

# Hope and Despair at the Kantian Chicken Factory

## Moral Arguments about Making a Difference

*Andrew Chignell*

### 11.1 Kant and Kantians on animals

Kant famously—notoriously—claims that we have no *direct* duties to non-human animals. Like plants, animals<sup>1</sup> are ‘*things*’ with only ‘relative worth’, he says, and so it is permissible to treat them (their bodies, secretions, and eggs) as mere means to our ends of feeding, clothing, and entertaining ourselves. The impermissible actions regarding animals are those that disrespect their owners or desensitize us to the suffering of the sentient creatures who *do* morally count—namely, the rational ones. In short: for Kant, our direct duties are to persons; any duties to other animals are *indirect* and go via our duties to persons.<sup>2</sup>

Many contemporary philosophers—including many Kantians—regard Kant’s arguments about animals (like his arguments about gender,<sup>3</sup> race,<sup>4</sup> masturbation,<sup>5</sup> homosexuality,<sup>6</sup> and political revolution<sup>7</sup>) as tied to his cultural context in a way that makes them (in letter if not in spirit) obsolete and sometimes repugnant.<sup>8</sup> A few contemporary philosophers, most prominently Tom Regan and Christine Korsgaard, have developed alternative approaches that are *broadly Kantian* in flavour and yet come to quite different conclusions about our obligations to animals. They are *broadly* Kantian because they cite the kinds of considerations that Kant himself uses in establishing the dignity of human beings. They are not narrowly Kantian because they conclude that we do have direct duties to animals, and that we may not treat them merely as means to our ends.

Regan (1983) does this by arguing that animals are, like us, sentient ‘subjects of a life’ (this is a technical term for him) and that this is evidence that they have an

<sup>1</sup> In what follows I will use ‘animals’ to refer to non-human animals.

<sup>2</sup> G 4:427–8; *Con.* 8:110ff; MM 6:241, 443–4; EC 27:459, 710. <sup>3</sup> See Varden (2017).

<sup>4</sup> See Allais (2016). <sup>5</sup> See Kielkopf (1997). <sup>6</sup> See Altman (2010).

<sup>7</sup> See Flikschuh (2018).

<sup>8</sup> But for a positive assessment of Kant’s views about animals, see Wilson (2008).

irreducible ‘inherent value’ that confers moral standing and welfare rights. Regan’s view has been influential (often playing the role of the ‘Kantian’ or ‘deontological’ approach in animal ethics anthologies), but critics have wondered how to understand the metaphysics and epistemology of the obscure conception of value involved.

Korsgaard (2018) tries to avoid this concern by developing an explicitly ‘naturalistic’ Kantian-constructivist approach. She starts with the Aristotelian point that animals have ‘tethered goods’—goods *for* them, just as we do. One of the main goods for any sentient animal is to live in safe, comfortable environs; another is to have a *full* life—one that is not radically curtailed, even painlessly. Obviously we humans take, and require others to take, what is good *for us* as a defeasible basis for making moral claims on ourselves and others. Korsgaard regards such taking and requiring as conferring ‘absolute’ value on our tethered goods, and as ‘claiming standing’ for ourselves as ends rather than mere means (2018: 139). But there is no non-arbitrary basis, she says, for refusing to take the tethered goods of *other* sentient creatures into account in this way, even if they are not able to do so themselves. Thus

when we consider the reason why Kant thinks we must claim the standing of ‘end-in-itself’ for ourselves, we will see that we must claim that standing for the other animals as well. (2018: 130)

I have examined Korsgaard’s innovative argument in detail elsewhere (Chignell, forthcoming 2020), but one major problem involves the following premise (this is my paraphrase rather than a quotation):

*Parity:* If there is no difference between us and the other animals with respect to either the having of tethered goods or the importance of those goods, and if we treat our tethered goods as goods absolutely (thereby making them good absolutely), then we ought to treat other animals’ tethered goods as good absolutely (thereby making them good absolutely).

Korsgaard argues over the course of several chapters that the two conjuncts of the antecedent are satisfied, and then spends the central chapter of the book arguing for the conditional as a whole. But an opponent could plausibly suggest that the reason we take our tethered goods to be part of the absolute good is *not* because they are tethered goods *simpliciter*, but because they are *our* tethered goods. In other words, the opponent could plausibly suggest that we treat human goods (of a full life, health, safety, capacities, opportunity, etc.) as part of the absolute good because, well, they are goods *for us*. Interestingly, Korsgaard herself seems to articulate the premise this way in places: ‘I have no other reason for taking my end to be good absolutely, than the fact that it is good *for me*’ (2018: 144).

Instead of *Parity*, then, it looks like we have

*Parity\**: If there is no difference between us and the other animals with respect to either the having of tethered goods or the importance of those goods, and if we treat our tethered goods as goods absolutely *because they are our tethered goods* (thereby making them good absolutely), then we ought to treat other animals' tethered goods as good absolutely (thereby making them good absolutely).

Parity arguments are tricky, but now this looks invalid: the italicized phrase in *Parity\** adds something to the antecedent that makes the conditional as whole implausible. Genuine parity would require a consequent like '...then the other animals ought to treat their tethered goods as goods absolutely *because they are their tethered goods* (thereby making them good absolutely)' But that, unfortunately, is not something that our fellow creatures can do.

In what follows, my goal is not to rehearse Kant's narrow views about our duties to animals; nor is it to examine further these broadly Kantian efforts to improve on Kant's ideas.<sup>9</sup> Instead, I propose to draw on another part of Kant's philosophy altogether—namely, his moral psychology of hope and despair—to develop a different line of thought about our relationship to animal *products*.

For these purposes I will simply start by assuming that there are sound arguments (of one of these broadly Kantian varieties, perhaps, or of a more consequentialist variety) for the claim that:

*Don't Farm*: It is morally wrong to raise, kill, and harvest animals for the purpose of selling their bodies, secretions, and eggs if there are readily available alternatives for us to eat and wear.

Now consider an average North American consumer of middling means—call him Oppy—who likes meat and cheese and omelettes and leather. He would purchase these products if he could, but he believes that *Don't Farm* is true, and he assumes, like many of us, that *Don't Farm* implies:

*Don't Purchase*: It is morally wrong to purchase the bodies, secretions, or eggs of farmed animals if there are readily available alternatives for us to eat and wear.

Recently, however, Oppy has confronted some sobering empirical facts that seem to problematize this natural move from *Don't Farm* to *Don't Purchase*. The sobering empirical facts, roughly put, are that the sheer size and complexity of the

<sup>9</sup> For more of the latter, see Varden (this volume).

industrial system that brings 99% of the animal products to our refrigerators and closets—and, in particular, the oversupply and buffers in many of its supply chains—make it *exceedingly* unlikely that an individual's choice not to purchase will have any effect at all on the system, whether positive or negative. Despite his unwavering commitment to *Don't Farm*, then, Oppy has become an opportunist: he thinks it is okay on occasion to privately enjoy and wear some of his favourite farm animal products, since (as long as he is purchasing from the industrial system) doing so almost certainly doesn't make a difference.<sup>10</sup>

Michael Martin (1976) and R.G. Frey (1985)—and much more recently Mark Bryant Budolfson (2016, 2019) and Eliot Michaelson (2016a)—have argued that difference-making concerns like these are pressing for consequentialist advocates of *Don't Purchase*. This is because act consequentialism ('consequentialism' hereafter) endorses the following principle:

*Causal Consequences Principle (CCP):* Action A is morally wrong only if the causal consequences of A are worse than they would have been if some readily available alternative to A had been performed.<sup>11</sup>

'Worse' here is traditionally construed in a welfarist way that includes all sentient animals.<sup>12</sup> So given CCP, if a private<sup>13</sup> purchase of a chicken sandwich that is wrongfully produced almost certainly *doesn't* cause any additional loss in welfare, and clearly *does* lead to the goods of pleasure and nourishment, then it cannot be morally wrong. (It might even be morally required!) This sets up a trilemma for the consequentialist: give up CCP, give up *welfarism*, or give up *Don't Purchase* and get comfortable with opportunism. The whole problematic is now referred to under the label of the *inefficacy* (or *futility* or *causal impotence*) objection.

