

PRIORITARIANISM A (PLURALIST) DEFENSE

Shai Agmon and Matt Hitchens

PRIORITARIANISM is the distributive view that welfare gains matter more, morally, the worse off you are.¹ A common and intuitively compelling objection to prioritarianism is that it wrongly treats cases involving one person (intra-personal cases) like cases involving more than one person (inter-personal cases), when they should be treated differently. In a nutshell, the objection goes as follows. A person is allowed, when faced with an intrapersonal choice between someone else's possible futures, to reason prudentially when choosing on their behalf.² She is not required to give special moral weight to the future in which the person for whom she is choosing would be worse off. However, when choosing how to distribute goods between multiple people, prudential reasoning on behalf of the group is not justified, as the claim of the person who is worse off should matter more (call this the *Moral Shift*). Thus, prioritarianism, which as an aggregative, impersonal view is committed to treating intrapersonal and interpersonal trade-offs similarly, cannot explain the Moral Shift.

Opinion is divided among philosophers as to whether this objection is powerful enough to reject prioritarianism.³ In a 2009 article, Michael Otsuka

- 1 Throughout this article, for clarity and consistency, we refer to "welfare" rather than "utility" or "well-being." Our definition is not intended to differ in any meaningful way from other standard canonical formulations of prioritarianism. For examples, see Arneson, "Luck Egalitarianism and Prioritarianism," 340; Broome, *Weighing Goods*, 199; Parfit, "Another Defence of the Priority View," 401; O'Neill, "Priority, Preference and Value," 335; Segall, "In Defense of Priority (and Equality)," 344; Otsuka and Voorhoeve, "Why It Matters That Some Are Worse Off than Others," 176.
- 2 In all the cases to come, we assume that whatever set of relations is required for prudential concern (e.g., identity, psychological continuity, the prudential unity relations) holds, and holds to the same degree, between all the current and future people under discussion.
- 3 For example, McCarthy thinks this objection is sufficient to reject prioritarianism, but McKerlie and Rabinowicz do not. Interestingly, McCarthy, *contra* Otsuka and Voorhoeve, argues that prioritarianism should be rejected in favor of utilitarianism (rather than a pluralist egalitarianism), which does not recognize the Moral Shift. On McCarthy's view, the

and Alex Voorhoeve level a famous argument in favor of this proposition.⁴ They claim that prioritarianism's inability to explain the Moral Shift means it ignores the unity of the individual, and hence should be rejected altogether. In this article, we show that their argument has a significant weakness, as its only answer to a standard prioritarian response is self-defeating. To explain briefly: a natural way in which the (pluralist) prioritarian can explain the Moral Shift is to appeal to the autonomy of the person being chosen for, as the prioritarian is choosing as the person would choose for herself. In order to preempt this, Otsuka and Voorhoeve assume the decision maker is a morally motivated stranger who knows (almost) nothing about the preferences and attitudes of the person in question. This assumption, we argue, makes the examples on which their argument relies incoherent, unless they rely on assumptions that most prioritarians would (we believe) find implausible. Thus, we argue that most prioritarians can both continue to appeal to autonomy to explain the Moral Shift and maintain their prioritarian commitments.

The article is structured as follows. First, we sketch the commitments of the prioritarian, and set out the objection to her view that we will be discussing. Next, we spell out the pluralist prioritarian's natural response to the Moral Shift—an appeal to the autonomy of the person being chosen for. We show that Otsuka and Voorhoeve can only head off this response, and so continue to press their objection, on the basis of two controversial assumptions: that welfare consists in a list of objective goods, and—if one takes an unorthodox but plausible view of risk aversion—that there is only a narrow range of rational risk aversions (or, analogously, that the prioritarian's weightings are extreme). These assumptions, we argue, significantly limit the scope of Otsuka and Voorhoeve's argument, and make it applicable for only a limited range of prioritarians. Along the way, we address possible responses on Otsuka and Voorhoeve's behalf, and show that they are unsuccessful.

1. PRIORITARIANISM

As mentioned, prioritarianism is the distributive view that welfare gains matter

Moral Shift is not justified to begin with. In this article, we assume that the Moral Shift exists and needs justification, and try to explain it in prioritarian terms. We do not address the claim that there should be no Moral Shift whatsoever. See McKerlie, "Dimensions of Equality"; Rabinowicz, "Prioritarianism for Prospects" and "Prioritarianism and Uncertainty"; McCarthy, "Utilitarianism and Prioritarianism II."

4 Otsuka and Voorhoeve, "Why It Matters That Some Are Worse Off than Others." Otsuka has defended and expanded on this critique in Otsuka, "Prioritarianism and the Separateness of Persons" and "Prioritarianism and the Measure of Utility."

more, morally, the worse off you are; the prioritarian unit of moral value is priority-weighted welfare. Prioritarianism can take a variety of forms, depending on how steeply the moral weighting applied to a unit of welfare decreases as the person receiving the welfare becomes better off. One limiting case is utilitarianism, which is prioritarianism with a weighting of one (i.e., no matter how badly off you are, an additional unit of welfare matters equally). The other is something like a maximin principle, which is prioritarianism with an infinite weighting (i.e., no welfare gains to the better off, no matter how large they are, could justify making the worse-off worse off, even by a minute amount).⁵ We assume no one holds the view that welfare matters more if it goes to the better-off. Most prioritarians take an intermediate position, arguing that sufficiently large gains to the better-off can outweigh losses to the worse-off.⁶

2. OTSUKA AND VOORHOEVE'S ARGUMENT

Otsuka and Voorhoeve object to the priority view. Drawing on an example of Thomas Nagel's, they ask us to imagine two cases, structurally similar to the following cases.⁷

Two-Person with Certainty: There are two siblings, one who is disabled and one who is able-bodied. The disabled child would benefit from moving to the city, as he would have access to medical support; the able-bodied child would benefit from moving to the country, as she would be able to ride bikes and go on long walks.⁸ The welfare gained by the able-bodied child in moving to the country is greater than the welfare gained by the disabled child in moving to the city. You, as a morally motivated stranger, must decide what to do for the children (they must live together). This choice is represented in Table 1.⁹

5 Sen, "Rawls versus Bentham," 302.

6 Arneson, "Luck Egalitarianism and Prioritarianism"; Parfit, "Another Defence of the Priority View"; O'Neill, "Priority, Preference and Value," 332–34; Segall, "In Defense of Priority (and Equality)."

