine benefits the student, what we likely have in mind is that it benefits her *overall*. This overall evaluation seems so plausible because if the student had not got the vaccine, she would have been on balance worse off.

This case motivates the counterfactual comparative account (CCA) of overall harm and benefit. According to this account, an act overall benefits a subject when it leaves them better off than they would have been otherwise.

> *CCA*: Act α overall benefits (harms) subject S if and only if S's lifetime welfare would have been lower (higher) had α not been performed.[1]

CCA implies that getting the vaccine overall benefits the student since the student would have been on balance worse off had they not gotten it.

CCA is theoretically elegant, and it surely gives the right verdict in many cases. It is not an accident that CCA has amassed the broadest coalition of supporters in the literature. However, at the same time, problems for CCA have multiplied. CCA arguably gives the wrong verdicts in cases involving one harm preempting another as well as cases involving the overdetermination of harm, and it also leads to disharmony between which act overall benefits an agent and which act there is sufficient prudential reason for that agent to perform.

In this paper, I argue that to resolve these problems we should revise CCA so that its evaluations foreground the *relevant alternative set*. After presenting the problems of prudential disharmony, preemption, and overdetermination (Section I), I argue that they arise for CCA because it does not, for the purposes of evaluation, distinguish any set of alternatives as the relevant alternative set (Section II). I then argue that alternative sets have special relevance when their members are directly relevant to the prudential questions raised by a choice situation (Section III). This motivates a revision to CCA that resolves the problems of prudential disharmony, preemption, and overdetermination (Section IV).

[1] For defenses of CCA, see Feldman (1991), Parfit (1988: 18), Bradley (2009), Klocksiem (2012; 2019; 2022), Boonin (2014), Feit (2015; 2019; 2023), and Timmerman (2016). This principle is often formulated more generally in terms of *events*, but since my interest here lies in the practical significance of overall harm and benefit, I will be working with a narrower formulation in terms of acts. That being said, as emphasized by Bradley (2012: 394–395), it is important to keep in mind the ultimate goal of more broadly accounting for acts (individual and collective) and events in a unified framework. See note 22 below for more discussion.

I. Problems for CCA

Traditionally, it has been thought that CCA harmonizes nicely with the normative significance of overall harm and benefit.[2] But Carlson, Johansson, and Risberg (2024) argue that CCA does not capture the normative significance of overall harm and benefit.[3] To illustrate, consider the following case.

> *Drugs*: Arya has a painful disease that severely compromises her quality of life. She can take exactly one of three drugs. She can take the red drug A. This would fully cure her disease and lead to a life in which she flourishes. She could alternatively take a blue drug B. This would partially ameliorate the disease's symptoms, but it would result in a life that is not worth living. Finally, she can take a blue drug C that would be totally ineffective and leave her even worse off than if she had taken B. Since Arya likes blue, and she doesn't know the results of taking the drugs, she takes drug B. Had she not taken drug B, she would have taken drug C.[4]

In cases like this, the verdicts of CCA fail to harmonize with objective prudential reason. In particular, CCA conflicts with two attractive principles about what there is objective prudential reason to do and how this relates to what is overall beneficial.[5]

> *Avoid Evil*: In a choice situation in which act α is available to a subject S, if α would leave S at a negative well-being level while an alternative available act would leave S at a positive welfare level, then there is not sufficient prudential reason for S to do α.[6]

---

[2] For instance, see Parfit (1988: 18) who holds that CCA captures the morally relevant sense of 'harm,' and see also Bradley (2009: 69–72) who defends CCA on the grounds that what it considers overall beneficial (harmful) aligns with what is prudentially reasonable to (not) choose.
[3] See also Carlson (2019; 2020) and Carlson, Johansson, and Risberg (2021).
[4] This case is analogous to one discussed by Carlson, Johannson, and Risberg (2024).
[5] Throughout, I use 'reason' in the objective or fact-relative sense, and I will leave this qualifier implicit.
[6] This roughly corresponds to "Claim 3" in Carlson, Johannson, and Risberg (2024: 477), and Harmony is loosely inspired by their "Claim 4."

Harmony: In a choice situation in which act α is available to subject S, if α overall benefits (harms) S, then there is (not) sufficient prudential reason for S to perform α.

Avoid Evil will only apply to a limited number of choice situations, but for those choice situations, it clearly gives the right verdicts about what there is not sufficient reason to choose. Harmony expresses an attractive harmony between what overall benefits (harms) an agent and what there is (not) sufficient prudential reason for them to do; it captures one important way in which overall harm and benefit are normatively significant.

Following Carlson, Johansson, and Risberg (2024), we may show that CCA is inconsistent with jointly accepting Avoid Evil and Harmony. Since taking B leaves Arya at a negative lifetime welfare level whereas taking A would have left her at a positive lifetime welfare level, by Avoid Evil there is not sufficient prudential reason for Arya to take B. However, according to CCA, taking B overall benefits Arya since she would otherwise have taken C, and she would then have been even worse off. By Harmony, it follows that there is sufficient prudential reason for Arya to take B. This is inconsistent.

In addition to this recent problem of disharmony between CCA and prudence, it has been known for some time that CCA gives implausible verdicts about what is overall harmful and beneficial in cases of preemption and overdetermination. First, consider a case where one harm preempts another:

Preemptive Punch: Jack is a rising movie star. Jealous of his success, Jill viciously punches Jack, disfiguring him and derailing his career. After the attack, Jack only gets

roles as villains in B-tier movies. Jill could easily have left Jack alone, but if Jill had not punched Jack, she would in fact have shot Jack to death.[7]

Since Jack would have been even worse off had Jill not punched him, on CCA the punch overall benefits Jack. This is very counterintuitive.

Next, consider a case of an overdetermined harm:

Overdetermined Hit: Al Capone wants George "Bugs" Moran dead, but he doesn't trust his goons to do the job right. Capone shoots Moran with a revolver while simultaneously stepping on a TNT detonator. Moran dies from the resulting shooting and explosion. Had Capone not shot Moran, Capone would have still blown up Moran with the TNT. But had Capone neither shot nor blown up Moran, Moran would have gone on to enjoy a much longer and happier life.

Intuitively, Capone's shooting Moran overall harms him. But had Capone not shot Moran, Moran would have been no better or worse off. As a result, CCA implausibly says that Capone's shooting Moran is overall neutral for Moran.

II. The Source of the Problems

In considering a case, we have various options about what to evaluate. In Drugs, we could evaluate Arya's *taking B* or we could evaluate her *taking a blue drug* (i.e., the disjunctive act of taking B or C). In Preemptive Punch, we could evaluate Jill's *punching* Jack, or we could evaluate Jill's *attacking* Jack (by punching or shooting him). In Overdetermined Hit, we could evaluate Capone's *shooting* Moran, or we could evaluate Capone's *killing* Moran (by shooting or blowing up Moran).

[7] This case is due to Norcross (2005; 2020) and especially Carlson, Johansson, and Risberg (2021). Other cases of preemption have been discussed by Feldman (1991), McMahan (2002: 117), and Bradley (2009: 53; 2012) among others.

It will be useful to have on hand some terminology for describing the different things we may evaluate. Two acts are *alternatives* just in case they practically cannot be performed together. An agent's available act is *an alternative act* just in case there is an act available to the agent that is an alternative to it. A set of available acts is an *alternative set* just in case its members are jointly exhaustive and each pair of its members are alternatives. We assume that each alternative set contains at least two available acts. Alternative sets have special practical significance: for each alternative set for an agent, the agent will have to perform exactly one alternative in the set.

