

Probabilistic promotion and ability

Article

Published Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Open access

Elson, L. ORCID: <https://orcid.org/0000-0002-3013-8030>
(2019) Probabilistic promotion and ability. *Ergo*, 6 (34).
20200313. ISSN 2330-4014 doi:
10.3998/ergo.12405314.0006.034 Available at
<https://centaur.reading.ac.uk/75807/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.3998/ergo.12405314.0006.034>

Publisher: Michigan Publishing

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

PROBABILISTIC PROMOTION AND ABILITY

LUKE ELSON

University of Reading

We often have some reason to do actions insofar as they promote outcomes or states of affairs. But what is it to promote an outcome? I defend a new version of ‘probabilism about promotion’. According to Minimal Probabilistic Promotion, we promote some outcome when we make that outcome more likely than it would have been if we had done something (anything) else. This makes promotion easy and reasons cheap.

MANY theories of reasons can tell us which outcomes or states of affairs we have practical (normative) reasons to bring about. Some theories, for example, say that we have reason to bring about those states in which well-being or the satisfaction of our desires is maximised.

But telling us which outcome to bring about doesn’t yet tell us what to do. Let’s suppose that Rhys has a reason—of whatever sort—to obtain a nice new bike. Tomorrow, a fair raffle will be drawn. A raffle is a lottery where each ticket has a unique number, and all 500 tickets will be sold, though none has yet been sold. Rhys is first in line to buy tickets. The prize is a nice new bike. Setting aside tedious stipulations and complications, we can say that Rhys has some reason to win the raffle and thus the bike.

Let’s focus *only* on this reason. Again given some stipulations, if Rhys buys n tickets, then the probability that he wins the bike is $n/500$. We can see that he has a reason to buy all of the tickets—he has a reason to get the bike, and if he buys all of the tickets, then he’ll get the bike. Similarly, he has no reason (that we are focusing on) to buy *none* of the tickets—the only reason we are discussing is his reason to get the bike, and if he buys no tickets, then he won’t get the bike.

But what about the intermediate actions, buying somewhere between 1 and 499 tickets? The upshot of these actions is non-trivially chancy or otherwise

Contact: Luke Elson <luke.elson@reading.ac.uk>

probabilistic. All we have said so far is silent about such actions. It seems clear that Rhys's reason to get the new bike gives him some practical reason to buy more tickets rather than fewer: the more tickets he buys, the more likely it is that he gets the bike. He has most reason to buy all the tickets. He might of course have countervailing reasons to buy fewer tickets, but those are not our focus. To capture his reason for buying (say) 50 tickets, we need a theory of how reasons to bring about a state of affairs (Rhys wins a bike) engender reasons to do some action (buying 50 raffle tickets).

Somewhat stipulatively, let's say that his buying more rather than fewer tickets *promotes* the outcome that he wins the bike. And increasingly so, as he buys more rather than fewer. A theory of promotion tells us which actions you have some reason to do, and the relative strengths of those reasons: it gets us from reasons to bring about certain states of affairs, to reasons for particular actions.¹ Promotion tells us how much rational support there is for many actions (such as buying 257 raffle tickets).

This is not just a side issue. Philippa Foot famously claimed that 'irrational actions are those in which a man in some way defeats his own purposes, doing what is calculated to be disadvantageous or to frustrate his ends' (1972: 310). But if his end is that some outcome or state of affairs obtains, then which actions frustrate or defeat that end, and why? And which actions promote or guarantee that end?

This normative conception of promotion underlies my methodology. Promotion is, I claim, implicitly defined by two fairly obvious claims about promotion's connection to practical reasons.

The first claim is:

Absolute Reason-Promote. If S has a reason to bring about an outcome D, then: there is a promotional reason for S to A *iff* S's A-ing would promote D.²

Absolute Reason-Promote concerns *contributory* (or 'pro tanto') reasons, which count for or against some action (Dancy 2004: esp. 16–17). In particular, this principle says nothing about the strengths of the various reasons. An account of promotion that satisfies Absolute Reason-Promote will tell us which of Rhys's possible actions would promote the outcome that he wins the bike, and so which actions he has a reason to do. But it won't tell us which actions he has more reason to do than others (other than perhaps in trivial cases where there is no reason to A, and some reason to B).

1. An at-least-similar distinction crops up in many places, such as that between utility and expected utility.

2. Snedegar (2014: 46) defends a version of Absolute Reason-Promote, which he simply calls 'Promote'.

The second claim is:

Degree Reason-Promote. If S has a reason to bring about an outcome D, then: there is more promotional reason for S to A than to B *iff* S's A-ing would promote D to a greater degree than her B-ing would.

Degree Reason-Promote concerns comparisons between actions, in terms of how much reason there is to do them. In Dancy's terminology, it is *verdictive*, in that it concerns verdicts about the balance of reasons or rational permissibility, such as that 'we have more reason to do this than to do that, but most reason to do some third thing' (2004: esp. 16–17). Intuitively, we want our account of promotion to tell us that Rhys buying 400 raffle tickets promotes his winning the bike to a greater degree than his buying 300 tickets does. This will vindicate the thought that he has more reason to do the former than the latter.

My task in what follows is to find an account of promotion which can play its role in Absolute Reason-Promote and Degree Reason-Promote. Both of these principles are restricted to 'promotional' reasons: reasons which are grounded in the promotion of outcomes or states of affairs. I do not assume that all reasons are promotional, but since they are the topic of this paper, I will largely omit the qualifier.

'Promote' is a relatively commonplace word in English, and we might hope for an account of promotion that is at least recognisably similar to its everyday meaning. But my implicit-definition methodology could take us quite far from this everyday sense, since I'll be—via the two Reason-Promote principles—relying on claims about which reasons there are. I don't think this need be a problem, since there's little pressure to think that 'promote' is totally univocal, and we shouldn't be afraid to depart from the everyday in search of a good account of the normative. I'll argue that the best theory of our everyday talk of promotion and the best normative theory of promotion are recognisably similar, but they do differ in the details.

Except in one explicitly-marked case (in §7.1), I am concerned throughout with cases where there is only *one* outcome that provides or grounds promotional reasons. Stipulatively, the only outcome an agent S has some reason to bring about is D. So when I write, for example, that S has stronger reason to A than to B, this means that she has stronger reason to A than to B *in virtue of having reason to bring about the outcome D*.

Finally, often the outcome or state of affairs to be promoted is the satisfaction of a desire (and I will often write 'promote a desire' rather than 'promote the satisfaction of a desire' or 'promote the outcome or state of affairs in which a desire is satisfied'). But the discussion is intended to be fully general.

Cases such as the raffle tend to support 'probabilism about promotion', which says that to promote an outcome or state of affairs simply *is* to raise its probability.

Rhys has practical reasons to take actions—such as buying more tickets—that increase the probability that he wins the bike. Moreover, at least the way I presented it, promotion only became an issue when we made the upshot of actions probabilistic.

In this paper, I'll defend a new probabilist theory of promotion.

1. Probabilism and its Discontents

Probabilism is extremely intuitive, but has proved much easier to state as a general thesis than to defend in detail. Here it is in barebones terms:

Probabilism about Promotion. An agent S 's action A promotes the outcome D iff $p(D|A) > X$.

As Jeff Behrens and Joshua DiPaolo have argued, a probabilist needs two things. The first is an interpretation of the probabilities involved (Behrens & DiPaolo 2016). I won't discuss this question, but will simply talk of 'probabilities', where these should be understood as objective chances or some highly idealised credences.

The second requirement—the baseline, X —has proven more troubling. Promotion involves increasing the probability of an outcome. But compared to what? Each proposed baseline faces apparent counterexamples, many of which descend from cases described by Behrens and DiPaolo (2011).³ Without a baseline—which will allow us to say whether promotion is occurring—we cannot fill in Absolute Reason-Promote.

Mark Schroeder influentially suggested that the baseline is the probability of D conditional on the agent 'doing nothing—conditional on the status quo' (2007: 113). Here is probabilism with a 'do nothing' baseline:

Do Nothing. S 's A -ing promotes D iff $p(D|A) > p(D|S \text{ does nothing})$.

(I discuss the status quo below.) Evers (2009: esp. 60) makes two important criticisms of Do Nothing. First, it's not especially clear what 'doing nothing' amounts to. The same (in)action might variously be described as doing nothing, standing firm, being patient, and so on. I don't think this criticism need be decisive: perhaps imprecision or ambiguity in 'doing nothing' is reflected in our reasons.

But the second criticism is rather more worrying: according to Do Nothing, there can never be a reason to do nothing. The probability of a given outcome

3. For accounts of promotion and further counterexamples, see, e.g., Coates (2013), Sharadin (2015), Manne (2016), and Lin (2018).

conditional on you doing nothing would have to exceed the probability of that outcome conditional on you doing nothing. But that is clearly not going to happen. And intuitively, we can have a reason to do nothing: if I tell you that you will receive £1000 if you do nothing for ten minutes, but won't get any money if you do any (positive) action during that ten minutes, it's at least extremely plausible that this is a reason for you to do nothing for ten minutes (cf. Behrens & DiPaolo 2011: 4ff.).

This is an instance of a general problem for accounts of promotion, which I'll call the *sticky baseline*: since promotion involves exceeding some baseline, there can be no reason to merely reach the baseline, or to remain there. Much discussion of promotion consists of stickiness objections to proposed baselines.