7 Nagel, *Mortal Questions*, 123–24.

8 We use children in our examples throughout this article for consistency. We recognize that children are sometimes used in the literature as examples of agents without the capacity for autonomous choice, but we assume throughout that children do have this capacity, except where it is clearly indicated (by reference to the "Absence of Subjective Information Condition," which we introduce below).

9 Cf. Otsuka and Voorhoeve, "Why It Matters That Some Are Worse Off than Others," 187.

Table 1

	Units of welfare	
	Disabled child	Able-bodied child
Send them to the city	40	50
Send them to the country	20	90

One-Person with Risk: There is one child with a particular diagnosis—she is 50 percent likely to be disabled, 50 percent likely to be able-bodied. If she is disabled, she would benefit from moving to the city, as she would have access to medical support; if she is able-bodied, she would benefit from moving to the country, as she would be able to ride bikes and go on long walks. The welfare gained by the child if she is able-bodied in moving to the country (over the able-bodied child living in the city) is greater than the welfare gained by the child if she is disabled in moving to the city (over living in the country). You, as a morally motivated stranger, must decide what to do for the child before you find out whether she will be disabled or able-bodied.¹⁰ This choice is represented in Table 2.¹¹

Table 2

	Units of welfare	
	If she is disabled	If she is able-bodied
Send her to the city	40	50
Send her to the country	20	90

Otsuka and Voorhoeve’s argument relates to the Moral Shift between *Two-Person with Certainty* and *One-Person with Risk*. The prioritarian makes no distinction between these two cases. In both, sending the child(ren) to the country results in a higher total expected welfare than the other; but sending her/them to the city results in a higher total priority-weighted welfare.¹² Therefore, the prioritarian chooses the city in both cases, and does so for exactly the same strength of reason.¹³ Otsuka and Voorhoeve, on the other hand, argue that the cases should

10 Structurally similar cases can be constructed with different units of welfare in which the same will be true for any prioritarian set of weightings.

11 Otsuka and Voorhoeve, “Why It Matters That Some Are Worse Off than Others,” 188.

12 Otsuka and Voorhoeve, “Why It Matters That Some Are Worse Off than Others,” 179–82, 188.

13 We are here considering the standard interpretation of prioritarianism, on which it is actual, not expected, welfare that receives priority weighting. For consideration of a view that attaches weight to both actual and expected welfare, see Parfit, “Another Defence of the Priority View”; Otsuka, “Prioritarianism and the Separateness of Persons,” 368, 375–80.

be treated differently. This is because the outcomes in *Two-Person with Certainty* are possible futures of two different children, whereas the outcomes in *One-Person with Risk* are both possible futures of the same child. This is a morally relevant fact for Otsuka and Voorhoeve. They believe that it is permissible to reason prudentially on behalf of the child, and “trade off” the expected welfare of her different futures, without attending to the distribution of the expected welfare between those different possible futures. This is due to the unity of the individual: it matters, they claim, that two possible futures are part of the same person’s future.¹⁴

To motivate this view, they ask us to imagine that the one-child asked, having been moved to the country and then having turned out to be disabled, “Why did you choose to send me to the country?” The morally motivated stranger could reply, “to maximize your expected welfare.” This response is not available in *Two-Person with Certainty*. In that case, if the children were moved to the country and the disabled child asked why this decision had been taken, the morally motivated stranger would have to reply, “to make someone, who was already better off than you, even better off.” Prioritarianism is insensitive to the presence of the additional prudential justification, and thus, the argument goes, insensitive to the intuitive force of the Moral Shift. Otsuka and Voorhoeve believe that you have more reason to send the one child to the country than the two children.¹⁵ They contend that prioritarianism’s insensitivity to the prudential justification (i.e., to the Moral Shift) ignores the unity of the individual, and is therefore a serious moral shortcoming.

3. THE PRIORITARIAN RESPONSE: AN APPEAL TO AUTONOMY

However, if this were all there was to the example, the prioritarian would have an easy way out. Prioritarians need not be purists: they can be (and, we would argue, should be) pluralists, holding that more than one thing matters morally.¹⁶ A defensible pluralist version of prioritarianism values both priority-weighted welfare and autonomy.¹⁷

As Gerald Dworkin notes, the concept of “autonomy” is often used by differ-

14 Otsuka and Voorhoeve, “Why It Matters That Some Are Worse Off than Others,” 180.

15 And, going further, Otsuka argues that you can simply reason prudentially on behalf of the child, without giving any prioritarian weight to benefiting her if she turns out to be disabled. See Otsuka, “Prioritarianism and the Measure of Utility,” 5.

16 O’Neill, “Priority, Preference and Value,” 343–44.

17 Otsuka and Voorhoeve acknowledge this option; see “Why It Matters That Some Are Worse Off than Others,” 186.

ent people to mean different things.¹⁸ It is therefore worth spending some time defining what we mean by autonomy, as it forms a central part of the argument to come. Otsuka and Voorhoeve define autonomy for the purposes of their discussion as “a deference to [one’s] wishes regarding choices [one] has a right to make.”¹⁹ There are at least two ways in which “deference to her wishes” could be taken. For this article, we distinguish between two kinds of autonomy, each of which is a possible interpretation of Otsuka and Voorhoeve’s definition.²⁰ One is deference to actual (self-regarding) preferences—we call this “superficial autonomy.”²¹ Superficial here does not imply meaningless: superficial autonomy can really matter. Respecting actual preferences, even when they conflict with one’s underlying values or reasons, is respecting what Joseph Raz calls the *independence principle*. Raz suggests a range of cases for which it is applicable.²² For example, if someone chooses a particular romantic partner, even if I know for a fact that her interest in that partner rests on false beliefs, and I further know that another partner best aligns with her core values, if I am in a position to choose her partner for her it seems that I should choose the partner she herself chose out of respect for her superficial autonomy.