When evaluating an agent's acts, it matters quite a bit which alternatives in particular we consider. Arya's taking B overall benefits her on CCA, whereas Arya's taking a blue drug overall harms her on CCA. Similarly, on CCA Jill's punching Jack overall benefits him, but her attacking him overall harms him. Finally, on CCA Capone's shooting Moran is overall neutral for him, but Capone's killing Moran overall harms him. Relative to some alternative sets, these agents do things that are beneficial or neutral while relative to others they do things that are harmful. This raises a deep question: when engaged in evaluative inquiry, should we treat every alternative set the same, or should we somehow foreground some while backgrounding others?

II. A. The Relevant Alternatives Diagnosis

Loosely speaking, an alternative set is relevant just in case each of its members is apt for being evaluated in comparison to its other members. More exactly, let us say that an alternative set is relevant just in case its members have their harmfulness or beneficialness directly determined by how they compare to the alternatives in the set. Given this notion of relevant alternative sets, we may distinguish two parts of CCA. First is a part of CCA that is often left implicit—it treats each performed alternative as belonging to a corresponding relevant alternative set.

CCA's omnism: For any performed alternative α, α belongs to the relevant alternative set {α, ~α}.[8]

In other words, CCA's omnism implies that each performed act's harmfulness or beneficialness is directly determined by how it compares to what would have been done instead of it. The second part of CCA is its evaluative criterion, its evaluative core. This criterion evaluates relevant alternatives by comparing the lifetime welfare that would result from them.

CCA's criterion: For any performed act α in a relevant alternative set {α, ~α}: α overall benefits (harms) S iff α leaves them on balance better (worse) off than ~α would have.

Now that we have pulled apart CCA's omnism and its criterion, it is natural to wonder how they share the blame for the problems CCA faces. It may be tempting to automatically place all the blame on the criterion, but I will argue that omnism is what is to blame. In my view, to employ CCA's criterion we must identify some alternative sets as uniquely relevant: we must deny CCA's omnism.

In fact, it is natural to think that some alternative sets but not others are relevant. For instance, Arya's taking B seems to harm her overall even though it is better for her than what she would have done otherwise (namely, taking C). This suggests that {Arya's taking B, Arya's not taking B} is not a relevant alternative set for Arya in Drugs. In contrast, Arya's taking a blue drug seems to harm her overall because it is worse for her than what she would have done otherwise (namely, taking the red drug A). This suggests that {Arya's taking a blue drug, Arya's not taking a blue drug} is a relevant alternative set for Arya in Drugs. In this case, some sets of alternatives but not others are relevant.

[8] The "omnism" label is, as far as I know, due to Douglas Portmore who has employed it to describe normative and deontic theories. See, for instance, Portmore (2017). The classic study of relevant alternatives, in the context of deontic theory, is Bergström (1966).

The same story goes for Preemptive Punch.[9] Jill's punching Jack leaves him better off than what she would have done otherwise (namely, shoot him). But this does not suffice to make Jill's punching Jack overall beneficial for him. This suggests that {Jill's punching Jack, Jill's not punching Jack} is not a relevant alternative set. In contrast, Jill's attacking Jack seems to harm him because it is worse for him than what she would have done otherwise (namely, leave him alone). This suggests that {Jill's attacking Jack, Jill's not attacking Jack} is a relevant alternative set.

Finally, perhaps most clearly of all, Capone's shooting Moran does not seem neutral for Moran even though it is no better or worse for Moran than what Capone would have done otherwise (namely, blow up Moran without shooting him). So the alternative set {Capone's shooting Moran, Capone's not shooting Moran} is not relevant. But clearly Capone's killing Moran harms him since it leaves him worse off than what Moran would have done otherwise. So the alternative set {Capone's killing Moran, Capone's not killing Moran} is relevant. Altogether, these observations cast serious doubt on CCA's omnism.

II. B. Broadening the Diagnosis

So far I have suggested that some sets of alternatives have special relevance for evaluating what is overall harmful and beneficial. The evaluations of overall harm and benefit that we have been interested in so far use "harmful" and "beneficial" in their so-called *positive* form. But notice that evaluations which use "harmful" and "beneficial" in their *superlative* forms—evaluations of whether an act is overall *most* beneficial or overall *most* harmful—are also normatively

[9] Feldman (1991: 225) foreshadows my diagnosis of the preemption problem for CCA. He suggests that some alternatives are more appropriate to focus on: in discussing McMahan's (1988: 45) case where a young soldier's death at time t preempts a slightly later death, Feldman (1991: 225–226) writes "The real tragedy here is not that he died exactly at t [...] the real tragedy is that he died so young. Thus, [his dying so young] should be the focus of our attention." I agree with this assessment. But Feldman did not go on to blame CCA's omnism; he thought instead that if we carefully distinguished different alternatives, then we could uniformly use CCA to evaluate them.

significant. If an act is overall most beneficial (harmful) for an agent, then there is most prudential reason for the agent to (not) perform that act.[10]

To broaden our diagnosis of the problems facing CCA's analysis of what is overall harmful or beneficial, it will be instructive to consider whether it extends to analogous problems concerning what is most overall harmful or beneficial. As written CCA does not tell us whether a given act is overall most beneficial or harmful. But we can formulate a cousin to CCA that analyzes the superlative.

> *CCA_{Most}*: α is overall most beneficial (harmful) for S just in case α leaves S at a higher (lower) level of lifetime welfare than any alternative would have.

CCA_{Most} is close to being adequate. But we again run into problems when we fail to identify the relevant alternative set. Consider what CCA_{Most} says about the disjunctive act of Arya's taking A or B. The only alternatives to taking A or B involve taking C, and each of these alternatives would leave Arya worse off than taking A or B. It follows by CCA_{Most} that taking A or B is overall most beneficial for Arya. But this is wrong. Taking A or B is clearly not overall most beneficial for Arya. As a matter of fact, Arya takes A or B, and the outcome that results (namely, Arya's life being not worth living) from this is far from the best available outcome.

The fix is again to privilege a set of relevant alternatives. The problem arose because we were working with the alternative set {taking A or B, taking C}. But when it comes to evaluating what is overall most beneficial for Arya, we should consider Arya's taking B rather than taking A or taking C. That is, we should work with the alternative set {taking A, taking B, taking C}. Since taking A is an alternative to taking B, and since taking A would have left Arya better off than taking B, CCA_{Most} entails that taking B is not overall most beneficial for Arya.

[10] Thanks to an anonymous reviewer for helpful discussion on these points.

This again suggests that to properly apply the counterfactual comparative evaluative machinery, we need to identify the set of relevant alternatives. As with CCA, we may distinguish between two commitments in CCA_{Most}.

> *CCA_{Most}'s omnism*: Each performed alternative α belongs to a relevant alternative set.
>
> *CCA_{Most}'s criterion*: For any performed act α in a relevant alternative set: α is overall most beneficial (harmful) for S just in case α leaves S at a higher (lower) level of lifetime welfare than any alternative would have.

We should again doubt CCA's omnism rather than CCA_{Most}'s criterion.

I have argued that we should revisit the assumption of omnism in theorizing about what is overall (most) beneficial or harmful. But if we are to make progress in revising omnism, we must develop an account of the relevant alternative set. I turn to this task in the next section.