The claim that we can never have a reason to do nothing is counterintuitive—though we might get pushed into it by the failure of alternatives. Later, I'll suggest that Do Nothing has something to be said in its favour, especially as an account of our ordinary talk of promotion.

So what alternatives are there? Many of the most important proposed baselines are what I'll call *disposition-sensitive*: whether *S*'s A-ing promotes *D* depends on *S*'s dispositions to act. For example, Finlay (2006: 8 fn. 19) proposes that you promote an outcome by some action if the outcome is more probable given that action, than it would have been if you hadn't done that action:⁴

Finlay's Counterfactual Baseline. *S*'s A-ing promotes *D* iff $p(D|A) > p(D|\neg A)$.

This baseline is intuitive: if some outcome is more likely if I buy a ticket than if I don't buy a ticket, then surely in buying a ticket I promote that outcome? But here is an apparent counterexample—another instance of the sticky baseline:

Buttons. Debbie has some desire. There are three buttons in front of her. If she pushes either button A or button B, her desire is guaranteed to be fulfilled. If she pushes button C or does nothing, her desire will not be fulfilled. Debbie in fact pushes A. Had she not pushed A, though, she would have pushed B instead.⁵

Let's suppose that—as seems clear—in pushing A, Debbie *does* promote the satisfaction of her desire.

Finlay's baseline $p(D|\neg A)$ is disposition-sensitive. This baseline's verdict depends on what would (probably) have happened had Debbie not pushed A,

4. Following Behrens and DiPaolo (2011: 1 fn. 1), I'll interpret Finlay's 'is conducive to' in terms of promotion.

5. Adapted from Behrens and DiPaolo (2011: 2). I have adapted the case by explicitly clarifying what happens if Debbie does nothing.

including what Debbie would (probably) have done instead. The core idea of Buttons is to rig this baseline, by stipulating that had Debbie not pushed A, then she would have pushed B, and her desire would have been satisfied anyway. So Debbie's desire is not more likely to be satisfied if she pushes A than if she does not push A.

Buttons has structured much of the promotion literature. Finding a plausible baseline which vindicates the intuitive verdict—that Debbie promotes her desire—has proven surprisingly troublesome.

Another influential baseline proposal is Schroeder's status quo. But how to understand this baseline in probabilistic terms? Most naturally, as the current or antecedent probability of D. David Sobel interprets Schroeder this way—'a desire for p explains a reason to A iff A-ing makes p more likely than it already is' (Sobel 2016a: 305)—and Eden Lin has recently argued that a version of this view may be defensible:

Simple Probabilistic Promotion (SPP). *S's doing A promotes D iff it makes D more likely to obtain than it was prior to the occurrence of A.*
(Lin 2018: 363)

SPP says that one promotes D by *causing* (making) an increase in the probability of D. If $p_0(\cdot)$ describes the probabilities before S acts, and $p_1(\cdot)$ describes them after she acts, then D is promoted by her action only if $p_1(D) > p_0(D)$. (Not 'if and only if', because the probability could have changed without S causing that change.)

In Buttons, Debbie wouldn't have pushed C, had she not pushed A (she would have pushed B instead). But, as Lin (2018: 367) argues, beyond this counterfactual claim, the case is underspecified: was there *any probability* that Debbie would push C, had she not pushed A? If there was No Probability that she would, then $p_0(C) = 0$ and D was antecedently certain. So $p_0(D) = 1 = p_1(D)$, and SPP says that Debbie's pushing A did not promote D. But if there was Some Probability that she would, then $p_0(D) < 1$ but $p_1(D) = 1$, so according to SPP, Debbie's A-ing promoted her desire. Lin claims that both of these verdicts are correct.

Thus, SPP is disposition-sensitive, because whether pushing A promoted Debbie's desire depended on this counterfactual probability, which is in part determined by Debbie's propensities to act. But, I'll argue, here lies the problem for SPP. Between the No and Some Probability versions of the case, the *only* difference lies in Debbie's dispositions to act (let's suppose). Nothing about her desire that D, or the causal consequences of pushing each button, changes. Yet whether pushing A promotes D—in other words, whether she has a reason to push A—does change, according to SPP.

To see the implausibility of this, let's set aside button B and doing nothing, and consider a version of the case where the only buttons are A and C. Call this

case Two Buttons.

SPP says that Debbie has a reason to push A if and only if $p_0(A) < 1$. This antecedent probability can be arbitrarily close to 1, and Debbie has a reason to push A. But if $p_0(A) = 1$, then $p_0(D) = p_1(D)$ and her reason to push A evaporates. Simply by altering Debbie's dispositions to act, we can make her reason evaporate. Thus according to SPP, there can never be a situation where the causal upshot of A is certain, S has a promotional reason to A, and S is certain to act in accordance with that reason. Reasons are ephemeral—such probabilistic certainty dissolves the reason—which ironically undermines the possibility of giving a normative explanation for why the agent was certain to act as such: we cannot say that it was because she was certain to act in accord with that reason.

Lin of course sees the problem of cases where 'p will obtain if and only if you do A, and there is a probability of 1 that p will obtain because there is a probability of 1 that you will do A' (2018: 377). Intuitively, such an agent would have a reason; SPP cannot say this. His response is that it is unclear whether there are any such cases. But this response is unsatisfying: stipulated cases are common in this dialectic, and such a case is no more artificial than Buttons with No Probability. The stipulated probability is not so odd or distorted that the case can be dismissed as not probative. And the problem extends beyond cases where the causal upshot of A-ing is certain: if $p_0(D|A) = 0.7$, and Sarah is *certain* to push button A, then $p_0(D) = p_1(D) = 0.7$, and so her pushing button A doesn't promote D.

What I'll call 'three-action cases' illustrate a general problem with disposition-sensitive accounts.⁶ Suppose that Rhys could buy ten raffle tickets, or one, or none. He currently has none. Does buying one ticket promote his goal of winning the bike? Intuitively, yes. According to Do Nothing (which is not disposition-sensitive), yes.

According to SPP, we don't yet know whether it does. If Rhys was antecedently very likely to buy no tickets, then buying one ticket promotes the outcome; if he was antecedently very likely to buy ten tickets, then buying one ticket doesn't promote the outcome. According to SPP, whether an agent has a promotional reason to do some action depends on which actions are available to him, what the consequences would be of each action, and how likely he is to do each action. But the last of these seems just plain irrelevant. Lin (2018: 377–378) argues that such cases need not be a decisive objection to SPP—it may be worth biting the bullet if a convincing account of reasons can be built around the view.

On the other hand, many contrastivists—who reject the existence of an invariant baseline—take the opposite lesson from such cases. For Snedegar and Nathaniel Sharadin and Finnur Dellsén, three-action cases refute disposition-sensitive accounts—though they do not put things in those terms—and motivate

6. See also Lin (2018: 377–378) and Snedegar (2017: 74 n. 18) for such cases.

a move away from an invariant baseline, to contrastivism.⁷

I wish to steer a middle course. Three-action cases are a problem for invariance, and cannot simply be ignored, but don't yet push us towards contrastivism, which has its own problems.

In fact, a last-ditch defence can be made on behalf of SPP. Let's appeal to the distinction between the contributory and the verdictive, and extend SPP to include *dis-promotion*, and associated reasons against actions:

Dis-Promotion for SPP. S's doing A dis-promotes D if and only if it makes D less likely to obtain than it was prior to the occurrence of A.

Absolute-Reason-Against-Promote. There is a reason against A-ing if and only if A dis-promotes D.

Putting these claims together, in No Probability Buttons, pushing A doesn't promote D, and Debbie has no reason to push A. But pushing C would dis-promote D, and so she has a reason against pushing C. So she has no reason either for or against pushing A. But she does have a reason against pushing C. From these contributory claims, a verdictive implication seems undeniable: of the two buttons, she has most reason to push A, and least reason to push C; she ought (of most reason) push A, and she ought not push C. SPP can thus rescue the verdictive claim that Debbie has most reason to push A. But we must still accept the implausible contributory upshot that she does not have a reason to push A, even though pushing A will ensure the satisfaction of her desire.

The strategy can be applied to three-action cases. Of the three options, Rhys has *most* reason to buy ten tickets, and *least* reason to buy none. He has some intermediate amount of reason to buy just one ticket. His dispositions may determine the contributory reasons for and against each action, but they are irrelevant to the verdictive claims.

This defence is, however, not worth the price of such ephemerality of contributory reasons. Ephemerality highlights a philosophical tension in disposition-sensitive accounts. On the one hand, there are genuine contributory reasons for and against actions—if not, why specify a baseline?—but on the other, these reasons are ephemeral, and depend on how the agent is disposed to act. If we are in the business of providing an invariant baseline, and so dealing with non-contrastive contributory reasons, we should look for a baseline that doesn't lead to such ephemerality.

I have argued against disposition-sensitive accounts, especially SPP. But SPP has manifest intuitive plausibility. What could be more natural than saying that to promote an outcome is to make it more probable than it already is?

7. I discuss contrastivism further in Section 2.1.

Ultimately, however, the forward-looking nature of probabilities undermines this plausibility. Because the antecedent probability of D is in part determined by the future, including perhaps the fact that an agent S *will* ensure that D obtains, it is unsuitable as a baseline for determining whether S's actions promote D.