Michaelson (2016b) and Budolfson (unpublished) have argued that the inefficacy objection also makes things difficult for some non-consequentialist efforts to move from *Don't Farm* to *Don't Purchase*. For instance, it makes things difficult for views on which *only* acts that involve or lead to the infringement or violation of rights count as morally wrong. That's because the sobering empirical facts indicate that a choice to purchase the products of a gigantic, insensitive system is

<sup>10</sup> For the term 'opportunistic carnivore', see Almeida/Bernstein (2000). For more discussion, see (Chignell 2016).

<sup>11</sup> Compare Almeida/Bernstein (2000: 206) who state this as a biconditional. I use a necessary condition here, since some consequentialisms will have different ways of filling out the account of moral wrongness. See Sinnott-Armstrong (2015).

<sup>12</sup> Some welfarist views restrict the morally relevant outcomes to human animals. Joseph Raz calls this 'the humanistic principle': 'the explanation and justification of the goodness or badness of anything derives ultimately from its contribution, actual or possible, to human life and its quality' (1986: 194). Here I am working with a welfarist picture that is not speciesist in this way.

<sup>13</sup> The privacy condition is meant to deflect concerns about the expressive significance of one's purchasing behaviour and the influence it might have on the behaviour of others.

extremely unlikely to infringe or violate any rights. And most non-consequentialists of this variety will allow, in other contexts, that it is permissible to perform actions that are extremely unlikely to infringe or violate rights. Driving to the store, for instance, may end up in an accident that puts other people in a hospital; provided this result is extremely unlikely, however, it is still permissible to drive to the store.

Other varieties of non-consequentialism (certain forms of virtue theory, symbolic or expressive value theory, divine command theory, and even Parfit's neo-Kantianism<sup>14</sup>) are not as obviously vulnerable to the inefficacy objection, however. The decision privately to purchase the products of a morally wrongful system may be vicious, symbolically unacceptable, expressively noxious, divinely proscribed, or categorically irrational even if it has no other negative outcomes.<sup>15</sup>

In what follows, I will look a bit further at the sobering empirical facts and how they problematize the move from *Don't Farm* to *Don't Purchase* (Section 11.2). In Section 11.3, though, I want to turn our attention away from Oppy in order to focus on another character: Hope. As we will see, Hope is someone who likes animal products just as much as Oppy does but takes herself to have good reasons for thinking that *Don't Purchase* is true. These reasons might be grounded in one of the non-consequentialist theories just mentioned, and in any case do *not* hang on whether her action makes a causal difference to welfare outcomes. All the same, Hope often—and quite understandably—finds her efforts to adhere to *Don't Purchase* threatened by her awareness of the sobering empirical facts. And she's not alone in this: for all but the most impressively stoical among us, the awareness that our boycotting efforts almost certainly make no causal difference can pose a serious *psychological* threat to our ongoing moral resolve, even if our reasons for boycotting do not *depend* on the boycott making a causal difference.

Here is another way to put the point. The usual version of the inefficacy problem is *moral-conceptual*: it starts with the intuitive idea that purchasing a certain product is wrong, and then shows that various ethical theories imply the opposite in circumstances where individual choices are almost certainly inefficacious. The version of the problem that I develop in Section 11.3, by contrast, is *moral-psychological*: it, too, starts with the conviction that purchasing is wrong, and then points out that people who both recognize this and are still inclined to do it will often be psychologically demoralized under conditions of perceived inefficacy such that they are unable to resist.

<sup>14</sup> See Parfit's reformulations of the Categorical Imperative in response to the 'problem of imperceptible differences' in his 2011 (341).

<sup>15</sup> I call these non-consequentialist theories, but the relevant factors can obviously be 'consequentialized' in a non-welfarist axiology such that the increases in exercises of virtue, symbolic value, obedience to God, practical consistency, etc. *become* the consequences with respect to which we ought to seek the optimistic result. Thanks to Daniel Rubio and Ryan Darr for conversation here. See Dreier (1993).

Interestingly, Kant has something much more useful to say about this second set of concerns than he does about our direct obligations to animals. Or at least that's what I will suggest below. In Section 11.4, I reconstruct one version of his famous 'moral proof' of God's existence and the future life of the soul, one that he developed in the 1790's towards the end of his career (just as his hopes regarding the French Revolution were giving way to despair). I call it the 'moral-psychological proof' in order to highlight the fact that—unlike the better-known versions of the moral proof—it relies on empirical premises about our tendencies to hope and despair, and the concomitant effects on our psychological resolve. Like the earlier and more famous versions of the moral proof, the conclusion of the moral-psychological version is not that God and the future life exist, but rather that we are morally justified in having faith (the German term is '*Glaube*'<sup>16</sup>) that God and the future life exist. After making this detour through Kant's theistic moral psychology, my ultimate goal (in Section 11.5) is to see whether a secular analogue of the proof applies to people in contemporary industrial contexts who, like Hope, have trouble sustaining their resolve in the face of the sobering empirical facts.<sup>17</sup>

## 11.2 Inefficacy at KCF

Inefficacy concerns arise in many industrial contexts (leather, coffee, fabrics, energy, etc.), but I propose to focus on the chicken system in a world very much like our own: one in which 66 *billion* chickens are processed each year (this means that around 10,000 birds will have been harvested while you were reading this very sentence).<sup>18</sup> Let's also suppose that key parts of the massive supply systems in this world are 'lumpy' rather than linear. In other words, in these parts of the chain, supply responds to demand in large lots or 'lumps' rather than in one-by-one adjustments.<sup>19</sup> This means that there will be demand thresholds—upper and lower—that must be crossed in order to trigger *any* change in supply at all. It also

<sup>16</sup> Kant uses this term in a variety of contexts, and there is no good English translation of it. In other work I have used the Germanic-looking term 'Belief' (see 2007b); here, however, I'll use 'faith', since that is standard terminology in discussions of Kant's moral proof of the rationality of theism.

<sup>17</sup> I have tried to set up Oppy's situation more precisely here than I did in my (2016) and (2018), but some of the argumentation is the same. I am grateful to the editors of both volumes for letting me incorporate revisions of that material in what follows.

<sup>18</sup> That was the statistic in our world for 2017, according to the Food and Animal Organization of the United Nations. Data can be found here: <http://www.fao.org/faostat/en/#data/QL>, selected for World + (Total), Producing Animals/Slaughtered, Meat, Chicken > List, 2017. The number has certainly gone up since then.

<sup>19</sup> See Michaelson (2016a) for a discussion of the extent to which the inefficacy objection hangs on contingent facts about the actual food supply. For doubts, partly from an economist's point of view, about how much insensitivity there is in the actual world's system, see Halteman and McMullen (2018). But even if systems in the actual world are *not* insensitive in this way, they certainly *could* be. More importantly for present purposes, they could reasonably be *perceived* to be that way by would-be moral agents.

means that the system will be insensitive to slight changes in demand, and thus that an individual's occasional choice privately to purchase a small poultry product is *very* unlikely to have an effect on the conditions of the animals and workers in the system, or the environment and wildlife affected by it.

Now let's consider a massive global fast food company in this nearby world—one that is run by people who share the narrow Kantian view about our direct duties to animals (*viz.* that we have none). This company sources chicken from vast industrial operations that carefully comply with human health and worker regulations (like good Kantians) but uses intensive factory techniques on the birds (dark crowded indoor sheds, beak-clippers, battery cages, automated throat-cutters, scalding tanks, etc.) in order to maximize yield. We can call this company the *Kantian Chicken Factory*: 'KCF' for short.