III. Choice Situations and Questions

A choice situation is determined by an agent, the acts available to them, and the outcomes that their acts decide between. I assume that the acts available to the agent are legion and lie at many different levels of specificity. For instance, in Drugs, Arya could take a drug, she could take a blue drug, she could take B, she could take B apprehensively, and so on. To understand a choice situation, we would like to *resolve* the acts available to the agent into an *alternative set* **A**.

It will be important to keep in mind the distinction between the acts *in* an alternative set **A** and the acts *available* to the agent in the choice situation. In each of our examples, we only consider relatively small alternative sets that highlight a few of the available acts. But there are many more available acts than these. Many, but not all, of the available acts will be ways, or instances, of performing the acts in the alternative set. When an act α is a way of performing

another act β (i.e., when it is practically impossible to perform α without performing β), we say that α entails β.

An act is forced in a choice situation just in case there is no available alternative act to it. For instance, in Drugs we imagine that for whatever reason Arya has to take a drug. Because they are not alternatives to any available act, forced acts will neither be in the alternative set, nor will they be ways of performing the acts in the alternative set.

As Brown (2018) observes, the available acts in a choice situation form a partially ordered set. At the top of the order, we have the act forced in the situation. This act is entailed by alternative acts, so as we move down the order, we split the forced act into its alternative ways.

Here is one way, among others, of graphing some of the available acts in Drugs. The graph may be useful to refer to since it provides a concrete interpretation of the issue of specifying the alternative set. How to specify the alternative set can be approached in terms of how to split an act into alternatives and when to stop splitting.

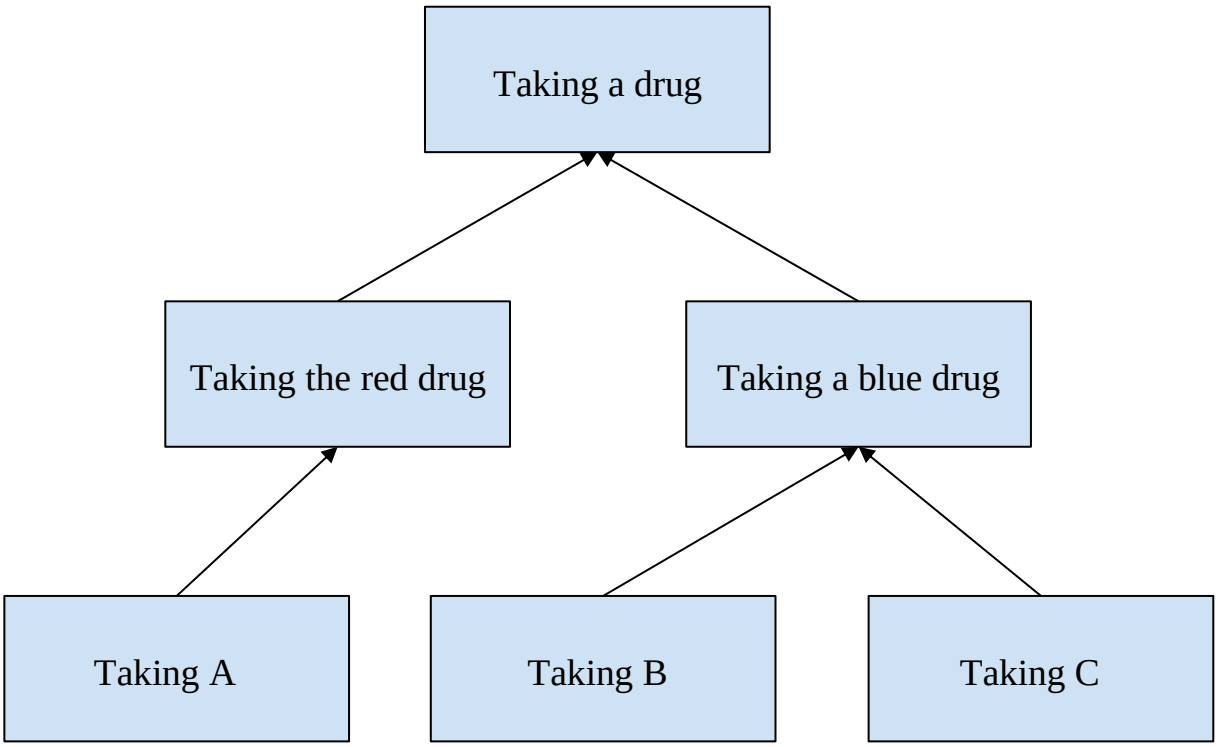

Moving on to a choice situation's outcomes, we will assume that the outcomes an agent's available acts decide between are possible worlds. We call the set of possible worlds that an agent decides between in a choice situation *the outcome set* **W**.[11] We assume that the possible worlds in **W** do not go beyond what would result from the agent performing an available act in the situation: for each world w in **W**, we assume that there is an act available to the agent such that if the agent performed that act, w would result. Although this act must be available to the agent, we do not require that this act is included in **A**. (When it is more felicitous, I will often talk of the available *lives* for a given subject in a given choice situation. This is just a shorthand way of referring to the choice situation's available *possible worlds* in which the subject exists.)

For example, consider Arya's choice situation in Drugs. We can model the outcome set $\mathbf{W_{Drugs}}$ as simply $\{w_A, w_B, w_C\}$, the set consisting of w_A (the world where Arya takes A and has a life that is worth living), w_B (the world where Arya takes B and has a life that is not worth living), and w_C (the world where Arya takes C and has a life that is even worse).

As in any choice situation, in Drugs there are distinct alternative sets. One is $\mathbf{A_1}$ = {taking A, taking B, taking C}. A second is $\mathbf{A_2}$ = {taking the red drug, taking a blue drug}. A third is $\mathbf{A_3}$ = {taking A or B, taking C}. It will be important to keep in mind that what defines a set of alternative acts is its members, not how they are described. For example, $\mathbf{A_2}$ = {taking the red drug, taking a blue drug} = {taking A, taking B or C} = {taking the red drug, not taking the red drug}. As this example suggests, any two member alternative set can always be written in the form $\{\alpha, \sim\alpha\}$, for some available act α.[12]

[11] This is intended to evoke Stalnaker's (1975: 274) conception of a set of worlds that is live for the purposes of conversation. For us, the outcome set specifies what is *practically* open in a choice situation.

[12] I am assuming here that if α is an available act with at least one available alternative, then $\sim\alpha$ is also an available act.

Only some alternative sets seem to be relevant. My proposal is that these are exactly the alternative sets that are relevant *to the prudential questions raised by a choice situation*. To work this out, we need to look more closely at the nature of a question.

On the standard view in linguistics and philosophy of language, a question partitions logical space.[13] For instance, the question of *Whether I am at least six feet tall* partitions logical space into the worlds in which I am at least six feet tall and the worlds in which I am not at least six feet tall. In general, Yes/No questions always partition logical space into two cells—the cell in which they are answered Yes and the cell of worlds in which they are answered No. Similarly, a multiple-choice question like *Which of A, B, or C* partitions logical space into the A-worlds, the B-worlds, and the C-worlds.[14]

III. A. Two Prudential Questions

One prudential question is the fine-grained question *How well does the subject fare on balance?* This is a prudentially important question since prudentially it matters how well the subject fares on balance. This question partitions logical space into cells that correspond to levels of lifetime welfare. Since there are many such levels, this is a potentially very fine-grained partition. I will call this question 'the fine question' and its partition 'the fine partition.'[15]

There is a second prudentially important question, one that is coarser-grained: *Is the subject's life close enough to the best available to be worth wanting for them*?[16] This is a coarse-grained question since it distinguishes at most two cells of worlds—the cell of worlds in which

[13] This goes back to Groenendijk and Stokhof (1982). For a review of the literature, see Groenendijk and Stokhof (1997). See also Lewis (1988a; 1988b) and Yablo (2014) who use partitions of logical space to give a theory of subject matter with broad philosophical implications. I am indebted to Hoek (2022) and especially McNamara (2024) for their discussions of questions in practical contexts.