I've argued that each of the proposed probabilistic baselines succumbs to the problem of the sticky baseline. As a preview for later, I'll say that if we make the baseline as low as is reasonably possible, its stickiness is no longer a problem. The baseline will be so low—and hence promotion so cheap—that we will (almost) never have cause to say that though some action intuitively promotes an outcome, the baseline fails to respect this.

Of course, one lesson that could be drawn from the problems with SPP and the other probabilistic baselines is that probabilism is misguided. I'll now consider and reject a recent non-probabilist account of promotion.

2. Promotion as Fit

Consider impossible desires. These are desires whose probability of satisfaction is zero, where that probability cannot be raised above zero. (As I'll argue below, these two claims do not quite amount to the same thing.) Probabilism about promotion says that an outcome can be promoted only if its probability can be raised against some baseline. Probabilism implies that impossible desires cannot be promoted, because their probabilities cannot be raised at all. So impossible desires can't ground (promotional) reasons.

In a series of papers, Nathaniel Sharadin argues that an outcome can be promoted even if its probability cannot be changed.⁸ If this is right, then probabilism is false. Sharadin and Dellsén argue instead that promoting a desire is a matter of increasing the *fit* between the world and that desire, where fit is the 'match between the desire's content and the way the world is' (2019: 1277). The fit account allows for the promotion of impossible desires insofar as the actual world can be made to more closely (albeit partially) fit such a desire's content.

There are two issues here—the promote-the-impossible attack on probabilism, and the fit account in particular—but they are obviously connected. The apparent promotion of impossible desires is at the heart of both. After spelling out the Contrastive Expected Fit view in more detail, I'll argue that we cannot promote impossible desires, and offer some criticisms of the fit view.⁹

8. See, for example, Sharadin (2015) and Sharadin (2016).

9. My notation differs slightly from that of Sharadin and Dellsén, to be more consistent with the rest of this paper.

2.1. Contrastive Expected Fit

Sharadin and Dellsén give the fit relation an almost entirely intuitive characterisation, together with a few examples. The intuition is fairly clear: some possible worlds fit (a desire that) D better than others do, and some pairs of worlds fit D equally well. This lets us define an equivalence relation on the set of possible worlds:

D-outcome. A D-outcome o_i is a set of possible worlds such that all worlds in o_i fit D equally well.

We then define a fit function $F(o_i, D)$ from D-outcomes to the interval $[0, 1]$, which describes how well the worlds in that outcome fit D (Sharadin & Dellsén 2019: 1281–1282).

But if I do some action A, it might be non-trivially probabilistic which outcome will obtain—non-trivially probabilistic how well the world will fit D. For example, if A is buying one raffle ticket, and D is winning the raffle, then there is some probability $p(o_1|A)$ that I win the raffle, and D-outcome o_1 , where $F(o_1, D) = 1$, obtains. There is a somewhat larger probability $p(o_2|A)$ that I lose the raffle, and D-outcome o_2 obtains, where $F(o_2, D) = 0$. To cope with this, we define the expected fit $EF(A, D)$ between an action A and a desire that D. Let O be the set of all D-outcomes. Then:

$$EF(A, D) = \sum_{o_i \in O} p(o_i|A) \cdot F(o_i, D)$$

This function weights the D-outcomes by how probable they are if I do A. Now, we can state Contrastive Expected Fit account:

Contrastive Expected Fit. A-ing rather than B-ing promotes D *iff* A's expected fit with D is greater than B's expected fit with D. (Sharadin & Dellsén 2019: 1282)

Suppose that I have some reason to bring about D. Then if o_X is the outcome of my doing X, the fit account says that as $F(o_X, D)$ increases, so does the strength of my reason to X. In particular, and in more precisely contrastive terms, if the causal upshot of A-ing better fits D than the causal upshot of B-ing does, then I have more reason to A than to B.

The account has implications concerning desires which are *not* impossible. Plausibly and correctly, given normal assumptions, Contrastive Expected Fit says that buying one ticket rather than buying no tickets promotes my desire to win the raffle.

But Contrastive Expected Fit also implies that there can be no desires which provide 'all or nothing' reasons, where there is a very strong reason to A *iff*

$F(o_A, D) = 1$, but no reason whatsoever to A iff $F(o_A, D) \neq 1$. The account also implies that there can be no uncanny valley, where there is strong reason to completely fit some desire or leave it completely unfitted, but weak or no reason to fit it moderately well. These consequences are not obviously incorrect, but are at least in need of defence.

Contrastive Expected Fit is contrastive. We may compare actions by how well they promote a state of affairs, but there is no invariant baseline for comparisons. Contrastive Expected Fit does not (strictly speaking) answer cases such as Buttons: it does not give a definitive answer to the question of whether Debbie promotes her desire by pushing A. Whether she does depends on the contrastive baseline which is specified by context (Sharadin & Dellsén 2019: esp. 1273). Though this might be a pill we can swallow, it is a disadvantage for the view, because it seems intuitively clear that Debbie does promote the satisfaction of her desire, and without some serious argument an appeal to context here seems unsatisfying.

Of course, Sharadin and Dellsén do provide arguments for a move to contrastivism. These arguments amount to the claim that an invariant account of fit—one with a baseline—would ‘face exactly the same sorts of trouble faced by non-contrastive *probabilistic* accounts of promotion’ (Sharadin & Dellsén 2019: 1283, esp. n. 45). In other words, Buttons-type and three-action cases. Similarly, Snedegar (2014: 49; 2017: 89) defends a moderate contrastivism (where reasons are contrastive, but the promotion relation is invariant) by appeal to the apparent failure of invariantism about reasons. So the primary motivation for contrastivism is the apparent failure of invariantism. If we can provide a successful invariantist account, the main extant motivation for contrastivism is undercut.

2.2. *The Impossible Promotion Argument*

Sharadin and Dellsén rely on examples like the following, to argue that impossible desires—which are impossible to satisfy—can nevertheless be promoted:

Extreme Ascetic. The desire that none of your desires are satisfied. (Sharadin & Dellsén 2019: 1274)

Let’s accept for the sake of argument that this fairly unusual desire cannot be satisfied. Sharadin and Dellsén argue that Extreme Ascetic provides reasons for action (which implies that it can be promoted). In particular, they argue that if an agent Agatha has Extreme Ascetic, then if she is offered a desire-frustration pill which will ensure that fewer of her desires are satisfied than at present—but not none of them, which is impossible—then she has a reason to take the pill.¹⁰ If they are right, then this looks fatal to probabilism, which says that Extreme Ascetic cannot be promoted, and thus that it provides no reason to take the pill.

10. I’m grateful to Robert Mason for catching a mistake in a previous version of this case.

But the fit account could explain such a reason rather well: Extreme Ascetic cannot be satisfied, but the world can be made to better fit the desire, thereby promoting it. So if Extreme Ascetic can be promoted, then not only does this show that impossible desires can be promoted, sinking probabilism, but also that fit is a promising theory of their promotion. So if it can be promoted, then Extreme Ascetic does double duty in defence of the fit account.

So how can a probabilist account for any reasons in the ascetic case? Extreme Ascetic won't do the job, so if there are reasons, then they must be grounded in some other desire. (I write that a reason to A is 'grounded in' some outcome if what explains the reason is that it would promote that outcome.) In particular, the probabilist is likely to say that the reason for her to take the pill is grounded in a 'neighbourhood desire' such as:

Comparative Ascetic. The desire that as few of her desires as possible are satisfied. (Cf. Sharadin & Dellsén 2019: 1275; DiPaolo & Behrends 2015: 5)

So if Agatha has *only* Extreme Ascetic (and not the neighbourhood desire Comparative Ascetic), then probabilism says that she lacks any reason to frustrate some of her desires.

To further draw out the implications, suppose that Agatha has 100 desires, and is offered a choice between *two* pills. The better pill would frustrate 51 of her desires; the worse pill would frustrate only 50. If impossible desires can be promoted, and especially if this can be done by fit, then Agatha plausibly has a stronger reason to take the better pill than the worse one. Probabilism, however, denies that she has a reason to take either pill, and says that the reasons she has to take either pill are equally strong (i.e., not at all). Sharadin and Dellsén find this 'bewildering' (2019: 1276); for the record, I find it wholly natural.

This is a stalemate, so I'll now pursue a different line of argument.

2.3. *The Structure of the Examples*

The examples used by Sharadin and Dellsén are very friendly to the fit account. They involve desires which are typically—in psychologically normal agents—accompanied by neighbourhood desires.

Sharadin is of course not blind to the possibility of a neighbourhood desire explanation, and rightly asks on what grounds the probabilist can assume that one must exist.¹¹ But he doesn't consider the answer I have in mind: that in psychologically typical agents, Extreme Ascetic would normally have some kind of explanation—an explanation that also explains the neighbourhood desire. In other words, there is a third desire or other psychological state, which is a common

11. See DiPaolo and Behrends (2015: 5) and Sharadin (2016: 3ff.) for discussion of this move.

explanation of Extreme Ascetic and Comparative Ascetic.

I'll argue that this is so, by working through the cases.

Why would someone hold Extreme Ascetic? A very common explanation would be a proto-Buddhist religious conviction that desire satisfaction is to be avoided. But it would be odd to have Extreme Ascetic on such grounds, without also having the neighbourhood desire Comparative Ascetic: if one thinks that desire satisfaction is to be avoided, this will likely motivate one to desire that as few as possible of one's desires are satisfied.