Let's suppose that in this world (which, again, is very much like our own) the regional KCF has a policy of ordering 1,000,000 chickens every month from its supplier in order to meet an average monthly demand just shy of 20 million chicken sandwiches. (I am estimating that a fast-food chicken patty contains around 1/20<sup>th</sup> of the edible parts of a chicken.<sup>20</sup>) More specifically, the policy says that *if* the demand in any given month is between 19,900,000 and 20 million sandwiches, then KCF won't change its usual order: their supply chain managers are prepared to waste (or donate, when they can) that much oversupply in a month if they think the average demand will continue to be around 20 million sandwiches.<sup>21</sup> That's because, given how cheap the government-subsidized chicken fed on government-subsidized corn is in this nearby world, and how expensive it would be to introduce mechanisms that would allow them to track demand more closely, and the unpredictability of product recalls, the KCF executives figure it is worse to face a shortage in a given month than it is to throw out some extra supply. Moreover, average monthly demand in this region has been within these two thresholds every month for the past ten years. Still, the policy is not entirely insensitive to market changes. If the number of sandwich orders in a given month falls *below* the 19,900,000 threshold, this will trigger a 'lump' reduction in their order for the following month: they will order 995,000 chickens

<sup>20</sup> The average broiler chicken weighs around five pounds. The average chicken sandwich is a quarter-pound or less. So that's around twenty sandwiches per bird. But a good portion of the chicken isn't edible, and a lot of fat is lost in the cooking process. On the other hand, a chicken sandwich patty is not entirely made of chicken. The Canadian Broadcasting Company recently sponsored efforts by DNA researcher Matt Harnden (of the Trent University Wildlife Forensic DNA Laboratory) to analyse the contents of various fast food chicken sandwiches. They ranged between having 43% chicken DNA content (Subway) and 89% chicken DNA content (Wendy's). The rest was a mixture of soy and other additives. <https://www.cbc.ca/news/business/marketplace-chicken-fast-food-1.3993967>. So it seems like 1/20 is a decent estimate.

<sup>21</sup> KCF is a fictional company, but a real-world company with a similar name processes around 2.5 million chickens *per day* (i.e. 850 million chickens per year). So although the numbers I'm using here may seem large, they are probably much too small.

instead of 1 million from the supplier. They have similar thresholds every 100,000 sandwiches below that.

So now, finally, consider Oppy, who is morally opposed to farming chicken generally, and thus finds the industrial chicken system morally atrocious. He has walked into his local food court for a private lunch, and is trying to decide whether to purchase one of the delicious KCF spicy mesquite chicken sandwiches. Given our stipulations, if the number of *other* chicken sandwich orders during the month ends up being more than 19,899,999 but less than 20 million—as it has for the past decade—then Oppy’s decision won’t make any causal difference to the number of actual chickens harvested. For if he abstains, they will just dispose of or donate the patty that he would have purchased, and they won’t change their monthly order from the supplier.<sup>22</sup> Likewise if the total number is below 18,999,999 but above 17,999,999, then his decision also won’t make a difference. And so on. It is *only* if Oppy happens to be at the food court during a month in which the number of *other* sandwich orders, not counting his, is *precisely* 18,999,999 (or 17,999,999, or 16,999,999, etc.) that his order will have any effect at all. But if he *is* on such a threshold, and if he orders something other than a chicken sandwich, then his action will trigger a big lump reduction and KCF will order 5,000 fewer chickens from the regional supplier next month.<sup>23</sup>

The situation here is a difficult one for Oppy given his commitment to *Don’t Farm* and his taste for spicy mesquite. But here we can offer him a standard reply from the literature on expected utility. For even if the supply chain is lumpy like this, a consumer can never *know* when he is at a pivot or threshold point, and so he should just equate the *expected effect* of his single choice with the *average effect* of all the choices between thresholds. We have already stipulated that every ‘lump’ order is for 5,000 chickens, and that these are only triggered every 100,000 sandwich purchases. So the *average effect* of each purchase is  $1/100,000 \times 5,000 = 1/20$ . In other words, the effect of Oppy’s individual sandwich purchase is precisely what one would predict, given the amount of chicken in a patty—namely,  $1/20^{\text{th}}$  of a chicken. It’s then an empirical question whether the thought of occasionally saving  $1/20^{\text{th}}$  of a chicken (or of saving one chicken every 20 times he chooses the spicy Beyond patty) is sufficient to motivate Oppy’s abstention. Many philosophers who offer a reply like this argue that it should be.<sup>24</sup>

<sup>22</sup> I’m setting aside the fact that Oppy might order a salad instead, and thereby signal increased demand for plants instead of animals.

<sup>23</sup> This is essentially a large-scale version of the ‘Three-in-a-Boat’ game (named after a J.K. Jerome short story): If only one of us rows, the boat goes in a circle; if the two of you row, then my rowing won’t make a difference. So I should only bother rowing if *exactly* one of you is going to join me. Bovens (2015) argues that the ‘tragedy of the commons’ is best modelled as this kind of game rather than a Prisoner’s Dilemma. Thanks to Bovens for helpful discussion here.

<sup>24</sup> For versions of this argument, see Singer (1980), Norcross (2004), Kagan (2011), and Norcross (*draft*). Despite some similarities, this chicken sandwich scenario poses problems for both rationality and resolve that are somewhat different from real-world voting situations. When voting for president,



The expected utility reply would make sense in a system that was perfectly or at least highly efficient. But the reply is threatened by the sobering empirical facts about KCF's tolerance for oversupply and the decade-long trend in the region. Remember, KCF has received between 19,900,000 and 20 million chicken sandwich orders every month for ten years now, and so it is *extremely* likely that this will happen again this month. They also have a standing policy of throwing out or donating any oversupply between thresholds. This means there is a tolerance for *up to* 99,999 sandwiches' worth of oversupply in the system—a *buffer* that keeps them from ever falling short of demand. But it also means that an individual decision to purchase or abstain in a given month is likely to be drawing from that buffer, rather than moving the system along to the next threshold. This in turn suggests that it cannot simply have the 'average effect' of a purchase *between* thresholds. In a system with buffers, we have to consider not just what it takes to get from one threshold to another, but also the size of the buffer and the regional trend.

Some people may find it implausible that a real-world analogue of KCF would tolerate a buffer as large as 99,999 patties per month.<sup>25</sup> But even if that's correct and the buffer is smaller, it is almost certainly not the only one. In many industrial systems, even in those striving for vertical integration, there will be a *series* of buffers of this sort—a series of points along the chain where waste is tolerated or absorbed. KCF's suppliers, for instance, will presumably also have some sort of lumpy threshold policy vis-à-vis the farmers who raise the chickens, and be prepared to absorb some decrease in demand before reducing their order. Similarly, the farmer gets the chicks from an incubator facility, which itself will have a threshold policy that tolerates a certain amount of oversupply. Moreover, in a market as big as the chicken market, the suppliers almost certainly have other customers who will be able to pick up the slack in a given period, and so may not tell the farmer to produce 5,000 fewer chickens after all, even if KCF occasionally crosses the relevant threshold; likewise at the level of farmer and incubator, and so on: it's a network of buffers all the way down.

If this description is coherent, then it looks like there is not a 1/20 chance but rather a *much* smaller chance that Oppy, standing at the other-worldly food court that day, is at a point in the industrial supply chain that will make his choice

for instance, one might not make a difference to the outcome but still be able to signal support, or increase the 'manifest normative mandate' of the winner (see Guerrero (2010)). Something similar cannot be said about someone who abstains from a chicken sandwich in a lumpy supply system where her purchase would simply be taken from the buffer. For this reason I think assimilating inefficacy cases to voting cases (in the way that, e.g. Norcross (*draft*) and Nefsky (2017) do) can be misleading.