[14] When I talk of partitions throughout, I always mean a *non-trivial* partition (i.e., not a partition of set X that is simply {X}).

[15] I am very grateful to an anonymous reviewer for encouraging me to discuss the fine partition.

[16] This uses 'worthy' in a relative, satisficing sense. On satisficing, see Slote (1984) and on relative versus absolute satisficing, see Hurka (1990).

the subject's life is close enough to the best available to be worth wanting and the cell of worlds in which their life is not close enough to the best available to be worth wanting. I will accordingly call it 'the coarse question' and its partition 'the coarse partition.'[17]

I will argue that the fine question and its fine partition as well as the coarse question and its coarse partition have important roles to play in our account of overall harm and benefit. The former determines the relevant alternatives for evaluations of what is overall *most* harmful or beneficial. The latter determines the relevant alternatives for evaluations of what is overall harmful or beneficial.

The two figures below illustrate the fine partition and the coarse partition of $\mathbf{W_{Drugs}}$. The fine partition distinguishes each world in $\mathbf{W_{Drugs}}$ since Arya's level of lifetime welfare is different in each. But since w_B is not close enough to w_A to be prudentially worth wanting, the coarse partition lumps together w_B and w_C as both too far away from the best to be prudentially worth wanting.

[17] Moreover, when it is convenient, I will say that a world w is worth wanting *out of* **W** when w is close enough to the best available in **W** to be worth wanting.

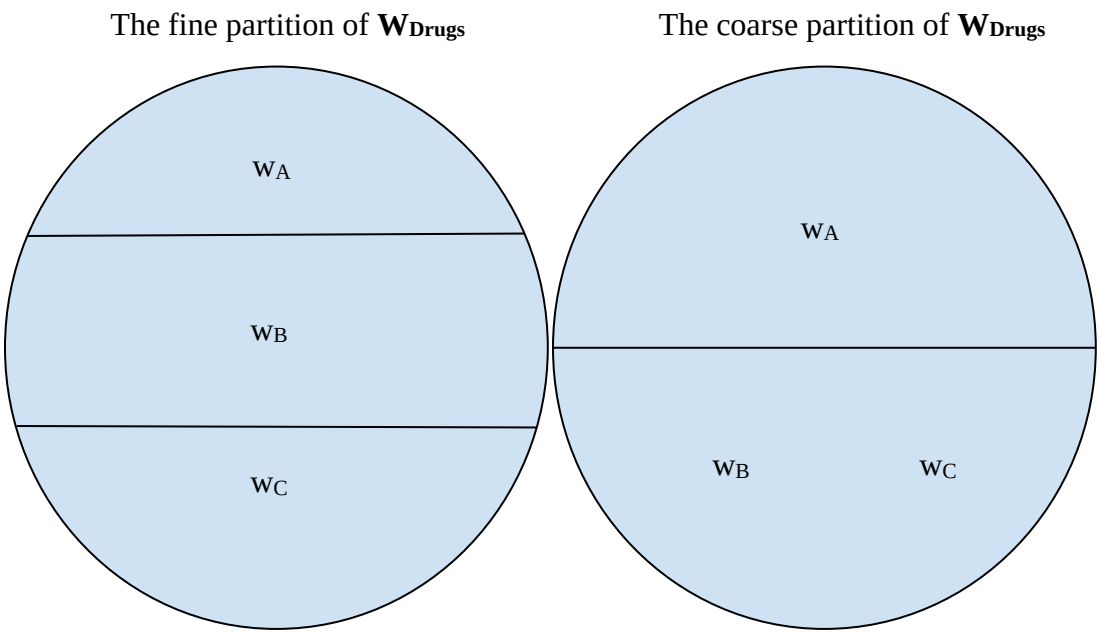

Being close enough to the best to be worth wanting is not the sort of concept that lends itself to a fully precise analysis. It comes with an intrinsic degree of vagueness, and I will not pretend to exorcize this vagueness. But that does not mean that we cannot illuminate this concept. To do so, we lay a foundation for the concept that is not perfectly solid but rather sways with the concept.[18]

The foundation I propose is *prudential similarity*. Two worlds are prudentially similar to the extent that a subject fares similarly well in them. There are many respects in which a subject might fare similarly well in two worlds. One respect of prudential similarity between worlds is lifetime welfare. This respect is surely the most important because it captures how the subject fares *on balance over the course of their entire life*. But more local respects are also important: two worlds are prudentially similar in a respect when the subject receives a similar bit of well-being (the intrinsic prudential good) or a similar bit of ill-being (the intrinsic prudential evil) in them. For instance, supposing that pleasure is intrinsically prudentially good, when the subject

[18] I am echoing Lewis (1973: 92) here.

receives the same pleasure in two worlds, this is a respect in which the two worlds are prudentially similar. This is a specifically synchronic respect of similarity. There are corresponding diachronic respects of similarity: when the subject has similar trends in welfare in two worlds, these worlds are prudentially similar in this respect. For instance, when a subject receives a similar increase in pleasure over time in two worlds, this is a diachronic respect in which these worlds are prudentially similar.

The coarse partition lumps together some worlds with the best available. By definition, these are the worlds that are close enough to the best to be prudentially worth wanting. But when is a world close enough to the best to be prudentially worth wanting? Roughly speaking, a world is close enough to the best to be prudentially worth wanting when it is not too prudentially dissimilar from the best.

More exactly, a world w in **W** is close enough to the best in **W** to be prudentially worth wanting just in case w is more prudentially similar to the best world(s) in **W** than the worst world(s) in **W**.

Coarse Partition (definition): For any outcome set **W** with worlds that have different levels of lifetime welfare, the coarse partition of **W** is {{w in **W**: w has more lifetime welfare for S than L}{w in **W**: w has less lifetime welfare for S than L}}, for the level of lifetime welfare L such that:

- every world in **W** with at least as much lifetime welfare as L is more prudentially similar to the best world(s) in **W** than the worst world(s) in **W**;
- every world in **W** with less lifetime welfare than L is more prudentially similar to the worst world(s) in **W** than the best world(s) in **W**.

The best worlds are necessarily more prudentially similar to themselves than the worst worlds. This guarantees that L falls between the best and worst worlds for any outcome set **W** that has worlds with different levels of lifetime welfare. One upshot of this definition is therefore that the coarse partition always distinguishes the best worlds from the worst worlds, placing the best in one cell and the worst in another cell. Another upshot of this definition is that the fine partition *refines* the coarse partition, in the sense that two worlds belong to distinct cells of the coarse partition only if they also belong to distinct cells of the fine partition.[19] In other words, the fine partition remembers every distinction drawn by the coarse partition.

The coarse question is prudentially important since it matters prudentially whether a life is close enough to the best to be worth wanting. If a life is close enough to the best to be worth wanting out of those available in **W**, then there is *sufficient prudential reason* (though not necessarily *most* prudential reason) to want this life and to choose it out of **W**. So the coarse partition tracks an important prudential difference between worlds.