Similar points can be made about an earlier example of Sharadin's:

Live Forever. Suppose an agent desires to live forever and is offered a pill that will extend her life by a thousand years. The agent has a reason to take the pill. And this is so because taking the pill promotes her desire to live forever. (Sharadin 2015: 1379)

Why desire immortality? This is not an easy philosophical question, but standard explanations include a desire to avoid the pain and indignity of dying, and a desire to enjoy as much as possible of what life has to offer. Both of these explanatory desires would also clearly explain a neighbourhood desire for a thousand extra years of life. The former explains it via temporal discounting—if we must undergo something we fear, the further in the future, the better—and the latter explains it more directly: if we wish to enjoy what life has to offer, then a thousand extra years of that is most welcome.

Granting such commonplaces about the psychology of normal agents, we have independent grounds to posit explanatory desires and accompanying neighbourhood desires in these cases. One could in principle have Extreme Ascetic without Comparative Ascetic, for example, but this is highly unlikely in normal agents. Now we can of course try to stipulate that the relevant neighbourhood desire is absent. But it would be an odd person who held Extreme Ascetic as a basic or unexplained desire, and our intuitions about such people are not likely to be trustworthy or even widely shared. I think this is at the root of the stalemate.¹²

A better strategy is to construct cases where it is independently plausible that there is no neighbourhood desire, without resorting to excessive stipulation.

I'll choose a dystopian case. I am in Montreal, and have a strong desire to see the Pacific Coast of Canada, because I wish to see the sun set over the ocean, which of course I can only do on the West coast. In particular, I wish to see the Pacific sunset after a long roadtrip—after driving across the country. (Background psychological fact: I like driving, but find it very tiring, and so I'm particularly fond of ocean sunsets after long drives.)

12. Compare Sobel (2016b) on amoralism and subjectivism about reasons: agents with no subjective reasons to treat others well would be quite unusual.

Unfortunately the world has run out of fuel, and I have the last supply. I only have enough to get me as far as Saskatchewan. Assuming it is certain that the desire is impossible, there is clearly no reason to drive to Saskatchewan. Without some other justification it would be irrational, since doing so takes time and effort, especially during the apocalypse. So there seems to be no promotion here. Moreover, the analogue of the better versus worse pill argument doesn't get any purchase either. I have no more reason to drive to western Saskatchewan than to eastern Saskatchewan.

In this new example, we have no promotion and no reason. The original cases of impossible desires brought with them (in psychologically typical agents) neighbourhood desires which everyone agrees can be promoted. But the new case is constructed to avoid this.

For probabilism, driving to Saskatchewan promotes my desire only insofar as doing so raises the probability that the desire will be satisfied. In the actual world, this is likely the case: fuel is widely available in rich countries. So setting off from Montreal with insufficient fuel (inevitable, if I have a normal vehicle with a normal fuel tank) raises the probability that I'll make it to the West coast, because there's a good probability that I'll be able to buy more along the way.

But in the fuel-depleted apocalypse, this is not so. If the petrol stations are all deserted, and there is *no* probability that I will be able to get my car from Saskatchewan to the West coast, then I have no reason to drive halfway across Canada. This seems correct—doing so would be pointless, futile. So probabilism neatly explains why in the new example there is no reason to drive halfway across Canada. There is a disparity between the old and the new cases, which probabilism can neatly explain in terms of neighbourhood desires.

I don't wish to claim that the fit account *cannot* explain this disparity (or explain it away), but here are some reasons to be skeptical that this can be done. Much would depend on the details of the fit relation—one option would be to spell out that relation in such a way that the drive wouldn't increase fit. But if the fit relation is simply the similarity of possible worlds, then it seems hard to deny that a world in which I drive halfway across Canada better fits a world in which I drive all the way across Canada to see a sunset than a world in which I remain in Montreal does. Any spelling out would need to leave untouched the fit verdicts that Sharadin and Dellsén appeal to—for example, that frustrating half of one's desires increases fit with an outcome in which none of one's desires are satisfied.

Sharadin and Dellsén could instead say that the intuition concerning the sunset is wrong, and there is a reason to drive to Saskatchewan. (Or in more contrastively-respectable terms: that a world in which I drive to Saskatchewan better fits my desire to see a sunset over the ocean than does a world in which I stay in Montreal.)

But this strategy becomes less promising once we see the trick to constructing

new cases of this general sort. We simply specify a goal—that I reach another star in my lifetime; that I gather 100 billion people in a room—which is acknowledged to be impossible. Then consider actions which would go *some way* towards fitting that goal—climbing up a tall building; gathering 100 people in a room—and ask whether the goal provides any reason to do that.

Suppose that the desire to visit another star in one’s lifetime is the strongest, most important desire in your life. If you *could* get to another solar system, then you would have very strong reasons to do many things in pursuit of that goal. But the goal is known to be impossible. Suppose that at the moment, the Earth’s position is such that if you climb up a tall university building, you will be marginally further away from the Sun, and closer to some other star. Does your desire give you any reason to go up in the lift to the top floor? I think, clearly not. But this case has a similar structure to Extreme Ascetic (getting a few hundred feet closer to another star is structurally similar to frustrating a few of your desires, though of course on a different scale). The fit account can bite the bullet and say that there are reasons in such cases, but this means that there are very many reasons indeed, to do many things in pursuit of knowably impossible goals.

Probabilism is in a dialectically better situation here. Probabilism’s problem was the lack of reasons in certain cases; this can be remedied, as we have seen, by adding background or neighbourhood desires. The reader will have to judge for herself how plausible my general psychological claims about the presence of such desires are. But the fit account faces a problem of there being too much promotion: promotion and attendant reasons in cases where this is implausible. And this problem is not something that can easily be tweaked by adding or removing desires, because it rests on fundamental features of the account. (This line of argument may seem a little ironic when I come to defend my own view, below. Needless to say, I think that’s different.)

3. Causal Impotence and Promotion

I’ve just criticised a non-probabilist account of promotion, and argued that impossible desires cannot be promoted—something which probabilism neatly explains. This seems to count in favour of probabilism. But before that, I criticised a recent probabilist account. So I’ll now defend a new probabilism about promotion.

First, I’ll explain the underlying picture. I will appeal to a notion of those actions—including inaction, or doing nothing—that the agent *can* do, which are available to her in some sense. Let’s suppose that at each moment, there are a number of (exclusive) actions that each agent *S* can do. Her ‘ability set’ $F = \{A_1, A_2, \dots, A_n\}$ contains these actions. *F* will typically include a member that corresponds to her doing nothing.

F's membership need not be wholly determinate. Does it include actions that someone is prevented from taking by compulsion or mental illness? The reality is no doubt messy, but such messiness is familiar in this area: ought (can) the addict avoid alcohol? It seems hard to give a general answer to this question.

It might also be wondered how the members of F are individuated. For example, does Debbie's ability set have as a subset just {push A}, or (also) {push A instead of B, push A instead of C}? Clearly a full discussion of the individuation of actions is beyond the scope of this paper, but it's plausible that this question doesn't matter: as probabilists, we care only about conditional probabilities, and $p(D|A) = p(D|A \text{ instead of } B) = p(D|A \text{ instead of } C)$.

These are just a quick responses, and I acknowledge that hidden dangers may lie here. But we are most interested in cases where F's membership is stipulated. With such caveats, let's grant the notion of an ability set. Ability is not a probabilistic notion, and it differs extensionally from probabilistic notions. We can see this by considering the biconditional 'necessarily, S can A iff the probability that S does A is nonzero'.

The right-to-left direction seems clearly true. It's hard to see how there could be non-zero probability that S As, if S can't A.

But the left-to-right direction is false in cases where S can do something, but is probabilistically certain not to. We saw this in the No Probability version of Buttons. I'm not sure whether there are any actual such cases, but they seem at least coherent, and it's hard to see a principled argument against the possibility. We want our account of promotion to render plausible verdicts in such cases—if possible.

If S cannot affect the probability of some outcome D, then we'll say that she is 'causally impotent' with respect to D:

Causal impotence. S is causally impotent with respect to some outcome D iff for all A_m and A_n in her ability set F, $p(D|A_m) = p(D|A_n)$.

In other words, S is causally impotent iff the probability of D does not depend on what she does now, of the things she can do. Obvious instances of causal impotence involve impossible desires—whatever I do, the probability of their satisfaction remains 0—as well as outcomes concerning necessary truths and (absent backwards causation) facts about the past. More prosaically, I am causally impotent with respect to the decay of some distant atom in the next 30 seconds, the explosion of the Moon, and the outcome that the next person to order soup in central London will order minestrone. (I am not in London.)

I claim that in general, S can promote an outcome D iff she is not causally impotent with respect to D. S can promote D iff there are actions A and B in F, where $p(D|A) \neq p(D|B)$. Promotion requires causal potency. The thought behind this is—remembering that 'promote' here is stipulative and defined by

its connection with reasons—if the probability of D is the same whatever S does, then how can there be a reason to do any particular action rather than another? S cannot influence whether D comes about, so D should be irrelevant to her deliberation.

I argued above that impossible desires can't be promoted; this also follows from the general claim. That general claim is perhaps most likely to run into resistance in cases of inevitable outcomes, where *whatever* the agent does, some outcome has probability 1.

There are three possibilities—and in none does promotion occur.