<sup>25</sup> Again, for more discussion of the situation in the actual world, see Halteman and McMullen (2018).

efficacious. If we add in other variables<sup>26</sup> and supply-chain ‘noise’, then the chance that he will buy a threshold sandwich becomes, in Budolfson’s words, ‘infinitesimal’ (2016: 208). Calculating the expected effect for each customer over the course of the month would be difficult, given all these factors, but it is clearly *nowhere near* as high as 1/20<sup>th</sup> of a chicken.<sup>27</sup>

This is an interesting thought-experiment in itself. But our actual world also seems regrettably similar to Oppy’s. If that’s right, then the sobering empirical facts problematize the intuitive move from *Don’t Farm* to *Don’t Purchase* in our world too. Rabbi Yanklowitz exemplifies this intuitive move in *The Jewish Vegan* when he states that ‘it is not hard to buy a leather-free belt to spare animals from unnecessary suffering’ (2015: 27). Forty years earlier, Peter Singer likewise suggested that:

[b]ecoming a vegetarian is not merely a symbolic gesture... Becoming a vegetarian is a highly practical and effective step one can take toward ending both the killing of non-human animals and the infliction of suffering upon them.  
(1975: 168–9).

Most of us share Yanklowitz’s and Singer’s sense that we have strong moral reasons to boycott products of wrongful activity on the part of others, i.e. that there is a way to defend not just *Don’t Farm* but also *Don’t Purchase*. However, if what morally matters is making a causal difference to the welfare of the creatures involved, then in our world—as in Oppy’s—lots of individual purchases of industrial animal products may be neither here nor there.<sup>28</sup>

<sup>26</sup> For instance: cultural trends can make the demand for certain *cuts* of poultry go up (turkey breast at Thanksgiving, chicken wings during years when the local football team does well, etc.). But animals come in wholes, and so if wings are in high demand this year, a decision not to eat a chicken sandwich (which is not made of wings) is even less likely to have any effect on how many chickens are produced and processed, since retailers are demanding as many wings as possible. Thanks to conversations with Cornell food scientist and poultry expert Joe Regenstein here; see also Parcell and Pierce (2000).

<sup>27</sup> Halteman and McMullen (2018) acknowledge these facts about buffers in the real-world poultry supply, but then note that the consumer will ‘know little to nothing about the distance to a threshold at their particular retailer, and even less further down the supply chain’. This is true, but in the noisy, buffered market situation described here, a consumer’s ignorance of precisely where she is relative to the next threshold seems moot—she only needs to recognize that the chance of being on any threshold is extremely small. Someone like Oppy can easily come to know that the vastly most likely outcome, given the buffer, the noise, and the trend, is that his decision to abstain will increase that month’s wasted oversupply by precisely one patty.

<sup>28</sup> Note that the sobering empirical facts also suggest that the choice to purchase small, family-farmed products is *much* more likely to make a difference. So if someone like Oppy thinks that *Don’t Farm* holds across the board, he should only indulge his opportunism with products from the industrial system.

### 11.3 Difference-making and Despair

The reasoning that motivates the negative conclusion we just reached (and the reasoning that Oppy exploits when he goes opportunist) is underwritten by something like this thought:

*Badness as a Reason to Abstain:* The badness of a production practice is a reason to abstain from purchasing its products *only if* abstaining has a non-negligible chance of making a positive difference with respect to the badness of that practice.

The more general principle here is:

*Badness as a Reason to Act:* The badness of an outcome O is a reason to perform action A *only if* A has a non-negligible chance of making a positive difference with respect to O.

In keeping with CCP, we are also construing what it is to ‘make a difference’ in a causal way:

*Difference-Making, Causal:* An action A makes a difference with respect to outcome O just in case performing A causally leads to a change in O.

It would be worth thinking more about whether consequentialist ethical theories can rebut the conceptual version of the inefficacy objection or rule out opportunistic purchasing without giving up one or more of their core principles. But I don’t propose to go further into that debate here.<sup>29</sup> Instead, I want to turn to the *psychological* version of the inefficacy problem, and the broadly Kantian argument from the moral need to avoid demoralization and despair.

So let’s now consider Hope, and suppose that she has been convinced by one of the many arguments in the literature that seek to bridge the gap from *Don’t Farm* to *Don’t Purchase* without appealing to causal difference-making. One family of such arguments invokes the symbolic value of ‘standing with the good’ and avoiding ‘moral taint’ by not consuming the products of a morally objectionable practice (see Hill 1979; Appiah 1986; Boey 2016). Another invokes a theory of ‘complicity’ or a ‘don’t benefit from wrongness’ principle that prohibits purchasing even if it makes no causal difference (Martin 2016; McPherson 2016). Another family of arguments appeals to virtue-theoretic or psychological considerations that don’t require any causal difference-making (Nobis 2002; Halteman/Halteman

<sup>29</sup> I attempted some of that in Chignell (2016). See Nefsky (2018) for an overview of such efforts.

Zwart 2016; Bramble 2016). Yet another emphasizes that an individual can be a *joint cause* of bad outcomes even without making a causal difference, and that the badness of being a joint cause can sometimes be a reason to abstain (Parfit 1984; Harman 2016; Albertzart 2019).

Difference-making concerns are not so easily dispatched, however. Even if Hope accepts *Don't Purchase* on one of these other grounds, the sobering empirical facts according to which her food choices with respect to the industrial system—even *over an entire lifetime*—are unlikely to make a significant difference can still have a profound psychological effect on her moral resolve. In other words, recognition of the sobering facts might lead to a very natural kind of *demoralization*:

(D1) General discouragement and dejection in the face of the very long odds of making a significant positive difference with respect to the outcomes in question.

Often this leads to another and more serious kind of *de-moralization*:

(D2) Loss of psychological resolve required to do what the agent *still* takes herself to have moral reason to do.

Let's call the state of being doubly demoralized in this way *despair*. When we fall into despair, we are liable to give up boycotts and revert to purchasing whatever is most convenient, tasty, and affordable—despite any ongoing moral qualms.<sup>30</sup>

Kant was acutely aware of our psychological vulnerability to perceived inefficacy in the face of large-scale evil and injustice. In one of his most lyrical passages in third *Critique*, he considers the psychology of an atheistic but 'righteous man (like Spinoza)' who

does not demand any advantage for himself from his conformity to the moral law, whether in this world or another; rather, he would simply and unselfishly bring about the good to which that holy law directs all his powers.

<sup>30</sup> This is clearly anecdotal and conjectural. To make the case we would need to look at sociological studies of consumer and activist behaviour in response to perceptions of inefficacy. Interestingly, the Meat Institute itself cites recent CNN and USDA survey data showing that of the approximately 5 percent of Americans who claim to be vegetarians, around 65 percent will confess to having eaten meat products in the past 24 hours. The Meat Institute's conclusion? 'Bottom line: meat is amazing—and irresistible'. (North American Meat Institute, n.d., 2) <https://www.meatinstitute.org/index.php?ht=d/sp/i/101931/pid/101931>, accessed 1 November 2019.