The second kind of autonomy is what we will call “deep autonomy”—deference to one’s actual (self-regarding) values, expressed in the form of one’s ideal preferences. Ideal preferences are here understood as those that would have been autonomously formed on the basis of one’s deeply held values by an individual under conditions of full information, deliberation, and rationality (this differs from Otsuka and Voorhoeve’s conception of “ideal preferences,” as will become clear). For example, if I know that you are committed to living a healthy lifestyle, but mistakenly think that eating sugary cereal is the best way to achieve that goal, I am respecting your deep autonomy if I serve you oatmeal rather than Coco Pops for breakfast. This is because what I choose for you is what you would have chosen in ideal deliberative conditions, even if it is not what you would have actually chosen—I am deferring to your ideal preferences rather than your actu-

18 Dworkin, *The Theory and Practice of Autonomy*, 6.

19 Otsuka and Voorhoeve, “Why It Matters That Some Are Worse Off than Others,” 186.

20 Our distinction builds on a standard view in the literature on personal autonomy, according to which (roughly speaking) an agent acts autonomously when she acts as she does due to some mental state of her own regarding the matter. See Frankfurt, *The Importance of What We Care About*, 80–94, and *Necessity, Volition and Love*, 155–80; Shoemaker, “Caring, Identification, and Agency”; Jaworska, “Caring and Full Moral Standing”; Watson, “Free Agency.”

21 This is what Arneson calls “the simplest formulation of subjectivism”; see Arneson, “Autonomy and Preference Formation,” 42. We limit the discussion to self-regarding preferences and values for clarity and to align with Otsuka and Voorhoeve’s approach.

22 Raz, “The Problem of Authority,” 1014.

al preferences. This notion is rightly called “autonomy” because it still involves respecting what the individual themselves would choose: each of our value sets are different, and so we still need to know about someone’s values, or reasons, before we can choose on their behalf on the basis of deep autonomy.

Respecting both kinds of autonomy involves deferring to someone’s preferences, whether actual or ideal, and therefore—given the assumption that both kinds of preferences are self-regarding—reasoning prudentially on their behalf in one-person cases. Respecting superficial autonomy means reasoning prudentially from actual preferences, and respecting deep autonomy means reasoning prudentially from ideal preferences. It is for this reason that the pluralist prioritarian has a straightforward response to Otsuka and Voorhoeve’s example, as originally formulated: we have reason to do that which maximizes welfare in one-person cases because this is what respecting autonomy demands. In other words, the Moral Shift can be explained because autonomy can be appealed to in the one-person case, but not in the two-person case. The pluralist prioritarian could therefore agree with Otsuka and Voorhoeve that, all things considered, the child should be sent to the country, but could reject the contention that this means they must give up their prioritarianism, as it is compatible—due to their pluralism—with such an all-things-considered conclusion.

One might worry that this way around the argument commits prioritarians to always choosing for the person what they would choose for themselves in one-person cases. This would mean autonomy must be a dominant consideration for them in such cases. But this is not the case, as Otsuka and Voorhoeve recognize.²³ They acknowledge the fact that, even if the prioritarian denies that, all things considered, one should send the child to the country in *One-Person with Risk*, she must still contend with the fact that some shift in moral weighting is required when moving from two-person to one-person cases—there must be some Moral Shift. Even if it is not decisive, the fact that this prudential justification exists in one-person cases, where it does not in two-person cases, must make some difference to the moral calculus. And this can be accounted for by the pluralist prioritarian through an appeal to autonomy. Though autonomy need not always be decisive, it must always have some impact on deliberation. To the extent that there is a Moral Shift when moving from two-person to one-person cases, that weighting can always be accounted for by an appeal to autonomy, when one’s prioritarianism is pluralist. It is this subtler version of the pluralist defense that can save the prioritarian from the objection outlined above.²⁴

23 Otsuka and Voorhoeve, “Why It Matters That Some Are Worse Off than Others,” 179.

24 For similar discussions regarding a subtler version of pluralist egalitarianism, see Parfit, “Another Defence of the Priority View,” 399; Temkin, “Egalitarianism Defended,” 780.

4. THE "ABSENCE OF SUBJECTIVE INFORMATION" CONDITION

Otsuka and Voorhoeve foresee the response of appealing to autonomy, and want to isolate considerations of priority-weighted welfare from this (legitimate) pluralist position. So they assume that the child is too young to have well-informed and rational preferences regarding the outcomes.²⁵ They believe that this condition, which we will call the *Absence of Subjective Information Condition* (ASIC), rules out considerations of autonomy. They do not clearly set the condition out as a complete absence of subjective information, but we believe this is the most plausible way to interpret it: if we only have completely unreliable and ill-informed subjective preferences (whether actual or ideal), then we effectively have no subjective information with which we can deliberate. Indeed, this is precisely the point of the condition, otherwise it would not isolate prioritarian considerations from considerations of (superficial or deep) autonomy. With this condition in place, there are no preferences (whether actual or ideal) to appeal to, and so, in this modified case, "there is ... no rival autonomy-based justification" available to the prioritarian for moving to the country.²⁶ Otsuka and Voorhoeve believe that, with this condition, they can continue to appeal to the notion of the expected welfare of the child with the ASIC in place. This can be done, they argue, with reference to what they call the child's "ideally rational and self-interested preferences."²⁷ These are not ideal preferences in the sense discussed thus far, i.e., preferences that would have been formed under ideal deliberative conditions, but something else: they idealize away not only misinformation, irrationality, and the like, but also any difference in underlying values, such that we can know what a person's "ideally rational and self-interested preferences" are without any subjective information. Since the preferences are completely idealized/abstracted in this way, they do not rely on any information about the actual child, and so the notion of expected welfare that depends on them can survive the introduction of the ASIC in a way that the pluralist prioritarian's appeal to autonomy cannot. In this way, then, Otsuka and Voorhoeve believe that they can sidestep the pluralist prioritarian defense, and press the case that prioritarianism ignores the unity of the individual.