First, suppose that whatever S does, the outcome D obtains, and her actions are outside the causal chain leading to D. It should be fairly uncontroversial that there is no promotion in such cases. For example, we do not promote the outcomes that necessary truths obtain, or that facts about the past continue to obtain.¹³

Second, suppose that D will occur whatever S does, but it's under her control whether any action of hers is part of the causal chain that leads to D. For example, consider the outcome that Smith is dead by the end of the century. If you shoot him, he dies today. If you do not shoot him, he dies of old age later this century. So the probability of his death this century is 1 whatever you do. I find it wholly obvious that your shooting him does not promote his death this century, though of course you might promote other outcomes, such as that he dies *today* or at your hand.

But you can be the proximate cause of his death, and it might be objected that clearly if one's action is the proximate cause of how an outcome obtains, then one promotes that outcome. I think here we simply have a clash of intuitions: if you accept this objection, then you are not a probabilist. SPP and Finlay's counterfactual view agree with mine on this case, for what it's worth. Indeed, Lin explicitly rejects his *Sufficiency Principle*: 'S's doing A promotes p if it causes p to obtain' (2018: 367). For Do Nothing, whether you promote his death this century depends on whether his death was caused by a positive act of yours, or simply by your doing nothing—another implausibility for that baseline. Depending on one's moral theory, one might still be blameworthy for shooting Smith—and there might also be non-promotional reasons not to shoot him. These are separate questions.

Thirdly, and finally, suppose that *whatever* S does, her actions are part of the causal chain leading to D. I claim that in such cases, S doesn't promote D. The upshot is that we cannot be forced to promote an outcome: if we are put in a situation where whatever we do, the probability of the outcome is the same—

13. This, I think, is a counterexample to Sharadin's earlier 'Behrends-DiPaolo Constraint', which implies that nearly everything we do promotes such outcomes. See Sharadin (2015: esp. 1374).

perhaps because it is certain—then we do not promote that outcome. This claim is at least somewhat counterintuitive (to some, if not always to me). Once again, I think, a probabilist simply has to accept the verdict. But we can mitigate its surprisingness by considering more causally distant examples. It is plausible that any action I can do now is part of the causal chain which leads to the heat death of the universe. Did my pacing the room just now promote the heat death of the universe? I hope not.

I'll now defend a probabilistic account of promotion.

4. Ranking Probabilism about Promotion

When S is not causally impotent with respect to D, then she can promote D, and in such cases we can rank the things she can do by D's conditional probability:

Ranking Probabilism about Promotion. A_m promotes D to greater degree than A_n does iff $p(D|A_m) > p(D|A_n)$.

There will often be ties, where $p(D|A_m) = p(D|A_n)$. But if the case is not one of causal impotence, then not every member of the ability set will be tied.

Ranking Probabilism implies that S has more reason to A_m than to A_n if and only if $p(D|A_m) > p(D|A_n)$. That is, if and only if A_m ranks higher than A_n in terms of the conditional probability of D. In particular, some action(s) will rank bottom, though they need not have conditional probability 0. S has least reason to do such actions. Similarly, S has most reason to do those action(s) which rank top—which need not have conditional probability 1.

We can now see why in the Extreme Ascetic case, Agatha has no more reason to take one pill than another: the probability of her desire's satisfaction is 0 whether she takes the better pill or the worse one.

My primary claim is that when we are considering only one desire or outcome, Ranking Probabilism provides all the verdictive normative facts: it tells us what we have most reason to do, what we have least reason to do, and for any act that we can do, it tells us how it compares to the other acts in our ability set. And—other than the definition of the ability set itself!—this is cashed out in purely probabilist terms. We could stop here, and simply adopt Ranking Probabilism (without a baseline). But this is too revisionary—a last resort—and so I'll instead try to defend a plausible baseline.

5. Flat Probabilism

With the idea of an ability set in hand, a natural baseline would assume that each action in the ability set is equally likely to be done:

Flat Probabilism. S 's A -ing promotes D if and only if $p(D|A) > \frac{1}{|F|} \sum_{A_i \in F} p(D|A_i)$.

According to Flat Probabilism, the baseline probability of D is the mean of the probabilities $p(D|A_m)$, for all $A_m \in F$. Only actions which are above average (mean) in terms of the probability of D will promote D .

To my knowledge, nobody defends Flat Probabilism, though Stephen Finlay defends an account of 'ought' which excludes 'any information about the agent's dispositions to choose one means over another'. Instead, each action is assumed to be equally likely: 'relative to such a background each alternative has equal probability' (Finlay 2014: 73).¹⁴ We can easily construct counterexamples to Flat Probabilism as an account of promotion. (To be clear, this is not a problem for Finlay: the counterexamples to follow involve sub-optimal actions, and as such don't affect Finlay's view about 'ought', which concerns optimal actions.)

First, in Rhys's raffle, assume that Rhys can buy any number of tickets, up to 500. Flat Probabilism says that buying (exactly) 10 tickets doesn't promote his goal of winning the raffle, because there was antecedently equal probability of him buying every number of tickets. But it's intuitively implausible that buying 10 tickets doesn't promote his winning the raffle.

Second, multiplying actions near the bottom of the ranking affects whether actions near (but not at) the top of the ranking promote some outcome. Flat Probabilism is disposition-insensitive, so it doesn't face a three-action case problem, but the structure here is similar. In everyday cases, this can lead to absurd results. Suppose that I give you £100, £20, or £1 to buy lottery tickets, in decreasing order of the conditional probability of D . If the probabilities are appropriately specified, then Flat Probabilism implies that whether giving you £20 promotes D —whether I have a contributory reason to give you £20—depends on how the £1 option is individuated. But why should whether giving you £20 promotes some outcome depend on whether I have just a £1 coin in my other pocket, or also a pair of 50p pieces, or those plus five 10p pieces?

The central problem is that for Flat Probabilism, the baseline depends on *how many* means there are to reaching some goal. I'll now turn to my own preferred baseline, where promotion depends on *whether* there is some lower-ranked means available.

6. The Minimal Baseline

Here is my preferred baseline:

Minimal Probabilism about Promotion. S 's A -ing promotes D if and only if there is some B in S 's ability set such that $p(D|A) > p(D|B)$.

14. I am indebted to Stephen Finlay for invaluable discussion and comments.

The intuitive picture behind the baseline is this: as in Ranking Probabilism, the actions *S* can do may be ranked by the conditional probability of *D*. Any action which does not rank *bottom* promotes *D*.

Minimal Probabilism isn't disposition-sensitive. As such, it avoids SPP's ephemerality problem in No Probability cases: recall that $p_0(D) = p_1(D) = 1$ does *not* imply that *S* is causally impotent with respect to *D*. There could be some *B* in *S*'s ability set, such that $p(D|B) = 0$, but $p_0(B) = 0$. SPP says that in such a case, *S*'s *A*-ing doesn't promote *D*. Minimal Probabilism disagrees.

Now, let's turn to Buttons. If Debbie pushes *A* or *B*, then her desire will be satisfied, but if she pushes *C*, then it will not. For simplicity, I ignore all other possible actions (including doing nothing). Lin asked whether there was any probability that Debbie would push *C*. I claim that this is the wrong disambiguating question. We should ask: *can* Debbie push button *C*? Is *C* in her ability set?

In the canonical presentation of Buttons, we are not explicitly told that Debbie can push *C*, but it is heavily implied. Ignoring other possible actions, as we are, Debbie can push button *C* *iff* she is not causally impotent with respect to her desire, because only if she pushes *C* will her desire not be satisfied. If my account is correct, then whether she can push *C* is crucial to whether her pushing *A* promotes the satisfaction of her desire. Let's consider both possibilities.

First, suppose that Debbie can push *C*. This version—call it 'Buttons with Some Ability'—seems the most natural way of understanding the case. As I mentioned, it seems to be implied in the original description of Buttons.¹⁵ She pushes *A*; had she not pushed *A* then she wouldn't have pushed *C*, but she was able to do so. She wasn't causally impotent with respect to her desire. Ranking Probabilism says that she had more reason to push *A* than to push *C*; more reason to push *B* than to push *C*; and equally strong reason to push *A* as to push *B*. Because pushing *A* and pushing *B* each rank above the bottom in her ability set, Minimal Probabilism says that both of those actions *do* promote the satisfaction of that desire. This is the intuitively correct verdict for this version of Buttons, which we have been trying to capture.

Second, let's consider the other version—Buttons with No Ability—where pushing *C* is not in her ability set. Perhaps that button is in a locked glass case. Again ignoring options besides the three buttons, there are only two things that she can do (push *A* or push *B*). *Whatever* she does, her desire will be satisfied. As I argued above, in such cases of causal impotence, she doesn't promote the satisfaction of her desire. After all, there is nothing else she could have done that would have made her desire less than certain to be satisfied. Minimal Probabilism again captures this verdict.

This is a victory for Minimal Probabilism. Other than Do Nothing, it's the

15. Joshua DiPaolo (personal communication) tells me that this is how he understands Buttons.

only account we've seen which respects the intuitive verdict in Buttons, without introducing context-dependence or disposition-sensitivity.

Working through these cases illustrates an important asymmetry that Minimal Promotion introduces. Likelihood facts about how S will presently act should not be taken into account in determining whether that act would promote some outcome, but likelihood facts about how *other* agents will act do seem importantly relevant. The causal chain from S's act to the likelihood of D may go via the acts of other agents, and so their dispositions will determine whether her act promotes some outcome, and what reasons she has. But (at least if I'm right that promotion is not disposition-sensitive), her own dispositions to act at present do not determine her reasons. (What about her future dispositions? This issue is beyond the scope of this paper, but Minimal Probabilism allows for both sensitivity and insensitivity about S's own future dispositions. I lean towards the former.)