Such a person's resolve, Kant says, is still liable to be sapped by the perception that his own moral strivings, as well as those of others, are inefficacious against large-scale structural evil and injustice:

But his strivings (*Bestreben*) have limits . . . Deceit, violence, and envy always surround him, even though he is himself honest, peaceable, and benevolent. The other righteous people that he encounters at times will, in spite of all their worthiness to be happy, nevertheless be subject by nature, which pays no respect to that, to all the evils (*Übeln*) of poverty, illnesses, and untimely death, just like all the other animals on earth. It will always remain so until one wide grave engulfs them all together (whether honest or dishonest, here it makes no difference) and hurls them, the very ones who were capable of believing that they were the final purpose (*Endzweck*) of all creation, back into the abyss of the purposeless chaos of matter (*Schlund des zwecklosen Chaos der Materie*) from which they all were drawn. (CPJ 5:452)

In the end, when faced with the abyss in this way, the righteous Spinoza has two options: either he will 'certainly have to give up his end [i.e. of being righteous] as impossible' or 'he will have to accept (*annehmen*) the existence of a moral author of the world (*Welturheber*), i.e. of God, from a practical point of view.'<sup>31</sup> I read 'impossible' here psychologically: even someone as righteous as Spinoza simply cannot maintain his resolve without some source of hope that, ultimately, justice will prevail.<sup>32</sup>

These reflections on demoralization suggest that, in addition to the sobering empirical facts about the industrial food system, there is another empirical question that we need to consider in connection with inefficacy. It is an empirical-psychological question about which sorts of background commitments or mindsets do better by way of helping ordinary agents avoid despair in the face of *perceived* inefficacy. Kant clearly thought that someone who accepts the existence of a benevolent deity working behind the scenes to bend the arc of history towards justice—presumably also recording and rewarding our efforts—will have more resources to sustain those efforts, despite long-term setbacks and perceived inefficacy. Kant is not alone in this: supersensible commitments are clearly effective in both providing motivation and sustaining resolve, and many 'moral arguments' along these lines can be found throughout the nineteenth-century Pragmatist, Idealist, and Existentialist traditions. But Kant is the *éminence grise*: 'What may I hope?', he says in the *Critique*, is one of the three questions that motivate his

<sup>31</sup> 'Accept' here is '*annehmen*', the verb Kant often associates with the noun '*Glaube*' (faith).

<sup>32</sup> Camus (1942) famously rejects this: he considers the same predicament and says that we must both accept the demands of the moral law and embrace the absurdity of a world in which justice never prevails: 'One must imagine Sisyphus happy'.

entire philosophy. So in the next section I propose to look more closely at the structure of his moral-psychological proof. After that we can consider whether some analogous forms of reasoning can help with psychological inefficacy problems in contemporary consumer contexts.

### 11.4 Moral Arguments for Theoretical Conclusions

Kant's project in *Religion Within the Bounds of Reason Alone* (1793) is to discover the rational essence of religion—the doctrines and practices that can survive examination by the tribunal of universal rationality. The result is called 'moral religion', 'rational religion', or 'rational faith' (*Vernunftglaube*) in order to distinguish it from creedal or enthusiastic forms based in special revelation or alleged mystical experience. At one point, he offers the following as the rational essence of the doctrine of providence:

Each must, on the contrary, so conduct himself as if everything depended on him. Only on this condition *may he hope* that a higher wisdom will provide the fulfilment of his well-intentioned effort. (Rel. 6:101, my emphasis)

The passage displays what might be called a *consequence-dependent moral psychology*. Kant is famously not a consequentialist in ethics generally: 'The fulfilment of duty consists in the form of the earnest will, not in the mediating causes of success' (CPJ 5:451). Nor is he a divine command theorist: 'Morality... needs neither the idea of another being above him in order to recognize his duty nor as an incentive anything other than the law itself in order to observe it' (Rel. 6:3). But Kant also recognizes that it is often psychologically crucial to be able to *hope* that our actions will have their intended results, even if by way of a supersensible mechanism ('a higher wisdom').

This fits with a general pattern of sensitivity to embodied, empirical psychology in Kant's ethical thought—i.e. sensitivity to the fact that it is hard to keep soldiering on in the moral life without the hope that it will do some good. The best state for human beings is not (as the Epicureans say) one of mere desire-satisfaction, but neither does it consist (as the Stoics suggest) simply in duty and virtue alone. Rather, the best state—the 'highest good'—is a state in which happiness is perfectly proportioned to moral worth. But happiness itself is at least partly bodily: it involves the satisfaction of our inclinations, many of which arise from our sensible nature. And in the *Critique* he says that 'all hoping aims at happiness' (A805/B833). So despite the traditional picture of Kant as a rule-obsessed Prussian bachelor, he was in fact keenly aware that we are embodied 'creatures of need'—animals for whom happiness consists largely in bodily and emotional welfare (CPrR 5:61). Although he does not endorse the satisfaction of our inclinations willy-nilly (cf. his infamous remarks about 'defiling oneself by lust'), he

consistently characterizes happiness as the proper satisfaction of the ‘sum total’ of them (MM 6:424–6; G 4:399, 418).

Behind his theory of the highest good, then, is Kant’s recognition that prolonged experience of injustice—of a world-history in which the wicked prosper and virtue seems to make no difference—can chip away at our resolve. In the *Religion* passage quoted above (from 6:101), he acknowledges that one of the things we hope for—and thus one of the things that make us happy, in the context of ethical and political action—is having our ‘well-intentioned efforts’ fulfilled. For example: suppose the moral law demands that we try to help the disadvantaged by donating money to various causes; for Kant, this will be true regardless of actual outcomes. Still, he recognizes that most of us *also* have a ‘need’ for our altruism to be effective<sup>33</sup>: it helps us retain *not our rational motive but our psychological resolve* if we can believe (or at least reasonably hope) that the needy are genuinely benefiting from our efforts. So even in a Kantian context, a morally good person can reasonably care a lot about the goodness of the *consequences* of her willing, and not just about the goodness of her will.

This is a complicated piece of reasoning, so it is worth looking at it in stepwise fashion:<sup>34</sup>

- (1) I ought to do what is morally right. [Independent argument]
- (2) For me, it would be *demoralizing* in the (D1) sense (i.e. it would lead to discouragement and dejection) not to be able to have substantial hope that there is a moral order by which a just arrangement (i.e. a ‘moral world’) will come about, for then I would have to regard it as certain that the entire history of the world will not be good on the whole, no matter what I do. [Empirical premise, concept of ‘hope’]
- (3) Such demoralization has an enervating effect on my resolve, and is thus *demoralizing* in the second, (D2) sense: I will almost certainly no longer perform actions that I take to be morally good or required. [Empirical premise]
- (4) Despair of this sort is seriously morally undesirable. [From (1)–(3)]

<sup>33</sup> The ‘effective altruism’ movement introduces itself as follows: ‘Most of us want to make a difference. We see suffering, injustice, and death, and are moved to do something about them. But working out what that ‘something’ is, let alone actually doing it, is a difficult problem. It would be easy to be disheartened by the challenge’ (<https://www.effectivealtruism.org/articles/introduction-to-effective-altruism>, accessed 1 November 2019).

<sup>34</sup> Robert M. Adams sketched an empirical-psychological way of reconstructing Kant’s proof in his (1979). My presentation here is influenced by his, though also different in many key respects. In the first *Critique*, Kant presents an early version of the moral proof according to which hope for happiness seems to be part of the *incentive* for acting rightly. In the second *Critique*, he provides the more canonical articulation of the proof that starts not with (1) above but with the more controversial claim that we ought to will the highest good, and this involves, as a sort of presupposition, adopting moral faith in the existence whatever is required to make the highest good really possible (i.e. God and the after-life). For articulations of that form of the proof, see Wood (1970) and Willaschek (2016). Very little attention has been devoted to the moral-psychological argument from despair in the third *Critique* and *Religion*, although see Fugate (2015) and Ebels-Duggan (2016).