This being said, the coarse partition is arguably less prudentially natural than the fine partition.[20] The coarse partition forgets some prudential distinctions that the fine partition draws. Even when the subject has different levels of lifetime welfare in two worlds, the coarse partition may lump these worlds together as similarly worthy or similarly unworthy. Moreover, the coarse partition is also potentially vague since it depends on the balance of different respects of prudential similarity—a matter that cannot be expected to be fully precise. So while the coarse

[19] To see this, suppose that w_1 and w_2 belong to different cells of the coarse partition for any given outcome set. It follows by definition that there is a level of lifetime welfare L such that exactly one of w1 and w2 has more lifetime welfare than L. So w_1 and w_2 must have different levels of lifetime welfare, and it follows immediately by the definition of the fine partition that they must belong to different cells in the fine partition.

[20] On naturalness, see Lewis (1983).

partition is prudentially important, and even prudentially natural to a degree, it is not as prudentially natural or sharp as the fine partition.

III. B. Relevant Alternative Sets

Like a question, an alternative set induces a partition of logical space. To see how, associate with each alternative set a distinctive question—the question of which act in it is performed. Given an alternative set **A** = {α_1, α_2, …, α_n}, we then have the associated question of *Which of α_1, α_2, …, α_n is performed*. This question partitions a set of worlds **W** into {{w in **W**: w is an α_1-world}, {w in **W**: w is an α_2-world}, ..., {w in **W**: w is an α_n-world}. We will therefore call this partition ***A**'s partition of **W*** or *the partition of **W** induced by **A***.

To illustrate, recall the three alternative sets I mentioned earlier for Drugs: $\mathbf{A_1}$ = {taking A, taking B, taking C}, $\mathbf{A_2}$ = {taking the red drug, taking a blue drug}, and $\mathbf{A_3}$ = {taking A or B, taking C}. These alternative sets partition **W** differently. $\mathbf{A_1}$ induces the fine partition of $\mathbf{W_{Drugs}}$; $\mathbf{A_2}$ induces the coarse partition of $\mathbf{W_{Drugs}}$; and $\mathbf{A_3}$ induces a partition that is different from both the fine partition and the coarse partition.

We are now in a position to identify the relevant alternative sets. *How well off the subject is* has prudential importance. Corresponding to its fine partition of outcomes is one kind of prudentially relevant set of alternative acts: the one that induces the same fine partition of outcomes. For any outcome set **W** with worlds that leave the subject with different levels of lifetime welfare, there is an alternative act **A** such that **A** induces the fine partition of **W**. I will call this the *fine alternative set*.

The members of the fine alternative set are distinguished exactly according to how well off they could leave the subject. Which act in this set is performed thus exactly settles how well

off the subject is. An act from this set is overall most beneficial (harmful) just in case it leaves the subject better (worse) off than each alternative in this set.

The fine alternative set is the relevant alternative set for evaluating whether an act is overall *most* harmful or beneficial. But it is not the relevant one for evaluating what is overall harmful or beneficial. It distinguishes acts at a potentially very fine-grained level, and this is potentially too much fineness of grain for the purposes of evaluating what is overall harmful or beneficial. To evaluate overall harm and benefit, we care about the most important comparative differences between alternative acts, not every such difference. For instance, we care about the difference between Arya's taking A and her taking B or C—this difference is what makes Arya's taking A overall beneficial. But we do not similarly care about, say, the difference between Arya's taking B and her taking C—though B leaves Arya better off than C would have, this does not suffice to make taking B overall beneficial. In order to foreground the relevant alternatives for evaluating overall harm and benefit, then, we sometimes want a coarser alternative set than the fine alternative set.

The coarse question—whether the subject's life is close enough to the best available to be prudentially worth wanting—partitions the available worlds coarsely, distinguishing the worthy worlds from the unworthy worlds. There is a prudentially relevant set of alternative acts corresponding to the coarse partition: the set of alternative acts that induces the same partition. When **A** induces the coarse partition of a given outcome set, I will say that **A** is the *coarse alternative set*.

The coarse alternative set can always be written in the form $\{\alpha, \beta\}$ such that every worthy outcome is an α-world and every unworthy outcome is a β-world. On the one hand, the performance of α settles that the subject is in a worthy world rather than an unworthy world. On

the other hand, the performance of β settles that the subject is in an unworthy world rather than a worthy world. So the coarse alternative set distinguishes alternative acts exactly according to whether they leave the subject in a worthy world rather than an unworthy world or an unworthy world rather than a worthy world. These are the important differences made by overall harms and benefits. When an act from the coarse alternative set is better (worse) for the subject than the alternative from that set, it overall benefits (harms) the subject. The coarse alternative set is the relevant one for evaluating what is overall harmful and beneficial.

IV. The Erotetic Account

I have argued that the structure of our account of harm and benefit should foreground the alternatives that are relevant to the prudential questions raised by choice situations. If that is right, then we should reject the omnism of CCA and CCA_{Most}.

> *CCA's omnism*: Each performed alternative α is in a relevant alternative set (namely, {α, ~α}).
>
> *CCA_{Most}'s omnism*: Each performed act α belongs to a relevant alternative set.

For evaluating what is overall harmful or beneficial, the coarse alternative set is the only relevant one, and for evaluating what is overall most harmful or beneficial, the fine alternative set is the only relevant one.

This leads to the following question-centered, or *erotetic*, revision of CCA_{Most}.

Erotetic Account$_{Most}$:

- For any α in the fine alternative set: α is overall most beneficial (harmful) for S just in case α leaves S at a higher (lower) level of lifetime welfare than any other member of the fine alternative set would have.
- For any α not in the fine alternative set: α is not overall most harmful or beneficial for S.

For instance, $\mathbf{A_1}$ = {taking A, taking B, taking C} induces the fine partition of $\mathbf{W_{Drugs}}$, and so $\mathbf{A_1}$ is the relevant alternative set for determining which act is overall most harmful or beneficial for Arya. Since taking B does not leave Arya at a higher level of lifetime welfare than any other member of $\mathbf{A_1}$ would have, the erotetic account$_{Most}$ (like CCA$_{Most}$) implies that taking B is not overall most beneficial for Arya. But since taking A or B does not belong to $\mathbf{A_1}$, (unlike CCA$_{Most}$) the erotetic account$_{Most}$ implies that taking A or B is not overall most beneficial for Arya.

The erotetic account$_{Most}$ only gives the members of the fine alternative set a chance at being considered overall most harmful or beneficial. Given its second clause, all other acts are automatically not overall most harmful or beneficial. This includes acts that lump together acts that belong to the fine alternative set (e.g., taking A or B) as well as acts that are *non-trivial versions* of acts in the fine alternative set (e.g., taking A apprehensively).[21] Non-trivial versions of acts in the fine alternative set are too specific to be most overall harmful or beneficial since they have alternatives (e.g., taking A without apprehension) that are just as good.

When it comes to evaluating what is overall harmful and beneficial, the structure of our account should differ in this last respect. It is clearly possible for *versions* of acts in the coarse alternative set to be overall harmful or beneficial. For instance, $\mathbf{A_2}$ = {taking the red drug, taking a blue drug} induces the coarse partition of $\mathbf{W_{Drugs}}$, so it is the relevant alternative set for evaluating what overall harms or benefits Arya in Drugs. Taking B is a version of taking a blue drug, and taking B overall harms Arya. So we should not deny that all acts outside of the coarse alternative set are overall harmful or beneficial. But if non-trivial versions of relevant alternatives are overall harmful or beneficial, how is their status determined?

[21] An act α is a *non-trivial* version of an act β just in case α is a version of β but α is not identical to β.