That's the view. Now, we turn to its surprising implication.

7. What Reasons Do We Have?

The problem of the sticky baseline is that since promotion involves a probability exceeding the baseline in some way, promotion cannot occur in those cases where we stick to the baseline. According to Minimal Probabilism, the baseline is the (probabilistically) *worst* an agent can do with respect to an outcome. This is why it doesn't face a problem of the sticky baseline, except perhaps in cases of causal impotence: it's highly unlikely that we would wish to count doing the worst we can with respect to D as promoting D.

But such a low baseline might lead us to the opposite problem. Minimal Probabilism implies that we have very many reasons. Here's the example I'll rely on: I desire that D, to have \$20,000 in ten years. I now have \$10,000 in cash, and there are just five actions in my ability set, in rank order:

- A_1 : invest all the money. $p(D|A_1) = 0.7$
- A_2 : put all the money in a savings account. $p(D|A_2) = 0.3$
- A_3 : do nothing. $p(D|A_3) = 0.1$
- A_4 : give half the money to charity. $p(D|A_4) = 0.05$
- A_5 : give all the money to charity. $p(D|A_5) = 0.01$

Ranking Probabilism captures the verdictive 'ought' or 'most reason' facts: there is more reason to A_1 than to A_2 , than to A_3 , than to A_4 , than to A_5 . So far, so good.

Minimal Probabilism says that there are reasons for every action except A_5 . So a desire to have twice as much money in ten years gives me some reason to

give away half of my money now—making the probability that I double my money far lower than it would be if I simply do nothing. I'll consider two objections to this feature of Minimal Probabilism.

7.1. Cheap Reasons

The first objection is simply that the claim in question—that I have a reason to A_4 —is manifestly implausible. I don't find it so, but perhaps I've just got used to it. At any rate, it is a consequence of Minimal Probabilism that I will defend.

First, note that the disposition-sensitive baselines can also imply that I have reason to give away half of my money. In a three-action case, if I was antecedently likely to give away *all* of my money, then SPP says that giving away only half promotes my desire. Since Finlay's Counterfactual baseline is also disposition-sensitive, we can rig a similar case there, too: suppose that if I did not give away half of my money, then I would have given away all of it. But that is a negative, defensive mode of argument.

Here's a positive case for the existence of a reason to A_4 . I'll argue by dilemma. Are there only reasons for actions, or are there also reasons against actions?

The first horn is that there are only reasons for actions. Here, Minimal Probabilism seems clearly correct about the case. We can all agree that there is no reason to give *all* of one's money to charity. But it also seems clear that there is less reason to give all of the money to charity than to give half of the money to charity. Someone who did the former would be more rationally criticisable than someone who did the latter. (Remember that there is only one outcome in play—that of doubling one's money.) So there is no reason to donate all of the money, and less reason to donate all of the money than to donate half of the money. Since there can be no reasons against, we must say that there is *some* reason to donate half of the money. Otherwise, what explains the verdictive difference? If there is no reason to donate half of the money, then the verdictive difference rests on no contributory difference, which is implausible.

On the second horn of the dilemma, there are indeed reasons against actions. Discussion of this horn will take much longer.

Here is how reasons against can be accommodated by Minimal Probabilism. We again say that A-ing *dis-promotes* D iff A-ing promotes $\neg D$, and that there is a promotional reason against A-ing iff A-ing dis-promotes D. Taken together, these claims imply that there is a reason against A-ing when A promotes $\neg D$. This is so when there is some C in the agent's ability set such that $p(D|C) > p(D|A)$. The intuitive picture is that in accord with Ranking Probabilism, the actions that an agent can do may be ranked by the conditional probability of D. Minimal Probabilism says there is a reason for every action except that which ranks bottom. It also says that there is also a reason *against* every action which doesn't rank *top*. All of the intermediate actions promote both D and $\neg D$, and there are reasons

both for and against them.

Here is why this might look like a problem. If there are reasons against actions, isn't it much more plausible that there are (only) reasons against A_4 and against A_5 , and that the verdictive difference between the two actions is explained by the latter reason being stronger? Yes, this is an intuitively attractive picture. But it is also unsustainable. As we've seen, baselines which try to capture this verdict—most obviously, Do Nothing—run into the problem of the sticky baseline.

But just how implausible is the picture painted by Minimal Probabilism? Not so much as one might think! It is independently plausible—once one gets into the right frame of mind—that there is a reason to A_4 , as well as a reason against A_4 .

I'll appeal to T. M. Scanlon's uncontroversial claim that 'the concept of a reason is that of a consideration that "counts in favor of" something for an agent in certain circumstances' (Scanlon 2014: 44). This is a general conceptual claim which doesn't seem to be tied to any particular theory of reasons.

There *is* a consideration which counts in favour of A_4 : whatever makes it rank higher than A_5 . If I did A_4 , and you asked me to justify my actions, I would have something to say in my defence—that A_4 ranks higher than A_5 (at least I didn't give away *all* of my money!)—whereas had I done A_5 , I would have nothing to say in my defence. A reminder: we are concerned only with the outcome of doubling my money in ten years. I might have other things to say in my defence, but they are not relevant here.

In fact, I have nothing to say in my *defence* in either case. Nothing can get me off the hook for not doing A_1 , which is what I have most reason to do—what I ought to do. So perhaps it's better to say that if I A_4 , I have something to say in mitigation of my guilt, but nothing to say in mitigation if I A_5 . It's not a stretch to claim that anything I might say in mitigation of my guilt is a consideration that counts in favour of my action, and so is a reason for my action. (More precisely: anything I might say about the actions in question could be such a reason; appeals to my mental state and so on can be excuses or mitigations but might not be reasons.)

But, of course, claims such as 'a reason is a consideration which counts in favour' are often appealed to by those who claim that there can be no informative account of what it is to be a reason. For example, Scanlon defends Reasons Fundamentalism:

truths about reasons are not reducible to or identifiable with non-normative truths, such as truths about the natural world of physical objects, causes and effects, nor can they be explained in terms of notions of rationality or rational agency that are not themselves claims about reasons. (2014: 2)

But what if reasons are not fundamental? I'll also argue that two prominent and putatively informative theories of reasons—reasons as evidence, and its more

popular cousin, reasons as explanations—can accommodate a reason to A_4 .

According to the former view, a reason for S to A is evidence that S ought to A (see, e.g., Kearns & Star 2008). And the fact that A_4 ranks higher than A_5 (or whatever grounds that fact) is indeed evidence that S ought to A_4 . First, it excludes the possibility that A_4 is the worst option available to S. That this evidence is outweighed by reasons against A_4 —evidence that one ought not to A_4 —doesn't prevent that evidence being an (outweighed) reason. This is simply one way for the reasons-as-evidence view to handle outweighed reasons.

According to the reasons-as-explanations view, however, to be a reason to A_4 is to be (part of) what explains something. For example, here is Schroeder's Humean account, with some notational modifications:

Schroeder-Reason. For R to be a reason for S to A is for there to be some D such that S has a desire whose object is D, and the truth of R is part of what explains why S's doing A promotes D. (Cf. Schroeder 2007: 59)

In other words, if I have a reason to A_4 , this is a (perhaps partial) explanation of why doing A_4 would promote the outcome that I double my money. And there is such an explanation: if I A_4 , the chance that I achieve my outcome is much greater than if I A_5 , which I could also do. (Or perhaps the reason is the explanation of *this* fact. The point is that there is such an explanation in the area.) Since Minimal Probabilism says that A_4 promotes the outcome that I double my money, adding that there is an explanation for why A_4 promotes that outcome doesn't introduce any *additional* implausibility.

Clearly there is much more to be said about the nature of reasons, but I've argued that the apparently excessive cheapness of Minimal Probabilism—as exemplified by the claim that I have a reason to A_4 —fits naturally with a general conceptual claim about reasons, and with two popular accounts of reasons.

But if there are such reasons, why would it be so absurd to *say* that my desire to double my money supports giving away half of my money? Predictably, my answer appeals to pragmatic considerations. Such reasons are not normally worth mentioning and are potentially misleading. We can be fooled into thinking that they don't exist. In the present case, I ought not do anything except A_1 —invest all the money—and mentioning any reasons to do anything else could be seriously misleading, or at least pointless.

But they can be made worth mentioning. The reason to A_4 can be brought out by varying the case, and in particular considering a case of multiple competing reasons. Suppose that besides my wealth target, I have another outcome to promote. I'll put it in Humean terms. I have no desire to donate more than a token amount to my college which is—rightly or wrongly—a charity, but I do have a very strong desire to donate *something*. The college has a target that some proportion of its graduates donate to the annual fund, and I have a very strong

desire to help them meet this target (perhaps I simply desire that my name get on the list of those who have donated).

To be clear, this is a changed choice situation in one way—an additional desire and therefore reason is added—but the choice situation remains unchanged in that my ability set is unchanged, and the causal upshots of the members of that set are also unchanged. An objective is simply added.