- (5) Therefore, there is serious moral advantage for me in being able to have substantial hope that there is a moral world order. [From (2)–(4)]
- (6) Substantial hope that *p* requires the positive belief or faith that *p* is really, practically possible. [Theoretical premise]
- (7) Therefore, there is serious moral advantage for me in being able to believe or have faith that a moral world order is really, practically possible. [From (5)–(6)]
- (8) If a being or state of affairs is really practically possible, then there must be something in the actual world that can account for that fact. [Concept of ‘real practical possibility’]
- (9) The *actual* existence of God provides the only adequate account of the real, practical possibility of a moral world order. [Theoretical premise]
- (10) Therefore, there is serious moral advantage, for me at least, in being able to believe or have faith that God exists. [From (7)–(9)]
- (11) There are no good epistemic reasons either for or against the existence of God. [Results of examination of natural theology in Transcendental Dialectic]
- (12) Rational belief requires good epistemic reasons. [Kant’s evidentialism about belief<sup>35</sup>]
- (13) Therefore, belief in God’s existence or non-existence is irrational. [From (11)–(12)]
- (14) Rational *faith* (*Vernunftglaube*) does not require good epistemic reasons; it can instead be based on good moral or pragmatic reasons. [Conceptual truth<sup>36</sup>]
- (15) Therefore, faith (though not belief) that God exists is *prima facie* morally (though not epistemically) justified, for me at least. [From (10), (13), (14)]

The ‘for me at least’ sounds worrisome, but Kant is explicit: ‘I must not even say ‘It is morally certain that there is a God’, etc., but rather ‘I am morally certain’ etc.’ (A829/B857, original emphasis). Still, the proof does not count as an exercise in mere wishful thinking or self-deception, provided a few conditions are met:

- The justification is moral rather than epistemic
- The justification is defeasible
- The result is not the attitude that we would call ‘belief’ and Kant would call logical ‘conviction’ (*Überzeugung*). The result, rather, is *faith* (*Glaube*).

<sup>35</sup> Kant is a conceptual evidentialist about what *we* would call ‘belief’—the kind of holding-for-true (*Fürwahrhalten*) that can count, if true and justified, as knowledge (*Wissen*) (Kant calls it ‘conviction’ (*Überzeugung*), as do many contemporary German epistemologists). See Chignell 2007a.

<sup>36</sup> Rational faith, for Kant, is a voluntary state of holding-for-true (*Fürwahrhalten*) that, for non-epistemic reasons, a subject uses to guide deliberation, action, and assertion in certain contexts. See Chignell (2007b).



- The faith in question is about a state of affairs that is evidentially ambiguous (in this case theoretically undecidable; see discussion below)

Let's take a closer look at a couple of the key premises. Regarding (6): everyone agrees that there is some sort of modal constraint on hope, but it is clear that superficial hopes often involve merely taking possibility for granted (see Chignell 2014). Deep, life-structuring hopes that *p*, on the other hand, arguably do require moderately clear conceptions of how, given the things and powers that exist in the world, *p* could turn out to be true.<sup>37</sup> It would be extremely hard to hope in a serious way that there will be perpetual peace—to really structure one's life around this hope, for example—without having at least some positive conception of how we denizens of the actual world could really get there from here. This is what I mean by 'really, practically possible'.<sup>38</sup>

(9) is obviously another lynchpin; a friend of the argument would have to rule out other accounts of the real, practical possibility of a moral world: dialectical historical processes, political revolutions, a karmic system that ensures that justice will be done, liberal democratic institutions, and so on (compare Adams 1979). She would also have to explain why the full-blown classical deity, with all its omni-properties, is required to provide the relevant explanatory ground. Kant himself makes some efforts in this direction, but many readers have been unconvinced (see Michalson 1999).

(11)–(14) are also crucial: the existence-claim here must be *theoretically undecidable* in order for the faith to be rational. Large swaths of the *Critique of Pure Reason* are dedicated to undermining traditional efforts to *prove* or even render *probable* the existence of God on either demonstrative or empirical grounds. Kant also rejects all atheistic arguments, including the empirical argument from evil. So, in the famous phrase, *all* knowledge-claims about God's existence—theistic and atheistic—are 'denied' in order to 'make room' for the kind of moral faith that overcomes demoralization.

In the *Religion* passage quoted at the beginning of this section, Kant is not discussing a moral world order generally but rather hope for *specific* outcomes of *specific* actions. He clearly thinks that this hope, too, justifies faith in the existence of a supersensible mechanism: a providential 'higher wisdom' that makes the fulfilment of our well-intentioned efforts possible. Returning to the contemporary inefficacy problem: it seems clear that people who *already* have such a supersensible commitment can just focus their hopes around that. They can focus, in other words, on the possibility that this mechanism arranges things such that some of their choices make a significant positive difference—both unobservable

<sup>37</sup> For the claim that there is a distinction between 'superficial' hope and 'substantial hope', and that the latter has extra conditions on it, see Pettit (2004) and McGeer (2004).

<sup>38</sup> Compare Willaschek on 'practical possibility' (2016).

and observable. Their faith allows them reasonably to hope that their efforts *matter*—that they are recorded somehow, even if they seem to be inefficacious. And that hope would presumably sustain psychological resolve to keep doing what they regard as required. But is a version of this argument against demoralization available in secular contexts as well?

### 11.5 A Moral-Psychological Argument for Hope

Given that the industrial poultry system is deeply insensitive to slight changes in demand, any attempt on Hope's part to inculcate full-blown belief that *my abstention is going to make a significant positive difference* would be a miserable exercise in self-deception. *Faith* in that proposition also looks more like a Kierkegaardian leap than sweet Kantian reason. As we have seen, Kant anticipates James and others in the moderate pragmatist tradition in rejecting the idea that a moral argument can support a conclusion whose *negation* we have strong epistemic reason to believe.

With this in mind, let's consider a moral argument that is based, like Kant's own 'proof', in a claim about what's required for the substantial hope that sustains resolve:

- (1\*) *Don't Purchase*: It is morally wrong to purchase the bodies, secretions, or eggs of farmed animals if there are readily available alternatives to eat and wear. [From an independent argument or moral intuition]
- (2\*) It would be *demoralizing* in the (D1) sense (i.e. it would lead to discouragement and dejection) for me not to be able to have substantial hope that my abstention from animal products over time will make a significant positive difference with respect to animal welfare. [Empirical premise]
- (3\*) Such demoralization has an enervating effect on my resolve, and is thus *de-moralizing* in the second, (D2) sense: I will almost certainly fail in my efforts to abstain over time. [Empirical premise]
- (4\*) Despair of this sort is seriously morally undesirable. [From (1\*)–(3\*)]
- (5\*) Therefore, there is serious moral advantage for me in being able to have substantial hope that my abstention will make a significant positive difference with respect to animal welfare. [From (2\*)–(4\*)]

So far, so good: this looks like a simple argument for the practical rationality of *hope* that by abstaining I will somehow make a significant difference to the system in question. Note that 'significant' is left vague precisely because the amount of difference-making required to avoid despair will clearly differ from person to person.

But what are the conditions on having this sort of hope? We saw earlier that

- (6\*) Substantial hope that  $p$  requires the positive belief or faith that  $p$  is really, practically possible. [Theoretical premise]

But then from (5\*) and (6\*) we can infer:

- (7\*) Therefore, there is serious moral advantage for me in being able to believe or have faith that it is really, practically possible that my abstention will make a significant positive difference.

People of great willpower may find it psychologically possible to leave things there. Even if they dearly miss the taste of deep-fried spicy mesquite chicken, they will abstain in the stubborn hope that this behaviour, over time, will *somehow* make a significant positive difference. Despite the incredibly long odds (remember those 2000 birds per second), they are able to maintain belief or faith that it is really possible—and that is enough to preserve their resolve.

But Kant thought such moral saints are rare: even someone as righteous as Spinoza might despair in the absence of *some* sense of how justice might prevail. For the rest of us, clinging to the brute possibility that we just might make a difference won't be enough for long-term resistance, especially when the products in question are so convenient, tasty, and cheap. This is one reason why Kant's moral-psychological argument proceeds from a claim about what is really possible to a claim (in (8) and (9)) about the *actual ground* or *explanation* of that possibility: we need to have a sense, not just *that* it is really possible, but *how* it is so. What is it about the actual world that makes the hoped-for outcome really, practically possible?