We agree that the ASIC is necessary to rule out the pluralist prioritarian's ap-

25 Otsuka and Voorhoeve, "Why It Matters That Some Are Worse Off than Others," 187. For babies, or anyone completely without rational capacities, we can still identify universally pleasurable or painful goods, but these are basic, few in number, and do not significantly affect our argument; see note 29.

26 Otsuka and Voorhoeve, "Why It Matters That Some Are Worse Off than Others," 188.

27 Otsuka, "Prioritarianism and the Measure of Utility," 15.

peal to autonomy. However, it has further implications, which mean that Otsuka and Voorhoeve's argument causes no trouble for the majority of prioritarrians: the very move that blocks the appeal to autonomy also renders their argument, in most cases, toothless. This is because there is only one way of determining welfare for the outcomes with the ASIC in place, and that is through an objective-list account of welfare.

4.1. *The ASIC Implies an "Objective-List" Account of Welfare*

Three standard types of accounts of welfare appear in the literature: hedonistic, preference-satisfaction, and objective-list accounts.²⁸ To sketch an outline of these accounts: hedonistic, or mental state, accounts hold that a person's welfare consists in their balance of pleasure over pain. Preference-satisfaction, or desire, accounts hold that a person's welfare consists in the extent to which a person's preferences are fulfilled. Finally, objective-list accounts hold that a person's welfare consists in the extent to which they possess "objective goods." These goods are such that, even if the person does not derive pleasure from them and would prefer not to have them (i.e., does not benefit from them subjectively), their welfare is improved through their possession thereof. These accounts are called "objective-list" accounts because proponents of such a view owe us, and often provide, a list of the goods that make us better off, regardless of our subjective views on the matter. So, for example, someone who is given a good education, even if they are completely indifferent to it or would have preferred not to be given it, can be said to be made better off by having the good of education on these accounts.

Each of these accounts has been met with well-known objections, and it is not the purpose of this article to evaluate their independent merits. Instead, we want to draw attention to a (rather obvious) feature of hedonistic and preference-satisfaction accounts that is not shared by objective-list accounts: they are dependent on the person's subjective information. If we know nothing about someone's preferences, values, psychological characteristics, and dispositions (i.e., if the ASIC is in place), it is normally impossible to determine whether a certain state of the world will increase or decrease their welfare, on a hedonistic account.²⁹ For preference-satisfaction accounts, the case is even clearer: without

28 Kymlicka, *Contemporary Political Philosophy*, 13–19; Parfit, *Reasons and Persons*, 493. For a contemporary modified version of this classification, see Dorsey, "The Hedonist's Dilemma." For recent criticism of the standard classification of welfare accounts, see Fletcher, "A Fresh Start for the Objective-List Theory of Well-Being"; Woodard, "Classifying Theories of Welfare." For convenience, we use the standard classification—the differences between the classifications do not affect our argument.

29 There are two exceptions to this rule, neither of which poses a threat to our argument as a

any knowledge of what someone's preferences are, we cannot know what would satisfy them. For both these accounts, then, if we are deprived of all subjective information about the person in question, we simply cannot say what impact some change would have on their welfare.³⁰

It is worth noting that the connection between autonomy and subjective information is of a different character for the two subjective kinds of accounts of welfare. For preference-satisfaction accounts, the relationship is identical: what satisfies someone's preferences (whether actual or ideal) and what they would choose for themselves are one and the same. For hedonistic accounts, the relationship is more contingent: what will give someone pleasant mental states and what they would choose for themselves normally overlap. What is more, in actual cases, there is a strong epistemic connection, as the only way to understand what effect something would have on someone's mental states is normally to infer it from information about their preferences and values. In everyday life, we assess the welfare value of outcomes by comparing them with other possible outcomes: in action, by choosing an outcome (and therefore not choosing the other possible outcomes), and in speech, by expressing our enjoyment of an outcome (and therefore not expressing our enjoyment of the other possible outcomes).³¹ For this reason, we tend to infer that a friend would prefer the park over the cinema when we know that they would enjoy the park (and we know

whole. The first is for those hedonists who believe that pleasure can be measured completely objectively, for example by measuring the levels of dopamine in the subject's bloodstream. For these hedonists, we accept that ASIC does not prevent them from ranking outcomes by their welfare-improving properties, but we take this view to be implausible and rarely held. Further, these hedonists face the epistemic problem of determining what effects going to the city or the country will have on the child's dopamine levels if they do not use subjective information as a guide. The second is for universally pleasurable goods (though it is not obvious that anything falls into this category). These goods can be distributed by the hedonist without any subjective information, but the small number of goods, if any, in this category puts strict limits on the importance of this fact. For the purposes of this essay, we treat dopamine-type views as special cases of objective-list accounts, with just one good on the list—the implausibility of this view means its inclusion in objective-list accounts does not greatly improve their attraction. And we ignore the distribution of universally pleasurable goods, since the goods under discussion in all the examples in the literature fall outside this category.

- 30 The argument to come also holds true for hybrid and pluralist accounts of welfare, which we address in note 33. For examples of hybrid accounts of well-being, see Raz, "The Role of Well-Being"; Parfit, *Reasons and Persons*, 501–2; Kagan, "Well-Being as Enjoying the Good," 253–55. For an example of a pluralist account, see Fletcher, "A Fresh Start for the Objective-List Theory of Well-Being."
- 31 Often when we say we enjoy something we actually mean that we would prefer it, given a certain set of choices in front of us.

nothing, let us assume, about how much they would enjoy the cinema)—and, if asked, we would say that we are choosing for them as we think, according to the available information, they would choose for themselves.