Call this charitable outcome 'E'. Then $p(E|A_1) = p(E|A_2) = p(E|A_3) = 0$, and $p(E|A_4) = p(E|A_5) = 1$. So Minimal Probabilism says that I have a reason (grounded in E) to A_4 , and an equally strong one to A_5 , but no such reason to do any of the other actions.

I now have to act with two outcomes supplying reasons, so plausibly what I ought to do all-things-considered depends on the relative importance of those outcomes. Such a theory of weighing competing reasons is beyond the scope of this paper, so here I'll simply stipulate a verdictive outcome: the outcome E of donating to my college is so important, that when both outcomes are taken into account, A_1 , A_2 , and A_3 have far less rational support than A_4 and A_5 do.

So now, ought I donate half (A_4) or all (A_5) of my money to the college? Since both of these two actions involve a donation, the charitable reason to donate *something* is silent between them. But in terms of my wealth desire, the former does much better. Clearly, I ought to A_4 . A_5 would involve expected financial loss for no charitable gain. This brings out the consideration which counts in favour of A_4 : doing so keeps the probability of doubling my money higher than A_5 does, even though A_4 goes out of the way to reduce that probability.

Minimal Probabilism implies that such considerations are pragmatically swamped in ordinary talk of reasons—they are rarely mentioned because negligible or misleading—but the weighing case makes them salient. Clearly, there are things that my opponent can say in response. But I hope I've shown that it's not so implausible as one might think that there's a reason to A_4 .

7.2. *Restricted Exclusivity*

There's a more technical version of the cheapness worry, which rests on the fact that there is a reason both for and against each intermediate action A_2 , A_3 , and A_4 . According to Snedegar, such both-ways reasons are impossible.

First, a little background. Snedegar's Contrastive view of reasons—which rests on a non-contrastive promotion relation—says that given a (possibly non-exhaustive) set of actions Q, and where O is the objective to be promoted, the following principle holds:

Snedegar-Against. r is a reason against A-ing out of Q iff there's some O of the relevant kind such that r explains why B-ing better promotes O than A-ing, for some other alternative B in Q. (Snedegar 2017: 77)

Snedegar is here assuming a version of reasons-as-explanations. But Contrastivism doesn't rely on this. Snedegar-Against agrees with Minimal Probabilism about how cheap reasons against are: every action which doesn't *best* promote O has a reason against it. If cheap reasons against are a cost of Minimal Probabilism, then Contrastivism shares this cost: Snedegar's view faces its own version of the cheapness or 'liberality' worry when it comes to reasons against actions (Snedegar 2017: 80).

But while Minimal Probabilism also has cheap reasons for actions, Contrastivism makes reasons for actions very expensive:

Snedegar-For. r is a reason for you to A out of Q iff there's some O of the relevant kind such that r explains why A-ing better promotes O than B, for all the other alternatives B in Q. (Snedegar 2017: 78)

So where Minimal Probabilism says that there is a reason for every action except that which ranks bottom, Snedegar-For says that there are only reasons for those actions which rank top. Reasons against actions may be very cheap, but reasons for actions are very expensive.

Let's suppose that Q is everything Rhys can do. It seems implausible that unless A-ing is the *best* Rhys can do, then he has no reason to A. But Snedegar-For implies that Rhys has no reason to buy nearly all, but not all, of the tickets. Such a claim seems to me to confuse verdictive claims (about what one ought to do) with contributory claims (about what one has some reason to do).

Contrastivism also introduces an odd asymmetry between reasons for and against. Where Minimal Probabilism faces a cheapness worry in both directions—there is a reason both for and against buying 499 tickets—Contrastivism faces a cheapness worry in one direction (there is a reason against buying 499 tickets), but also its opposite in the other direction: there is no reason to buy 499 tickets. Minimal Probabilism has to explain cheapness; Contrastivism has to explain cheapness (of reasons against) and also its opposite (of reasons for).

It might be wondered whether Minimal Probabilism about Promotion can be reconciled with Snedegar's view.¹⁶ At heart, the views are irreconcilable. They disagree about the role of promotion in a theory of reasons: my Absolute Reason-Promote says that there is a reason to A *iff* A-ing promotes a relevant outcome, but Snedegar-For denies this: promotion is not sufficient for there to be a reason. We could force the views together, perhaps, but only by bleaching out their core theoretical commitments.

So by my lights, Snedegar-For is implausible. But of course Snedegar has good reasons for his view. He considers a view similar to Minimal Probabilism, that an outcome 'O can give you a reason to A relative to some set when your

16. I'm grateful to an anonymous reviewer for raising this question.

A-ing would better promote O than *some* other alternative in the set' (Snedegar 2017: 78). He rejects it because it leads to both-ways reasons—for example, both for and against A_4 —by appeal to the following principle:

Restricted Exclusivity. For all facts r , agents s , actions A , and objectives o , o cannot explain both why r is a reason for s to A and why r is a reason for s not to A . (Snedegar 2017: 31)¹⁷

Restricted Exclusivity says that when we are just considering one objective, a fact can provide a reason for some action, or a reason against it, but not both for and against it. We should admit that this principle is extremely intuitively plausible. Nevertheless, I'll argue that it doesn't ultimately undermine Minimal Probabilistic Promotion.

If there are no background conditions—no facts which enable a fact to be a reason, without themselves being part of that reason—then Minimal Probabilism doesn't violate Restricted Exclusivity. Though the same objective or outcome o explains both a reason for me to A_3 and a reason for me not to A_3 , the ' r ' in question varies. There is not one fact that is a reason with both valences.

The complete reason for A_3 is at least that $p(D|A_3) = 0.1$ but $p(D|A_5) = 0.01$; the complete reason against A_3 is at least that $p(D|A_3) = 0.1$ but $p(D|A_1) = 0.7$. So both the reasons for and against A_3 include the fact that $p(D|A_3) = 0.1$, but in each reason that fact is conjoined with another fact. Alone, $p(D|A_3) = 0.1$ is not a reason both for and against A_3 with respect to my objective or outcome, and so doesn't violate Restricted Exclusivity.

More trouble lurks if we suppose that there can be background conditions.¹⁸ It's of course clear that the same fact r can be a reason both for and against A relative to *different* sets of background conditions. But the trouble for Minimal Probabilism lies in a more restricted version of the Restricted Exclusivity:

Doubly Restricted Exclusivity. For all facts r , agents s , actions A , objectives o , and sets of background conditions Z , o cannot explain both why r is a reason for s to A with background Z , and why r is a reason for s not to A with background Z .

Here's why Minimal Probabilism violates Doubly Restricted Exclusivity. If the facts that $p(D|A_1) = 0.7$ and $p(D|A_5) = 0.01$ are in the set of background conditions, then the view says that $p(D|A_3) = 0.1$ is a reason both for and against A_3 relative to that set.

17. Since I'm here concerned with promotional reasons, the objective is always that some outcome obtains, so I'll use the terms interchangeably where it won't introduce confusion.

18. I am grateful to Emma Borg and Justin Snedegar for pressing this objection, and for extensive discussion.

But I'm going to argue that Doubly Restricted Exclusivity is false. Once we admit background conditions, the case for the principle looks very thin. To avoid question-begging, I'll illustrate this with a case that doesn't involve (probabilistic) promotion. Suppose that your only objective is finding strong coffee; the stronger, the better. The set of background conditions includes the following relevant facts: {there is very weak coffee from a gas station in the bathroom; there is very strong espresso in the kitchen}.

Now, consider the fact that there is fairly strong filter coffee in the study. Could that be—in violation of Doubly Restricted Exclusivity—a reason both for and against going to the study? I don't see why not. Suppose that you asked me for advice about where to go: is going to the study a good idea? I could cite the presence of fairly strong filter coffee in the study as a reason to go to the study (relative to the background condition that there is weak coffee in the bathroom), or I could cite the presence of fairly strong filter coffee in the study as a reason not to go to the study (relative to the background condition that there is strong espresso in the kitchen).

Of course, there is some awkwardness here. Because there is only one objective—getting the strongest coffee possible—it would be strange and irrational for you to do anything other than go to the kitchen for espresso. You ought to go to the kitchen. Nobody should deny this verdictive fact. This feature of the cases will be hard to escape without a move to multiple competing objectives, as in the charity case.

But here I am concerned with the contributory reasons. The same fact—that there is fairly strong filter coffee in the study—provides a reason for and against going to the study, but this reason is 'enabled' by different members of the set of background conditions. That there is fairly strong filter coffee in the study is a reason to go to the study given that there is weak gas station coffee in the bathroom; that there is fairly strong filter coffee in the study is a reason against going to the study given that there is strong espresso in the kitchen.

This, I think, points to the truth in Doubly Restricted Exclusivity. Define the set of *minimal background conditions* for a reason r (with a certain valence) as the smallest possible subset of Z relative to which r remains a reason with that valence.

In both violations of Doubly Restricted Exclusivity, r is a reason with both valences, but with a different set of minimal background conditions for each valence. As a reason for A_3 , that $p(D|A_3) = 0.1$ has facts about either A_4 or A_5 in its minimal background conditions, together with uncontroversial background facts. (r has two sets of minimal background conditions as a reason to A_3 : one which contains A_5 , and one which contains A_4 .) As a reason against A_3 , that $p(D|A_3) = 0.1$ has facts about either A_1 or A_2 in its minimal background conditions, together with uncontroversial background facts. (r has two sets of

minimal background conditions as a reason against A_3 : one which contains A_1 , and one which contains A_2 .)