The coarse alternative set contains the two alternatives that directly and exactly settle the question of whether the subject has a life that is prudentially worth wanting out of those available. A version of one of these two relevant alternatives *indirectly* settles this question. It indirectly settles this question by entailing the relevant alternative of which it is a version. For instance, Arya's taking B entails taking a blue drug, which belongs to the coarse alternative set $\mathbf{A}_2$. Since taking a blue drug leaves Arya worse off than she would have been otherwise, taking a blue drug overall harms Arya, and by extension taking B overall harms Arya.

We have covered how to evaluate acts that entail a relevant act in the coarse alternative set. What about acts (e.g., Arya's taking A or B) that do not entail any act in the coarse alternative set? These acts do not settle, even indirectly, whether the subject has a life that is prudentially worth wanting out of those available. They leave this question open. So even though they may be better or worse than an alternative, they are not overall beneficial or harmful.

So far, I have only discussed choice situations in which something is at stake for the subject—choice situations in which the subject might fare on balance better or worse. In every such choice situation, the best available outcomes leave the subject with more lifetime welfare than the worst available outcomes. It follows from the definition of the coarse partition that the best available worlds will be placed in one cell and the worst available worlds in another. As a result, there will be coarse alternative sets for these choice situations.

But what about choice situations in which the subject has the same level of lifetime welfare in every available possibility? Suppose I am choosing which color socks to wear—white or black? In this situation, suppose that nothing I can do makes a difference to how well I fare on balance. In this kind of case, there is no practically live prudential question raised by my choice situation. The fine and coarse questions are both practically closed since they are answered the

same in every available outcome. They both induce a trivial, one cell partition of my choice situation's outcome set. There is no *alternative* set that induces this trivial, one cell partition. (Since any alternative set must have at least two members that are available alternatives, it must induce a partition with at least two cells.) In choice situations in which every outcome leaves the subject with the same level of welfare, it follows that there is no relevant alternative set. In such situations, every act is overall neutral for the subject.

Altogether, I have argued that we should use CCA's criterion to directly evaluate the members of the coarse alternative set, and we should evaluate other acts based on whether they entail a member of the coarse alternative set. This leads to the following account of overall harm and benefit.

The Erotetic Account: For any choice situation with available act α and subject S:

- α would overall benefit (harm) S iff there is a coarse alternative set **A** and α entails the member of **A** that would leave S better (worse) off than its alternative in **A**.
- If there is no coarse alternative set, α would be overall neutral for S.[22]

The erotetic account resolves the problem of disharmony. Recall that the problem of disharmony arose because CCA is inconsistent with jointly accepting two principles.

Avoid Evil: In a choice situation in which act α is available to a subject S, if α would leave S at a negative well-being level while an alternative available act would leave S at a positive welfare level, then there is not sufficient prudential reason for S to do α.

[22] The erotetic account only speaks to what *acts* (performed by a single agent) would be overall harmful or beneficial in an agent's choice situation. Ultimately, however, we want to generalize its machinery to situations in which *events* occur, regardless of whether those events are the acts of single agents, collective acts of multiple agents, or impersonal events. I believe that the erotetic account can indeed be generalized in this way. But I will leave filling in the details of this generalization to future work.

Harmony: In a choice situation in which act α is available to subject S, if α overall benefits (harms) S, then there is (not) sufficient prudential reason for S to perform α.

Happily, given eminently plausible background assumptions, the erotetic account entails both of these principles. To derive Avoid Evil and Harmony, we need two background assumptions about sufficient prudential reason.

(*) For any choice situation with available act α and outcome set **W**: If the nearest α-world in **W** is (not) worth wanting for S out of **W**, then there is (not) sufficient reason for S to perform α.[23]

(**) For any outcome set **W** with worlds w and w*: if S's life is worth living in w and S's life is not worth living in w*, then w* is not prudentially worth wanting for S out of **W**.[24]

Both of these principles seem very plausible. As long as the erotetic account is consistent with (*) and (**), the erotetic account must also be consistent with Avoid Evil and Harmony since Avoid Evil and Harmony are entailed by the erotetic account, (*), and (**).

Consider more concretely how the erotetic account and prudential reason harmonize in Drugs. On the one hand, taking B implies taking a blue drug, a member of the coarse alternative set in Drugs. Since Arya would have been better off had she not taken a blue drug, the erotetic account implies that taking a blue drug harms her overall, and by extension it implies that taking B harms her overall. On the other hand, since taking B would result in a prudentially unworthy

[23] Assuming (*) and the erotetic account, we may derive Harmony as follows. Consider any choice situation with outcome set **W** in which S's performing α overall benefits S. Assuming the erotetic account, it follows that there is a coarse alternative set **A** = {α*, ~α*} such that α entails α* where α* leaves S better off than ~α*. Because {α*, ~α*} is the coarse alternative set, and α* leaves S better off than ~α*, every α*-world in **W** must be worth wanting for S out of **W**, and every ~α*-world in **W** must be not worth wanting for S out of **W**. Since α entails α*, it follows that every α-world in **W** must be worth wanting for S out of **W**. So the closest α-world in **W** is worth wanting for S out of **W**. By (*), it follows that there is sufficient reason for S to perform α. (The derivation of the respective claim for overall *harm* is analogous.)

[24] Assuming (*) and (**), we may derive Avoid Evil as follows. Consider any choice situation with outcome set **W** in which alternative acts α and β are available to subject S such that α would leave S at a negative well-being level while β would leave S at a positive welfare level. By (**), it follows that the nearest α-world is not worth wanting for S out of **W**. By (*), it follows that there is not sufficient prudential reason for S to perform α.

world, it follows by (**) that there is not sufficient prudential reason for Arya to take B. What overall harms Arya harmonizes with what she lacks sufficient prudential reason to do.

Moreover, taking A would result in a prudentially worthy world, and so by (*) there is sufficient prudential reason for Arya to take A. Happily, the erotetic account implies that taking A would overall benefit Arya. On the erotetic account, then, what would overall benefit Arya harmonizes with what she has sufficient reason to do in Drugs.

In addition to resolving the problem of disharmony, the erotetic account gives extensionally adequate verdicts in cases of preemption and overdetermination. Recall our central case of preemptive harm:

> *Preemptive Punch*: Jack is a rising movie star. Jealous of his success, Jill viciously punches Jack, disfiguring him and derailing his career. After the attack, Jack only gets roles as villains in B-tier movies. Jill could easily have left Jack alone, but if Jill had not punched Jack, she would in fact have shot Jack to death.

We may write the outcome set for Preemptive Punch as $\mathbf{W_{Punch}}$ $\{w_{Punch}, w_{Shot}, w_{Left\ Alone}\}$. There is an important similarity between w_{Punch} and $w_{Left\ Alone}$: in both, Jack's life is worth living. But since Jill brutalizes Jack in w_{Punch} and w_{Shot} alike, leading to similar declines in his welfare and relatively low lifetime welfare, w_{Punch} is more prudentially similar to w_{Shot} than $w_{Left\ Alone}$. As a result, only $w_{Left\ Alone}$ is prudentially worth wanting for Jack out of $\mathbf{W_{Punch}}$. The coarse alternative set can thus be written as {Jill leaves Jack alone, Jill punches or shoots Jack}. Since Jill's punching Jack entails her punching or shooting Jack, and Jack would have been better off had Jill not punched or shot him, the erotetic account entails that Jill's punching Jack overall harms him.