There are theoretical exceptions to this practical rule: most notably, Guy Fletcher's example of the ascetic, who prefers not to have pleasure because it would anger their deity, shows that hedonistic welfare and preferences can come apart in some cases (universally pleasurable goods might be another).³² But for the most part, because preferences and hedonistic welfare substantially overlap, and because hedonistic judgments rely on preference information in normal cases, the connection between hedonistic welfare and autonomy is strong enough for the purposes of this article. Because this connection holds contingently and epistemically, it is more precise to say that, for the preference-satisfaction theorist, the ASIC rules out subjective information and therefore autonomy, whereas for the hedonist it rules out the basis for both determining welfare and determining autonomy. But in both cases, the presence or absence of subjective information tracks the ability of the prioritarian to appeal to autonomy, whether necessarily or contingently, and therefore in both cases the ASIC undermines Otsuka and Voorhoeve's examples. This is not true of objective-list accounts. If education improves someone's welfare, then, regardless of their subjective experience or view of that education, they are made better off by it. The objectivity of such accounts allows their proponents to tell people, without reference to their subjective information, that something has made them better off. This is a special property of objective-list accounts: only they can tell us about someone's welfare in completely abstract conditions, when we are deprived of all subjective information about that someone. For the same reason, it is only with an objective-list account of welfare that Otsuka and Voorhoeve's notion of "ideally rational and self-interested preferences," which in turn is only required because of the ASIC, makes sense. It is a conception of preference entirely abstracted from facts about the individual who holds the preference, such that every person in a situation has the same such preference, and so it requires an objective-list account of welfare to have any content at all.

Consider again *One-Person with Risk*. With the ASIC in place, neither hedonistic nor preference-satisfaction accounts of welfare can make sense of this example, because the example relies on determinate units of welfare that could not be generated on the basis of available information. Without knowing anything about the child's attitudes toward riding bikes and going on long walks, how could the morally motivated stranger know that the country will give her a welfare of 90 if she is able-bodied, on either kind of account? On hedonistic

32 Fletcher, "A Fresh Start for the Objective-List Theory of Well-Being," 211–14.

accounts, there is no information about what will give her pleasurable mental states; on preference-satisfaction accounts, there is no information about her preferences. She could hate the outdoors and enjoy access to museums (which the city provides) much more; or love the outdoors far more than 90 suggests. Because there is no way to construct such a numerical measure of the child's welfare, when the ASIC is in effect, a proponent of either of the subjective accounts of welfare must reject the example as incoherent and need not respond to it.³³

Objective-list theorists, however, are pressed by the example. For them, it is entirely coherent to say that a world in which someone has a certain set of opportunities for outdoor activity is better for their welfare than a world in which they would have a certain other set of health states (depending, of course, on the exact constitution and weighting of their objective lists). If the child were to turn around, *post facto*, and say, "I would have liked the city much better," the objective-list theorist has the resources to respond, "Your welfare, as expressed by the satisfaction of your 'ideally rational and self-interest preferences,' is higher here, regardless of what you think or feel." For prioritarians who subscribe to an objective-list account of welfare, then, Otsuka and Voorhoeve's objection remains a telling one. So appealing to autonomy is not enough to reject Otsuka and Voorhoeve's argument altogether, but it is enough to severely curtail its scope.

4.2. Possible Responses: Ideal Preferences with the ASIC?

If Otsuka and Voorhoeve were to respond to our argument successfully, while avoiding assuming an objective-list account of welfare, they would need to show that there is some way of deriving their "ideally rational and self-interested preferences" without knowing anything about the actual preferences or values of the child, and without relying solely on objective facts about the choices. In their original article, Otsuka and Voorhoeve make an attempt in this direction, which we argue is unsuccessful. They try to ground the quantification of the expected welfare of the child in aggregated hypothetical preferences expressed by other people. As they write, in relation to a similar example:

Surveys indicate that people who imagine themselves in such a predica-

33 In the special case in which there is no subjective information about the child, hybrid accounts of welfare—which include both subjective and objective elements and therefore for which subjective information is fundamental—imply that there is insufficient information to decide which option is best for the child, and thus Otsuka and Voorhoeve's objection to them is unsuccessful. Pluralist accounts of welfare (essentially objective-list accounts in which desire satisfaction is one good) imply that there is sufficient information, as subjective information is not fundamental to their account—it is just one good on the list. Therefore, we should treat the pluralist-account theorist as a variant on the objective-list theorist, and address the same objections to her.

ment would typically be indifferent between (i) receiving the treatment that might raise one from a state of very severe impairment to this state of severe impairment and (ii) receiving the treatment that might entirely cure one of the slight impairment.³⁴

These impairments were fully described to those responding to the surveys, and—using the “standard gamble” method—indifference points were reached for the relevant changes in welfare. Otsuka and Voorhoeve rely on this fact in the construction of their later examples (on which ours are based), in order to fill out what they mean by the welfare of the various people in question. It might be thought that this process grounds the welfare of the child in such a way as to avoid the argument above, as there is no need to appeal to an objective-list account. By resting on the aggregate hypothetical preferences of other people, this view might be thought to avoid relying on the actual preferences of the child, but be subjectively determined and thus compatible with hedonistic or preference-satisfaction accounts—and so meet the objection outlined above.

There are two ways to interpret this attempt, and the distinction between what is evidence of welfare and what constitutes welfare is useful to understand them. Something can be used as evidence of welfare when it indicates that someone’s welfare has gone up, regardless of the underlying philosophical account of what welfare consists in. Mainstream economics often uses preference satisfaction in this way; the rate at which people smile might also be thought of a measure of welfare. Something is constitutive of welfare, on the other hand, when welfare consists in it—the standard debate about welfare, sketched above, is a debate about what constitutes welfare. Whether the survey data is understood as evidence of welfare or constitutive of welfare, the argument bottoms out as a poor way of understanding what would make the child better off. On the one hand, Otsuka and Voorhoeve might be using the hypothetical preferences of others as a way of approximating what the child would want: as a measure of their welfare. They might think that the fact that a representative group of people claims to be indifferent between two health states is good evidence for another person also being indifferent between those states. But this is just an appeal to autonomy, as preferences remain the constitution of well-being: it is an attempt to approximate the (actual/ideal) preferences of the child with reference to the (hypothetical) preferences of representative others—and an imperfect one, at that. If this is the way in which they seek to cash out this response, then they reopen the door for prioritariums to appeal to autonomy (and end up with a much

34 Otsuka and Voorhoeve, “Why It Matters That Some Are Worse Off than Others,” 172.

worse way of getting at the child's preferences than if they were to drop the ASIC), and therefore neutralize the response's purpose.