The truth in Doubly Restricted Exclusivity is that r cannot be a reason both for and against A_3 relative to the same set of minimal background conditions. I don't have space to argue for this weaker version of Restricted Exclusivity here, but can't think of any counter-examples. In any case, arguments for various versions of Restricted Exclusivity are hard to come by. At the risk of spawning too many named principles:

Triply Restricted Exclusivity. For all facts r , agents s , actions A , objectives o , and sets of minimal background conditions Z , o cannot explain both why r is a reason for s to A with background Z , and why r is a reason for s not to A with background Z .

A rich set of background conditions may include subsets which enable r to be both a reason against A , and a reason for A . But a set of *minimal background conditions* will not. At least, I think that this is plausible—and doesn't threaten Minimal Probabilism.

To sum up. If there are no background conditions for reasons, then Restricted Exclusivity is plausible but Minimal Probabilism doesn't violate it. If there are background conditions, then Minimal Probabilism violates Doubly Restricted Exclusivity, but that principle is false. Minimal Probabilism doesn't violate Triply Restricted Exclusivity, which may be plausible, though I've not argued for it.

This section has been largely defensive, so I will add a concessive note: if one accepts Doubly Restricted Exclusivity, and thinks that there are both reasons against actions and background conditions for reasons, then this would be—in the vein of Snedegar—a good reason to reject Minimal Probabilism. The price of the theory is too high for you. But the other views, including Contrastivism, have their own high prices.

8. Conclusion

I have argued that Minimal Probabilism deals correctly with perennial problem cases, most importantly Buttons and the three-action cases, and that it thereby undermines much of the motivation for contrastivism. Both the lack of promotion in cases of causal impotence and the very cheap reasons could reasonably be cited as reasons to reject the Minimal Probabilism. But I have argued that these are features that a probabilist should live with, or even embrace.

So where are we? There are four candidate views remaining, if one accepts that disposition-sensitivity and the fit account are just too implausible. These are Do Nothing, Ranking Probabilism (without a baseline), Contrastivism, and Minimal Probabilism.

You might be surprised to see Do Nothing on this list. But we now have the tools to make a more spirited defence of that baseline. The distinction between the contributory and the verdictive, supplemented with dis-promotion and associated reasons against actions, allows us to say that though there is never a reason to do nothing, there can often be *most* reason to do nothing. This is a contributory oddity, but perhaps an acceptable one—especially as a regimentation of our everyday talk of promotion. A claim that one promotes an outcome by doing nothing does stick in the throat a little; it's much more natural to say that one simply fails to dis-promote the outcome, as Do Nothing has it.

This addition to Do Nothing highlights an ambiguity or underspecification in the slogan that a reason is a consideration which counts in favour of some action: counts in favour *as opposed to what?* Finlay's counterfactual baseline says: as opposed to not doing it. Do Nothing says: as opposed to doing nothing. I have defended Minimal Probabilism: as opposed to something else you could do.

However, I think that Do Nothing is ultimately indefensible as a normative theory of promotion, as implicitly defined by Absolute Reason-Promote and Degree Reason-Promote. Minimal Probabilism has some quirks, to be sure, but these are systematic and not quite so damaging as they may first appear. To return to the financial case, it's overwhelmingly intuitive that there is a reason against A_4 . Minimal Probabilism simply adds that there is also a reason for A_4 . Any residual implausibility should be set against the fact that Minimal Probabilism correctly deals with Buttons, and avoids disposition-sensitivity, contrastivism, and three-action problems.

So ends my defence of Minimal Probabilism. We've seen that the relatively tricky and apparently niche issue of promotion has substantive implications for the theory of reasons. For example, one longstanding debate about reasons concerns their cheapness. It's a worry for Humeanism that it produces 'too many reasons', as Schroeder (2007: Ch. 5) puts it.¹⁹ Minimal Probabilism implies that promotional reasons are very cheap indeed, for the Humean and for everyone else. There are many more practical reasons than most of us suspected (as the financial example shows), and this mitigates cheapness as a worry for Humeanism. We can no longer say that by abandoning some versions of Humeanism we remove the problem of cheap reasons, unless we also adopt a different account of promotion.

One might think the promotion debate trivial because of how *recent* it all is. Isn't it just a relatively niche philosophical fashion? But there are good philosophical reasons for promotion's recent rise to prominence. Famously, as Scanlon (2014: 1) notes, 'reason'-talk has only been common (in meta-ethics at least) in the last few decades, replacing much talk of duties, obligations, and rightness.

19. As Evers (2009) makes clear, the cheapness of reasons is crucial to Schroeder's project of vindicating agent-neutral reasons.

Whereas obligations and duties concern which actions are required or optimal in a given situation, reasons—especially promotional reasons—naturally lend themselves to more scalar discussions, and to a concern for the degree of normative support that sub-optimal actions have. We talk not of actions that are obligatory and non-obligatory, but of those which are supported by stronger and weaker reasons. So the reasons programme in ethics naturally leads to a focus on the question of promotion. Granted, the promotion debate tends to use toy cases such as Buttons. But if the reasons programme fails in simple cases involving a single outcome, then that's not promising. I've argued that it doesn't fail. The best theory of promotion has some counterintuitive implications—but ones we can live with.

Acknowledgments

For discussion and comments, I am grateful to Emma Borg, Finnur Dellsén, Joshua DiPaolo, Daan Evers, Brad Hooker, Thomas Schmidt, Nate Sharadin, Justin Snedegar, several anonymous reviewers, and to audiences in Lund and Reading, including the Spring 2019 Reading Graduate Class. I'm indebted to Stephen Finlay for catching a serious error about his view—any remaining mistakes are, of course, my fault. Writing of this paper was funded by the European Union (H2020-MSCA-IF-2016 grant 748617, 'Austere Reasons') and the University of Reading, via a REF 2020 Fellowship. I'm extremely grateful to both institutions.

References

- Behrends, Jeff and Joshua DiPaolo (2011). Finlay and Schroeder on Promoting a Desire. *Journal of Ethics and Social Philosophy*, 6(1), 1–7. <https://doi.org/10.26556/jesp.v6i1.146>
- Behrends, Jeff and Joshua DiPaolo (2016). Probabilistic Promotion Revisited. *Philosophical Studies*, 173(7), 1734–1754. <https://doi.org/10.1007/s11098-015-0576-0>
- Coates, D. Justin (2013). An Actual-Sequence Theory of Promotion. *Journal of Ethics & Social Philosophy*, 7(3), 1–8. <https://doi.org/10.26556/jesp.v7i3.157>
- Dancy, Jonathan (2004). *Ethics Without Principles*. Oxford University Press. <https://doi.org/10.1093/0199270023.001.0001>
- DiPaolo, Joshua and Jeff Behrends (2015). Reason to Promotion Inferences. *Journal of Ethics & Social Philosophy*, 9(2), 1–9. <https://doi.org/10.26556/jesp.v9i2.176>
- Evers, Daan (2009). Humean Agent-Neutral Reasons? *Philosophical Explorations*, 12(1), 55–67. <https://doi.org/10.1080/13869790802635614>
- Finlay, Stephen (2006). The Reasons That Matter. *Australasian Journal of Philosophy*,

- 84(1), 1–20. <https://doi.org/10.1080/00048400600571661>
- Finlay, Stephen (2014). *Confusion of Tongues*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199347490.001.0001>
- Foot, Philippa (1972). Morality as a System of Hypothetical Imperatives. *The Philosophical Review*, 81(3), 305–316. <https://doi.org/10.2307/2184328>
- Kearns, Stephen and Daniel Star (2008). Reasons: Explanations or Evidence? *Ethics*, 119(1), 31–56. <https://doi.org/10.1086/592587>
- Lin, Eden (2018). Simple Probabilistic Promotion. *Philosophy and Phenomenological Research*, 96(2), 360–379. <https://doi.org/10.1111/phpr.12310>
- Manne, Kate (2016). Democratizing Humeanism. In Errol Lord and Barry Maguire (Eds.), *Weighing Reasons* (Vol. 6, 123–140). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199315192.003.0006>
- Scanlon, T. M. (2014). *Being Realistic about Reasons*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199678488.001.0001>
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199299508.001.0001>
- Sharadin, Nathaniel (2015). Problems for Pure Probabilism about Promotion (and a Disjunctive Alternative). *Philosophical Studies*, 172(5), 1371–1386. <https://doi.org/10.1007/s11098-014-0354-4>
- Sharadin, Nathaniel (2016). Checking the Neighborhood. *Journal of Ethics and Social Philosophy*, 10(1), 1–8. <https://doi.org/10.26556/jesp.v10i1.181>
- Sharadin, Nathaniel and Finnur Dellsén (2019). Promotion as Contrastive Increase in Expected Fit. *Philosophical Studies*, 176(5), 1263–1290. <https://doi.org/10.1007/s11098-018-1062-2>
- Snedegar, Justin (2014). Contrastive Reasons and Promotion. *Ethics*, 125(1), 39–63. <https://doi.org/10.1086/677025>
- Snedegar, Justin (2017). *Contrastive Reasons*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198785934.001.0001>
- Sobel, David (2016a). Subjectivism and Proportionalism. In *From Valuing to Value* (298–308). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198712640.003.0016>
- Sobel, David (2016b). Subjectivism and Reasons to Be Moral. In *From Valuing to Value* (16–42). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198712640.003.0002>