In Overdetermined Hit, the available outcomes involve Capone's shooting and bombing Moran or omitting to do these things to him. We can thus write the outcome set as $\mathbf{W_{Hit}} = \{w_{S\&B}, w_{S\&\sim B}, w_{\sim S\&B}, w_{\sim S\&\sim B}\}$. The only outcome that is worth wanting for Moran out of $\mathbf{W_{Hit}}$ is $w_{\sim S\&\sim B}$: this is the outcome where Moran avoids being shot or bombed and goes on to have a much longer and happier life, whereas in all the others his life is similarly cut tragically short. The coarse alternative set for Capone is thus {shooting or bombing Moran, neither shooting nor bombing Moran}. Since Capone's shooting Moran entails shooting or bombing Moran, and Moran would have been better off had Capone neither shot nor bombed him, the erotetic account entails that Capone's shooting Moran overall harms Moran.

Drugs, Preemptive Punch, and Overdetermined Hit are all cases with three acts in the fine alternative set. It is worth considering how the erotetic account applies to cases with four distinct acts in the fine alternative set.[25] Consider a four-act variation on Drugs:

> Drugs*: Arya has a debilitating disease that, so far, makes her life not worth living. Arya can take exactly one of four drugs for the disease. Drug A* would fully cure her disease and lead to a life in which she flourishes. Drug B* would significantly ameliorate the disease's symptoms and result in a life that is barely worth living. Drug C* would slightly ameliorate the disease's symptoms and leave her with a life that is barely not worth living. Drug D* would be totally ineffective and leave her with a life that is not worth living and much worse than C*. Arya takes B*.

We may model the outcome set for this case as $\mathbf{W_{Drugs*}} = \{w_{A*}, w_{B*}, w_{C*}, w_{D*}\}$. Arya has a life worth living in w_{A*} and w_{B*} alike. Moreover, both contain an incline in temporal welfare from a level at which Arya's life is not worth living to a level at which it is worth living. Because of

[25] I am grateful to two anonymous reviewers for encouraging me to address cases of this form.

this, w_{B*} is plausibly more prudentially similar to w_{A*} than w_{D*}. So w_{B*} is prudentially worth wanting out of $\mathbf{W_{Drugs*}}$. Since w_{C*} would leave Arya with a life that is not worth living while other outcomes would leave her with a life that is worth living, by (**) it follows that w_{C*} cannot be prudentially worth wanting out of $\mathbf{W_{Drugs*}}$. The coarse partition of $\mathbf{W_{Drugs*}}$ is therefore {{w_{A*}, w_{B*}},{w_{C*}, w_{D*}}}, and the coarse alternative set is {taking A* or B*, taking C* or D*}. Since taking B* entails the better act in this set, the erotetic account implies that taking B* overall benefits Arya.

By Harmony, it follows that there is sufficient reason for Arya to take B*. One might worry about the plausibility of this verdict.[26] After all, B* results in a life that is barely worth living, and if she had taken A*, her life would have been much better.

It is important to distinguish two issues here. One issue concerns whether w_{B*} is close enough to the best to be prudentially worth wanting. A second issue concerns what follows from an outcome's being close enough to the best to be prudentially worth wanting. The principles I have defended—the erotetic account and (*)—take a stand on this latter issue. But they do not require us to resolve the first issue in any particular way.

I am inclined to judge that w_{B*} is indeed close enough to the best to be prudentially worth wanting. It thus seems to me that there is sufficient reason to take B*, even though it is a distant second best. After all, taking B* improves Arya's life from one that is not worth living to one that is worth living whereas *at least one* of its alternatives would have left her with a life that is not worth living. In light of this, it seems to me that B* overall benefits Arya and that there is sufficient prudential reason for her to take B*.[27]

---

[26] Thanks to an anonymous reviewer for encouraging me to address this worry.

[27] This being said, there is not reason to take B* *rather than* A*. There is clearly more prudential reason to take A* than to take B*, and indeed there is most prudential reason to take A*. Happily, this verdict harmonizes with erotetic account$_{Most}$ which implies that Arya's taking B* is not overall most beneficial for her, and instead her taking A*

These judgments reflect an openness to a significant, bold degree of satisficing. Some may find this degree of satisficing too extreme. They are open only to less extreme instances of satisficing. Call them *reluctant satisficers*. Reluctant satisficers may insist that since w_{B^*} falls so far short of the best, it cannot be close enough to the best to be prudentially worth wanting.

The reluctant satisficer and I disagree over the first issue, the issue of whether w_{B^*} is close enough to the best to be prudentially worth wanting. But we can agree about the second issue concerning what follows from an outcome's being close enough to the best to be prudentially worth wanting. If they accept the erotetic account and (*), the reluctant satisficer will accordingly judge that taking B overall harms Arya and that there is not sufficient prudential reason for Arya to take B*. Though I do not share these particular judgments, they are not unreasonable.

Next, consider a case that has four acts in the fine alternative set with three resulting in declines in welfare.

> *Market*: Claire has built a successful hedge fund, but she now faces a critical choice. With the market looking shaky, she needs to decide between holding her positions, investing in company X, investing in company Y, or investing in company Z. If she holds, things will go very well and her quality of life will remain constantly high. If she invests in X, she will suffer a slight loss, leading to a slight decline in her quality of life. If she invests in Y, she will lose so much that her life will not be worth living on the whole. If she invests in Z, she will suffer an even worse decline, and her life will be even worse than if she had invested in company Y. She invests in X.

---

would have been most overall beneficial for her. Since taking A* would have been most overall beneficial for Arya, there is most prudential reason for her to take A*.

The outcome set for Market can be written $\mathbf{W_{Market}} = \{w_H, w_X, w_Y, w_Z\}$. On the one hand, Claire's lifetime welfare in w_X is very similar to w_H: in both she has positive and high levels of welfare. This makes w_X prudentially similar to w_H. On the other hand, w_X contains a decline in temporal welfare, and this makes w_X prudentially similar to w_Z in a respect. But because there is so much similarity in lifetime welfare between w_X and w_H, and because w_X contains only a slight decline, it seems to me that w_X is more prudentially similar to w_H than w_Z. If that is right, w_X is worth wanting out of $\mathbf{W_{Market}}$. The coarse alternative set is therefore {holding or buying X, buying Y or Z}. Since buying X entails holding or buying X, and this leaves Claire better off than buying Y or Z, the erotetic account says that buying X overall benefits Claire.

> Finally, consider a variation on Market where the second best is much less good.
>
> *Market**: If Claire holds, her quality of life will remain high. If she invests in X*, she will suffer a heavy loss, leading to a life that is barely worth living. If she invests in Y*, she will lose so much that her life on the whole will be not worth living. If she invests in Z*, she will suffer an even worse loss, and her life will be even worse than if she had invested in Y*. She invests in X*.

We may write this case's outcome set as $\mathbf{W_{Market*}} = \{w_{H*}, w_{X*}, w_{Y*}, w_{Z*}\}$. w_{X*} contains a lot less lifetime welfare than w_{H*}, and it contains a significant decline in temporal welfare similarly to w_{Z*}. As a result, w_{X*} seems more prudentially similar to w_{Z*} than w_{H*}. It is accordingly not worth wanting for S out of $\mathbf{W_{Market*}}$. The coarse alternative set must then be {holding, buying X* or Y* or Z*}. Since buying X* entails buying X* or Y* or Z*, and this leaves Claire worse off than holding, the erotetic account implies that buying X* overall harms Claire in Market*.