Well, as we saw at the end of the last section, one idea is just this:

- (9\*) The actual existence of God provides the only adequate account of the real, practical possibility that my abstention will make a significant positive difference. [Theoretical premise]

If this were correct, then the rest of the argument could run like the previous one and conclude with full-blown moral theism. Only such faith, this version of the argument might say, makes it psychologically possible to retain the hope that my abstaining (and/or my abstinence over time) will make a significant positive difference to a morally objectionable system. So that hope and that faith, together, allow me to avoid despair, and are *prima facie* morally justified as a result.

The problem here is obvious: (9) was debatable, but (9\*) is just patently false. Perhaps there is some plausibility in the idea that the existence of a supreme being is the *only* adequate account of the real, practical possibility of a *perfectly* just

world order (the highest good). But in the case of the more localized hope that my individual actions will make a significant positive difference with respect to the industrial chicken system, (9\*) itself is hopeless. We already saw back in Section 11.2 that there are naturalistic ways to account for the bare possibility here, even in a massive and lumpy supply system. It's just *barely* possible, for instance, that this is the month in which everyone else will purchase exactly 18,999,999 sandwiches, and thus that Oppy is indeed going to stand on one of those thresholds when he walks into the food court. If that were the case, then a choice on the threshold to abstain will make a significant positive difference (5,000 chickens saved!).

Given the availability of naturalistic ways of accounting for the real possibility that my abstinence makes a difference, then, there can be nothing like a moral *proof* that God, the 'universe', karma, or fate is arranging for individual efficacy. Again, if someone already has a supersensible commitment that can do the job, then it makes sense for *him* to fix his hopes by appealing to a premise like

(9\*\*) The actual existence of *supersensible mechanism X* provides an adequate account of the real possibility that my abstinence will make a significant positive difference.

But for someone who does not already have the supersensible commitment, there is no compelling pressure to adopt one.

No *compelling* pressure. There might be a little pressure, though, depending on how we think about the naturalistic options and the psychological constraints on hope. I just said that someone like Hope who is threatened by demoralization can try to focus on the 'infinitesimal' possibility of her own pivotality—she can give that possibility a prominent place in her mind every time her colleagues drag her to the bar on 'Shotz-n-Wingz Nite', and this way allow the idea of 5000 chickens saved to sustain her resolve.<sup>39</sup> But that's not going to be easy: given the sobering empirical facts, the chance that she is at a threshold point there in the bar is vanishingly small, and it may be hard for Hope to keep from sliding into despair (especially after a few shotz). Indeed, according to the account of hope that I favour, the key difference between hope and despair has to do with where the subject is disposed to place her mental attention or 'focus'. I might strongly desire something and regard it as possible, but if I keep focusing on it as massively improbable, I am in despair. If instead I am disposed to focus on it 'under the aspect of its possibility'—i.e. as having at least a chance of occurring—then I am

<sup>39</sup> For more on how 'Shotz-n-Wingz Nite' can destroy a person's resolve, see Halteman and Halteman Zwart (2016: 131). Those authors claim that such events occur regularly at a place called 'Baloneez', but Barnhill et al. (2015: 171) reference a sister establishment called 'Jimmy's You-Hack-it-Yourself BBQ'.

hoping.<sup>40</sup> But other things equal, the more unlikely a desired outcome is, the harder it is to stay mindfully focused on it *as possible*, rather than fixating on the overwhelming odds against. So while hoping for pivotality is an option for staving off despair, most of us will find it a hard psychological row to hoe.

In light of this, it would be useful to find some other naturalistic scenarios whereby Hope's abstinence might make a significant positive difference for animals. A candidate that comes to mind is this: Hope's decision might somehow be connected to the decisions of numerous other people such that, if she abstains, then a significant number of other people will or will be likely to abstain (where a 'significant' number is precisely what's required for there to be a 'significant' difference made with respect to animals—keeping in mind that what counts as 'significant' might be different across different subjects). This scenario divides into two: one causal, and one evidential. I'll discuss each in turn.

The causal version of the scenario would invoke this premise:

(9\*\*\*) The existence of a *causal connection* between my decision to abstain and the decisions of a significant number of other people to do the same provides an adequate account of the real, practical possibility that my abstinence will make a significant positive difference.

This seems true: such a causal connection would provide an adequate account. Following the logic of the moral-psychological argument, we would then arrive at:

(10\*) Therefore, there is serious moral advantage, for me at least, in being able to believe or have faith that such an interpersonal causal connection exists.

[From (7)–(9\*\*\*)]

But could such belief or faith be rational? That is, could it ever be rational to believe or have faith in the following?

*Interpersonal Causal Connection:* If I choose to abstain, my action will cause a significant number of other people to abstain, and if I choose to purchase, my action will cause a significant number of other people to purchase.

*Causal Connection* looks like a non-starter for people who aren't celebrities, dictators, or top-level 'influencers.' There is no reason to think that there is any such connection between one's person's actions and the actions of a significant number

<sup>40</sup> See Chignell (*draft*). Note that I am not claiming that hope just *is* the desire and the presumption of possibility plus a disposition to focus on the outcome as possible. But I do hold that substantial hope is at least *accompanied* by these states.

of others.<sup>41</sup> A quick look at Hope's number of Twitter followers provides decisive reason to think that there is *not* such a connection in her case. But then both belief and faith are ruled out on rational grounds (again, Kant is the forerunner of James rather than Kierkegaard here—a live option for moral faith has to be evidentially ambiguous).

There is a weaker version of the scenario, however, that looks more promising. Consider:

(9\*\*\*\*) The existence of a **strong evidential connection** between my decision to abstain and the decisions of a significant number of other people to do the same provides an adequate account of the real, practical possibility that my abstinence will make a significant positive difference.

(9\*\*\*\*) seems true, and so by the logic of Kant's moral-psychological argument, it would seem to give us *prima facie* moral justification for either belief or faith that:

*Interpersonal Evidential Connection:* If I choose to abstain, that is strong evidence that a significant number of other people are likely to abstain, and if I choose to purchase, that is strong evidence that a significant number of other people are likely to purchase.

Is this a principle for which we have any grounds one way or the other? Here I think things are less obvious than they were with *Causal Connection*. It seems reasonable for Hope to think that a significant number of people *could* be motivated by the same reasons against purchasing animal products that she is (that might be a part of what it is to take them to be *reasons*). But that's different from thinking that if she chooses to abstain, a significant number of people *will* be so motivated, or are *likely to be* so motivated. Given that most of these decisions will be made in distant places and times (and given that she is not a top-level influencer), it is *simply unclear* whether there is any evidence for *Evidential Connection*. Indeed: even if, over the short- to medium-term, she acquires evidence that other people around her are *not* following suit, her choice might still be strong evidence that *over the long haul* a significant number of people (here or elsewhere) will

<sup>41</sup> Obviously Hope could try to *establish* such causal connections (and thereby shore up her resolve) by signalling her choices publicly when she can, *attempting* to influence others, becoming part of a broader cultural movement, and so on. And obviously the question of whether her actions have made any causal difference must be considered over the longer haul. It's not clear how much this will help in private one-off moments under conditions of perceived inefficacy, however. Thanks to Elizabeth Harman and Renée Jorgensen Bolinger for discussion here. See also Lawford-Smith (2015).

come around and see the power of her reasons.<sup>42</sup> This means that in most epistemic contexts, anyway, belief either way would be out of rational bounds.