On the other hand, they might be arguing that the hypothetical preferences of others constitute welfare. But this is very implausible: it is hard to see why it would be the case that the aggregated views of others play a role in constituting what will make a different individual better off. This is an account of welfare that contradicts all three standard accounts. If I am a hedonist or a preference-satisfaction theorist, then what matters is how the states of health accord with the individual: the individual might have very different preferences from the group surveyed, so relying on their views will not help. But if I am an objective-list theorist, the survey results will not matter either (unless perhaps only objective-list theorists were surveyed): the situation itself tells me everything I need to know, and how people believe they would respond to it is neither here nor there. So falling back on this defense would imply an even stronger version of our original argument: that Otsuka and Voorhoeve's example has nothing to say about anyone who has taken a view on the "account of welfare" question. Deriving welfare from the hypothetical preferences of others therefore offers no support to Otsuka and Voorhoeve, whichever way it is interpreted.

4.3. *Possible Responses: A Weaker ASIC?*

Another possible response to our argument has been suggested by Otsuka in correspondence. Rather than assuming an objective-list account of welfare, he suggests the following assumption could be made:

We know facts about [the child's] psychological dispositions and character traits, including their general level of risk appetite. On the basis of this, we're able to idealize in a manner that yields determinate idealized preferences over gambles. We don't, however, know whether they would (in their non-idealized state) actually have these preferences if confronted with the particular choices among gambles. Maybe they would. Maybe they wouldn't.

In essence, this assumption is a weaker ASIC: rather than assuming away all subjective information, we simply assume away information about actual preferences, leaving enough information to determine what the child's ideal preferences (in this article's sense) would be. Determining ideal preferences in this way might be thought to make the examples coherent, while avoiding a reliance on an objective-list account of welfare.

This response is unsuccessful because the weaker version of the ASIC that it implies no longer performs the function required: it does not successfully rule

out considerations of autonomy. Although only knowing about the child's psychological dispositions and character traits, rather than their actual preferences, prevents the prioritarian from appealing to superficial autonomy, this weaker ASIC gives her precisely the information required to appeal to deep autonomy. With knowledge of the child's underlying values, assuming we are ideal deliberators, we can determine what they would choose if they were ideal deliberators (i.e., determine what their ideal preferences are). As we have said, both superficial and deep autonomy could matter. In section 3, we set out some cases in which superficial and deep autonomy come apart, and—following Raz—suggested that sometimes one should take precedence, and sometimes the other. But if we only have enough information to determine what respecting one kind of autonomy demands, then it is clear that someone who respects autonomy in general ought to respect that kind of autonomy. As an instance of this, in the one-child case, if we have enough information to determine what respecting the child's deep autonomy demands, then the pluralist prioritarian has a reason to choose in line with this: a reason directly derived from their respect for autonomy. This weaker ASIC therefore does not perform the task Otsuka and Voorhoeve demand of it, and so a stronger assumption—of an objective-list account of welfare—remains necessary to render their argument coherent. In fact, because both the determinate ideal preferences of the child and considerations of autonomy rely on the same subjective information, one will always track the other: if we have enough information to avoid assuming an objective-list account of welfare, we will have enough information to appeal to autonomy, and vice versa. However the strength of the ASIC is varied, then, it will not enable Otsuka and Voorhoeve to avoid this dilemma.

With these possible responses met, the conclusion of our argument remains undefeated. Without information about the subjective information of the child, there is no way to level the objection Otsuka and Voorhoeve describe without relying on an objective-list account of well-being to fill out their notion of “ideally rational and self-interested preferences.” If we are objective-list prioritarrians, the objection remains as compelling as before. If we are not, were we to be faced with a case like *One-Person with Risk*, we would have to respond that more information is required, that without knowing what the preferences of the child are we simply could not say what would maximize the priority-weighted utility of the child. And if more information were to be provided, then appealing to autonomy would be back on the table.

5. THE ASIC IMPLIES A RANGE OF RATIONAL RISK ATTITUDES

Section 4 narrowed the scope of Otsuka and Voorhoeve's argument to those prioritarrians who subscribe to an objective-list account of well-being; this section will argue that even among these prioritarrians, if we accept an unorthodox but plausible account of the normative relationship between risk aversion and welfare, most will have little to worry about.

"Orthodox" decision theory offers one way to represent an agent's preferences over prospects, according to which the value that a person assigns to an act can be determined by the conjunction of two functions: one that assigns utilities to its possible outcomes (a "utility function"), and one that assigns probabilities to those outcomes, conditional on each act being chosen (a "probability function"). On this view, we can construct cardinal rankings of outcomes (an ordering that puts numerical values on the utilities in each outcome) by referring to these two functions only. This idea is often formalized using the expected utility framework of von Neumann-Morgenstern.³⁵ To cardinalize preferences, in this framework, gambles are set up between the various outcomes. The gambles take the form, "You have an X percent chance of A or a Y percent chance of B . Which do you prefer?" With a sufficient number of such gambles, for various values of X and Y , the ranking can be cardinalized using the information resulting from the choices.³⁶

Choosing an option in these gambles requires, as an input, our risk aversion—if we have a very low chance of A , we might be more drawn to choosing B because we are averse to the risk of leaving empty-handed. So we cannot go through the cardinalizing process sketched above without relying on facts about our attitudes to risk. On the orthodox view, we can only come to numerical values for the utilities of outcomes if we have included risk aversion in the cardinalizing process. Risk aversion is not distinguished from any other desire attitudes. It is treated as any other preference, or, more accurately, it is taken into account in the process of the preference formation.³⁷ To be risk-averse with respect to expected utilities is to be irrational, on the orthodox view. If a proponent of Otsuka and Voorhoeve's argument holds some version of this view, we have nothing further to say in relation to risk aversion specifically—it is to be treated as just another kind of preference, and so the arguments in section 4 apply.