V. Conclusion

CCA is an attractive account of overall harm and benefit. It says that an act overall benefits (harms) a subject just in case it leaves them better (worse) off than they would have been had it not been performed. In spite of this account's attractiveness, it faces problems of prudential disharmony, preemption, and overdetermination. I have argued that these problems have a common source. All of these problems arise for CCA because it does not identify a privileged set of relevant alternative acts. Once we identify this set, we may apply CCA directly to its members to evaluate them as overall beneficial or harmful. Other acts inherit the status of being overall beneficial (harmful) just in case they entail the relevant alternative act that is overall beneficial (harmful).

I have defended an account of what determines the relevant alternative set. When evaluating what is overall beneficial or harmful, the relevant alternative set is the set whose members exactly decide whether the subject has a life that is prudentially worth wanting out of those available. This set can always be written as $\{\alpha, \beta\}$ such that the subject has a life that is prudentially worth wanting out of the available lives just in case α is performed, and they have a life that is not worth wanting out of the available lives just in case β is performed.

In addition to playing a foundational role in determining overall harm and benefit, which lives are worth wanting out of those available is also linked to what there is sufficient prudential reason to do. If an act would result in a life that is (not) worth wanting out of those available, then there is (not) sufficient prudential reason for the agent to perform it.

Since sufficient prudential reason and overall harm and benefit both depend on which lives are worth wanting out of those available, it is important to illuminate this notion. In my view, *prudential similarity* (i.e., similarity in how the subject fares) determines whether a life is

worth wanting out of those available. A life is worth wanting out of those available just in case it is more prudentially similar to the best available lives than the worst.

Prudential similarity depends on how well a subject fares at a time and over time. A subject's welfare is therefore the ultimate value to which overall benefit and sufficient prudential reason are sensitive. Welfare is the underlying third factor that establishes harmony between benefit and prudential reason.

Acknowledgements

This paper would not have been possible without the help of many others. For their comments and discussion, I am grateful to Luc Bovens, Ben Bradley, Errol Lord, Z Quanbeck, Sarah Stroud, Daniel Wodak, Alex Worsnip, participants in the Fall 2023 University of Pennsylvania Dissertation Research Seminar, and two anonymous reviewers. I am grateful most of all to Kara Damon—for everything.

Bibliography

Bergström, L. (1966). *The alternatives and consequences of actions*. Stockholm: Almqvist & Wiksell.

Boonin, D. (2014). *The non-identity problem and the ethics of future people*. New York: Oxford University Press. doi: 10.1086/684703

Bradley, B. (2009). *Well-being and death*. New York: Oxford University Press. doi: 10.1093/acprof:oso/9780199557967.001.1

---. (2012). Doing away with harm. *Philosophy and Phenomenological Research, 85*, 390–412. doi: 10.1111/j.1933-1592.2012.00615.x

Brown, C. (2018). Maximalism and the structure of acts. *Noûs, 52,* 752–771. doi: 10.1111/nous.12181

Carlson, E. (2019). More problems for the counterfactual comparative account of harm and benefit. *Ethical Theory and Moral Practice, 22,* 795–807. doi: 10.1007/s10677-018-9931-5

---. (2020). Reply to Klocksiem on the counterfactual comparative account of harm. *Ethical Theory and Moral Practice, 23,* 407–413. doi: 10.1007/s10677-020-10071-6

Carlson, E., Johansson, J., & Risberg, O. (2021). Well-being counterfactualist accounts of harm and benefit. *Australasian Journal of Philosophy, 99,* 164–174. doi: 10.1080/00048402.2019.1705363

---. (2024). Prudential problems for the counterfactual comparative account of harm and benefit. *Philosophical Quarterly, 74,* 474–481. doi: 10.1093/pq/pqad082

Feit, N. (2015). Plural harm. *Philosophy and Phenomenological Research, 90*, 361–388. doi: 10.1111/phpr.12033

---. (2019). Harming by failing to benefit. *Ethical Theory and Moral Practice, 22*, 809–823. doi: 10.1007/s10677-017-9838-6

---. (2023). *Bad things: The nature and normative role of harm*. New York: Oxford University Press. doi: 10.1093/oso/9780197660447.001.0001

Feldman, F. (1991). Some puzzles about the evil of death. *The Philosophical Review, 100*, 205–227. doi: 10.2307/2185300

Groenendijk, J. & Stokhof, M. (1982). Semantic analysis of 'wh'-complements. *Linguistics and Philosophy, 5*, 175–233. doi: 10.1007/BF00351052

---. (1997). Questions. In J. Van Benthem & A. Ter Meulen (Eds.), *Handbook of logic and language* (pp. 1055–124). Amsterdam: Elsevier.

Hoek, D. (2022). Questions in action. *Journal of Philosophy, 119*, 113–143. doi: 10.5840/jphil202211938

Hurka, T. (1990). Two kinds of satisficing. *Philosophical Studies, 59*, 107–111.

Klocksiem, J. (2012). A defense of the counterfactual comparative account of harm. *American Philosophical Quarterly, 49*, 285–300.

---. (2019). The counterfactual comparative account of harm and reasons for action and preference: Reply to Carlson. *Ethical Theory and Moral Practice, 22*, 673–677. doi: 10.1007/s10677-019-10025-7

---. (2022). Harm, failing to benefit, and the counterfactual comparative account. *Utilitas, 34*, 428–44. doi: 10.1017/S0953820822000243

Lewis, D. (1973). *Counterfactuals*. Cambridge: Harvard University Press. doi: 10.2307/2273738

---. (1983). New work for a theory of universals. *Australasian Journal of Philosophy, 61*,

343–377. doi: 10.1080/00048408312341131

---. (1988a). Statements partly about observation. *Philosophical Papers, 17*, 1–31. doi: 10.1080/05568648809506282

---. (1988b). Relevant implication. *Theoria, 54*, 161–74. doi: 10.1111/j.1755-2567.1988.tb00716.x

McNamara, C. (2024). Causal decision theory, context, and determinism. *Philosophy and Phenomenological Research*, *109*, 226–260. doi: 10.1111/phpr.13021

McMahan, J. (1988). Death and the value of life. *Ethics, 99*, 32–61. doi: 10.1086/293034

---. (2002). *The ethics of killing.* New York: Oxford University Press. doi: 10.1093/0195079981.001.0001

Norcross, A. (2005). Harming in context. *Philosophical Studies, 123*, 149–173. doi: 10.1007/s11098-004-5220-3

---. (2020). *Morality by degrees: Reasons without demands*. Oxford: Oxford University Press. doi: 10.1093/oso/9780198844990.001.0001

Parfit, D. (March 29, 1988). What we together do. Unpublished manuscript, 1–33. URL: https://philpapers.org/archive/PARWWT-3.pdf.

Portmore, D. W. (2017). Maximalism versus omnism about reasons. *Philosophical Studies, 174*, 2953–2972. doi: 10.1007/s11098-016-0819-8

Slote, M. (1984). Satisficing consequentialism. *Proceedings of the Aristotelian Society, 58*, 139–163. doi: 10.1093/aristoteliansupp/58.1.139

Stalnaker, R. (1975). Indicative conditionals. *Philosophia, 5*, 269–286. doi: 10.1007/bf02379021

Timmerman, T. (2016). Your death might be the worst thing ever to happen to you (but maybe you shouldn't care). *Canadian Journal of Philosophy, 46*, 18–37. doi:

10.1080/00455091.2015.1125249

Yablo, S. (2014). *Aboutness*. Oxford: Princeton University Press. doi: 10.1515/9781400845989