But here is where Kant's famous idea about denying belief to make room for faith comes in. When there is no clear evidence one way or the other, we might still have *moral* reason for adopting propositions in this other non-doxastic way. So if supersensible mechanisms aren't a live option for Hope, and if her own causal pivotality is too improbable to sustain her focus, she might still be defeasibly morally justified in taking *Evidential Connection* on faith. That is: she might be defeasibly justified in holding that a significant number of people (somewhere, sometime) will be similar enough to her to act on the same reasons in a broad range of cases. Thus, if she chooses to abstain now, she will have strong evidence that there will be a significant improvement in farm animal welfare. Faith like that would presumably suffice to stave off Hope's despair. However, if she think like this up to the moment of choice, but then defects and purchases (perhaps believing that all those other people will still act on the reasons that she has), she thereby *loses her evidence* that others will abstain. That is why it is crucial that *Evidential Connection* says that she only gets the relevant evidence if she *actually* abstains, and not just if she appreciates the reasons for abstaining.<sup>43</sup>

There is obviously more to be said about this last scenario, but here I simply want to note two of its most intriguing features. First, the reasoning is broadly Kantian in a manner that goes beyond the mere use of moral arguments. Hope relies on the idea that (other things equal) if she has a good moral reason to do something, then it is also good moral reason for others to do likewise. This is an implication of the universalizability of moral reasons—a theme that we typically associate with Kant. I have a good reason to act on a certain maxim *only* if it is a good reason for everyone else to do likewise in saliently similar circumstances. But the universalization goes the other way, too: if it is a good reason for *me* to act in such-and-such a way, then it is a good reason for everyone else in similar circumstances to do likewise. When Hope 'thinks for herself' about these matters, she also takes herself to be 'thinking from the point of view of others'—a maxim that Kant thinks of as common sense (CPJ 5:294). And so her moral faith in *Evidential Connection* invokes a broadly Kantian idea about the universalizability of reasons.

Second, and even more intriguingly, if faith in *Evidential Connection* is what sustains the hope to make a difference, then we seem to have arrived in a very roundabout way at a key principle of *evidential decision theory* (EDT). EDT says

<sup>42</sup> There are many dissimilarities between the cases, but just by analogy think of someone in late 17<sup>th</sup> century South Carolina who decides, on moral grounds, to free the slaves that he inherited from his family. Those reasons were good, and in order to sustain his resolve he might also have taken on faith that his choice was strong evidence that others would (at some point) do the same. But it took a long while.

<sup>43</sup> Thanks to Victoria McGeer for discussion here.

(roughly) that an agent ought to perform actions that are such that, if she performs them, then the chances of the desired outcome are maximized, regardless of whether there is any causal connection between her actions and the outcome. It is sometimes characterized as the choice to be guided by ‘auspiciousness over efficacy’; it can also be articulated as a theory of difference-making:

*Difference-Making, Evidential*: An action A makes a difference with respect to outcome O just in case performing A provides *evidence* that a change in O will occur.

According to this principle, no causal connection is required for an action to make a difference. That is why EDT is thought to support taking the one box in the Newcomb problem.<sup>44</sup> Even if there is no causal connection between your choice and the action of a near-perfect predictor of your choice (since the prediction has already been made), there is an *evidential* connection: what you do provides strong evidence regarding what the predictor did. So on the evidential conception of difference-making, your action does ‘make a difference’ and you should take just the one box (thus containing a million dollars). On the causal conception of difference-making, your action can’t make a difference and so you should be safe and take both boxes (thereby acquiring a thousand dollars).<sup>45</sup>

On the scenario we have just been considering, Hope abstains for independent reasons (involving e.g. symbolic value worries or concerns about some sort of complicity), but she also has a psychological need to hope that her action makes a difference, especially over time. If she takes *Evidential Connection* on as an item of moral faith, with *Difference-Making, Evidential* in the background, then she can retain faith that her action makes a difference even while recognizing that the sobering empirical facts strongly suggest that her action *doesn’t* make a causal difference.<sup>46</sup>

<sup>44</sup> Originally stated in Nozick 1969.

<sup>45</sup> Jon Elster cites empirical evidence for the claim that people use ‘diagnostic thinking’ to move from the assumption that they are ‘fairly typical members’ of a reference group to the conclusion that others ‘will tend to act like me’. Elster sees no flaw in this sort of reasoning *unless* explicitly causal claims are made—i.e. unless people start thinking that ‘my action will bring it about’ that others do the same. But the EDT conception of difference-making avoids that sort of ‘interpersonal magic’. See Elster (1985: 142–5). Thanks are owed here to John Pittard, Ewan Kingston (who pointed me to Elster), Philip Pettit, Richard Bradley, and Kian Mintz-Woo, as well as to the unknown author of a paper I recently refereed in which a connection between one-boxing in Newcomb and an evidentialist account of difference-making was made. Mintz-Woo and Bradley point out to me that this interpersonal ‘evidentialist’ approach to consumption and voting cannot rationally be extended to Prisoner’s Dilemma cases, however, since defecting is *always* the best thing to do in the latter, even on EDT. This is presumably why followers of EDT have resisted Lewis’s claim (in 1979) that ‘Newcomb’s Problem is a Prisoner’s Dilemma’. See also Pittard (2018).

<sup>46</sup> If this is correct, then the present broadly Kantian response to the psychological inefficacy problem has the added benefit of providing *prima facie* moral justification for faith in a principle that



## 11.6 Conclusion

Recent discussions of the inefficacy problem focus primarily on the conceptual side of the issue—that is, on how inefficacy poses a challenge to traditional consequentialism. In the bulk of this paper, I focused instead on the psychological side of the problem—that is, on how the *perception* of our almost certain inefficacy can lead to the kind of despair that, in turn, threatens to undermine moral resolve. I suggested that although Kant’s own narrow view regarding how we ought to treat animals has little appeal, his moral-psychological argument against despair can be applied to justify various resolve-sustaining strategies. These include focusing in hope on the fact that:

A. An extremely unlikely threshold event can cause at least one of my abstaining actions to make a significant positive difference to outcomes.

Alternative strategies involve adopting defeasible moral faith in one (or more) of the following claims:

B. A supersensible mechanism exists and can causally connect my abstinence to some significant positive difference in outcomes.

C. My choosing to abstain is strong evidence that there will be a significant positive difference in outcomes, even if there is no causal connection between the two.

Although it has a structure that is analogous to Kant’s moral-psychological argument for the existence of God, only one of the options here involves faith in something supersensible.

A final, Kant-scholarly point: the structure of the moral-psychological argument I’ve been considering here reveals an underappreciated way in which the practical has ‘primacy’ for Kant. The argument says that if we morally ought to act a certain way, and we are threatened by resolve-sapping despair, then we are *prima facie* morally justified in seeking strategies that will sustain our hope and thus our commitment to the ought in question. But although the moral commitment and the concomitant hope have *primacy*, Kant is not an advocate of irrational leaps. Theoretical reason does kick in at some point and require a coherent account of how the hoped-for scenario could really, practically come about. That in turn makes us *prima facie* morally justified in having faith in

supports one-boxing. There is something right about the prosperity gospel after all: faith can indeed make you rich...

whatever can adequately explain that practical possibility, as long as it is evidentially ambiguous. The faith in question can even involve a 'theoretical' issue: an existence-claim, for instance, or a principle in decision theory.<sup>47</sup>

<sup>47</sup> For feedback on these ideas as they developed, I am grateful to the editors of this volume and to Anne Barnhill, Renée Jorgensen Bolinger, Luc Bovens, Mark Budolfson, Gabriel Citron, Silvia De Toffoli, Tyler Doggett, Gabriele Gava, Matthew Halteman, Jakob Huber, Tania Lombrozo, Victoria McGeer, Kian Mintz-Woo, Philip Pettit, Michael Milona, Leigh Vicens, and Marcus Willaschek. I am particularly indebted to McGeer for lengthy written comments. I also thank audiences at Humboldt University in Berlin, Goethe University in Frankfurt, the Hungarian Academy of Sciences in Budapest, Georgetown University, Princeton University, and the unforgettable 'Kant and Animals' conference sponsored by University of Witswatersrand in the Krüger Animal Park, South Africa.