However, one of the major critiques of the orthodox view is that it fails to

35 Von Neumann and Morgenstern, *Theory of Games and Economic Behavior*.

36 For an elaboration on the orthodox view, see Stefánsson and Bradley, "What Is Risk Aversion?" 2; Buchak, *Risk and Rationality*, 1–2, 10–36.

37 Stefánsson and Bradley, "What Is Risk Aversion?" 2.

distinguish between preferences over goods and preferences over risk itself: risk aversion is treated as just another preference, albeit a general one that applies to all gambles.³⁸ Many argue that because of this failure, the orthodoxy does not capture the phenomenology of risk attitudes.³⁹ “Unorthodox views” in choice theory focus on this issue, and argue that risk aversion should be considered separately. Lara Buchak has suggested that a third parameter be added to the mix of the orthodox recipe: a separate function that measures the agent’s attitude to risk.⁴⁰ As a result, she proposes a “risk-weighted expected utility” theory, in which—because a risk function is included in the analysis of what matters—it is not irrational to be risk-averse with respect to expected utilities.⁴¹

We refrain from providing arguments for either of these views. Nevertheless, the distinction between orthodox and unorthodox views of risk aversion is relevant to our current endeavor, as on unorthodox views such as Buchak’s—which we believe have merit—Otsuka and Voorhoeve are forced to make a further assumption in order for their examples to work. This assumption narrows the scope of the applicability of their argument even further, as this section will argue.

On the unorthodox view, roughly speaking, in everyday life, if we are faced with a risky prospect, we take into account, in addition to our preferences for the possible outcomes, our general attitudes to risk. For example, when it comes to deciding whether to insure one’s house, two people might derive exactly the same utilities from the outcomes (equally happy if the house remains in good shape, equally annoyed if it burns down), and have the same understanding of the probability of a fire, but if the first is more risk-averse than the second, she might buy an insurance policy that the second thinks is not worth the money. Neither person would normally be condemned as irrational, because risk attitudes, like preferences, are seen as subjective. People tend to think that a variety of levels of risk aversion (defined here as how much negative weight is given to variance in outcomes) can be rational.

This fact is relevant to Otsuka and Voorhoeve’s argument. Even granting that we are objective-list prioritarrians, the ASIC means that we do not have access to any subjective information about the child, including their risk attitudes. To

38 Stefánsson and Bradley, “What Is Risk Aversion?” 2.

39 Watkins, “Towards a Unified Decision Theory”; Hansson, “Risk Aversion as a Problem of Conjoint Measurement.”

40 Buchak, *Risk and Rationality*, 49–50.

41 Other defenses of the view that it is rationally permissible to be risk averse with respect to one’s own good (whether in the form of one’s values or one’s hedonistic welfare) include Broome, *Weighing Goods*; Okasha, “On the Interpretation of Decision Theory.”

complete their examples, on the unorthodox view, they must tell us something about the risk aversion of the child—as it turns out, the only plausible assumption they could make regarding risk aversion further reduces the number of objective-list prioritarians who must respond to their argument.

We might have a very risk-averse child, Avril; we might have a child, India, who is indifferent to risk; or (less likely) we might have a daredevil, Darren, who is risk-seeking. The ASIC means Otsuka and Voorhoeve need to make an assumption about what a rational risk attitude looks like, otherwise the decision about the child's future will be lacking a vitally important piece of information. Since the examples are set up with the welfare in each of the outcomes cardinalized, the assumption needs to spell out what weightings it is rational to apply to expected welfare.

To illustrate the difference between Avril, India, and Darren, we can consider what they would choose in the following gamble.

Table 3

	Units of welfare	
	Heads	Tails
Choice A	1,000	0
Choice B	501	499

The expected value of both A and B is 500. For India, this is all there is to the gamble, so India would be indifferent between the two choices. For Avril, the expected value of B is given extra weight, because the variance between the outcomes is lower and the worst possible outcome (ending up empty-handed) is avoided. So Avril would choose B. Darren, finally, would be excited by the thrill of the gamble of A, and so would choose A.

So, what assumption about the rationality of risk attitudes do Otsuka and Voorhoeve have to make? For their examples to work, on the unorthodox view, they must assume that there is a range of rational risk attitudes: anyone whose level of risk aversion falls within this range is acting rationally, but anyone who is extremely risk-averse (or risk-seeking) is not. Indeed, Otsuka implies that he would agree with this assumption.⁴² Not only does this assumption allow the examples to work, it also chimes with common sense: someone who goes to extreme lengths to insure themselves against bad outcomes, paying thousands of pounds a month in premiums, would normally be considered irrational, fearing bad outcomes too much and being irrationally inattentive to the cost of avoiding

42 See Otsuka, "Prioritarianism and the Measure of Utility," 8, where he refers (in another context) to a risk aversion that is "low, yet not so low as to be irrationally so."

them. With this reasonable assumption, Otsuka and Voorhoeve's examples are complete: we are able to account for the fact that the child might be Avril, India, or even Darren (if any risk-seeking attitudes fall within the rational range).

This necessary assumption has consequences, however, that further weaken the overall argument against prioritarianism. First, since it introduces the notion of a rational range, there will be cases in which either the city or the country is in the rational self-interest of the child, and so both options will be permissible. How many such cases there will be depends on how widely the range is drawn, a matter that need not concern us here. In cases in which the expected welfare outcomes are sufficiently close, rational self-interest will be indeterminate. On reflection, this is how things should be: there are many choices in everyday life, like the prospect of moderately expensive house insurance, for which we might think rational self-interest is not decisive. So this does not pose a problem for Otsuka and Voorhoeve—though it means that, in such cases, their objection to the prioritarian is much weakened, as the prioritarian choice (choice B, for instance) is a permissible choice.

More troubling for their argument is the implication that, if a prioritarian's weightings are sufficiently moderate, they will always choose permissibly. Prioritarian weightings are a close analogue of risk-attitude weightings, as both are weightings applied to welfare to generate final value decisions, which (when the attitude is one of risk aversion) add weight to choices that avoid the worst outcomes. Risk aversion manifests, in cases like these, as a direct tendency to reduce the variance between outcomes, and prioritarianism, by adding more moral weight to the worse outcomes than to the better outcomes, does the same. Indeed, this analogy has been used by Derek Parfit to defend the rationality of prioritarianism.⁴³ If the prioritarian weightings are, by analogy, "within" the range of rational risk attitudes, then the prioritarian choice will always be a rational one, and Otsuka and Voorhoeve's objection to prioritarianism—that in some cases it chooses that which is not in the ideally rational self-interest of the child—will no longer hold true.

To see this, consider again *One-Person with Risk*:

Table 2

	Units of welfare	
	If she is disabled	If she is able-bodied
Send her to the city	40	50
Send her to the country	20	90

43 Parfit, "Another Defence of the Priority View," 423.

The expected value of sending her to the city is 45, and the expected value of sending her to the country is 55: this is what makes the example useful in Otsuka and Voorhoeve's argument. But the variance of sending her to the city is only 10, whereas the variance of sending her to the country is 70. If someone is risk-averse (within the rational range) they will therefore add weight to the option of going to the city, and this could mean, in this case, that the city is actually a more attractive prospect than the country. In a strongly analogous way, the prioritarian will add more weight to the lower-value outcomes, and less to the higher-value outcomes, and come to the same conclusion. Importantly, one could describe the prioritarian weightings and the risk-averse weightings using variants of the same formula: one would be over outcomes, and another would be over variances in outcomes, but each level of prioritarianism would track one and only one level of risk aversion. This means that, if the risk aversion is within the rational range, then so is the prioritarianism, and therefore the prioritarian will always choose a permissible option. Of course, in this example sending her to the country would also be rational (assuming that an indifference to risk is also within the rational range of risk attitudes), but this would be only one of the rational, or permissible, choices.

What does this imply, on the unorthodox view of risk aversion? It means that any moderate prioritarian who applies non-extreme weightings to welfare will always choose permissibly—by Otsuka and Voorhoeve's criterion of acting in the rational self-interest of the child—in these one-person cases. Of course, the notion of "moderate" is tied to the definition of the rational range of risk attitudes. If the range is small, then few prioritarians can be sure of always choosing permissibly. But as long as such a range is admitted, there will be a class of prioritarians who can safely ignore Otsuka and Voorhoeve's arguments.

Otsuka would dispute this conclusion (and in so doing, give an indication that he might be sympathetic to the unorthodox view of risk aversion). He has argued that risk aversion does not necessarily tell in favor of prioritarianism, so even if someone were being risk-averse when choosing for another they would sometimes choose the utilitarian over the prioritarian option. To demonstrate this, he gives the example of coming across a stranger in the wilderness who has just suffered an accident. He has a 50 percent chance of being disabled, and a 50 percent chance of full health. You have the option of what he calls the "risky treatment," which would make the stranger's condition less severe if he turns out to be disabled, but would leave him slightly disabled if he were to be in full health. To complete the symmetry with *One-Person with Risk*, the difference in welfare between "slightly disabled" and "full health" is greater than the difference between "disabled" and "less severely disabled." In this case, Otsuka claims, the

risk-averse person would not administer the “risky treatment.” This is because they will not be “*taking* any risk on his behalf if you leave him alone.”⁴⁴

We believe that Otsuka’s understanding of risk aversion as relying on the doing versus allowing distinction is misguided. As his emphasis shows, and as the wilderness example reveals more generally, for Otsuka the treatment is risky only insofar as you must actively choose it, whereas you merely allow the accident to take its course. However, when risk aversion is understood, as outlined above, as an aversion to variance in outcomes, the wilderness example does not tell against its identification with prioritarianism. Administering the treatment makes the least bad outcome less bad (at a cost to the value of the best outcome), and thereby reduces the variance in outcomes. When the distinction is drawn in this way, it is clear that the morally motivated stranger, truly choosing—with risk aversion—on behalf of the unconscious person, would choose the (not risky) treatment. It is possible that, independently, the additional worry about doing versus allowing would give them a separate reason to avoid the treatment; but it would not be risk aversion, properly understood, that would give them such a reason. If we are risk-averse in the sense of aversion to variance in outcomes, then a range of rational risk aversions will always make room for prioritarians, and the identification between the two views will remain consistent.

Otsuka and Voorhoeve might argue that, though the moderate prioritarian always chooses a permissible option, they do so for the wrong reason: they continue to wrongly ignore the unity of the individual, and treat the two possible futures of the person as distinct outcomes, when they should be considered as part of one life. This objection is unaffected by our argument in this section, but, in our view, the sting is taken out of it when it does not come with any impermissible implications. Further, prioritarians tend to subscribe to the view that one can coherently talk about something being good *sub specie aeternitatis* (this is often part of what it means to be a prioritarian), so an objection that calls this into question without providing further worrying counterexamples would be an external, rather than an internal, criticism.

6. CONCLUSION

These two objections, each an implication of the ASIC, work together to mean that Otsuka and Voorhoeve’s argument, and by extension the general objection to prioritarianism, has no force against the majority of pluralist prioritarians. Those who are vulnerable to the objection, and so must answer the charge of wrongly ignoring the unity of the individual, are only those who subscribe to

44 Otsuka, “Prioritarianism and the Measure of Utility,” 9.

both an objective-list account of welfare and either: (a) orthodox decision theory, or (b) an unorthodox decision theory in which the range of rational risk aversions is narrow (or, analogously, one coupled with an extreme set of prioritarian weightings). For these prioritarians, and there are likely to be some, the arguments of this article are no help: they continue to owe us a reason why the Moral Shift is not reflected in their moral calculus. For all others, the ASIC does not make sense, and so appealing to autonomy remains a decisive justification for the Moral Shift—their prioritarianism thus remains undefeated.⁴⁵

shaiagmon@gmail.com
matthitchens@live.co.uk

REFERENCES

- Arneson, Richard J. "Autonomy and Preference Formation." In *In Harm's Way: Essays in Honor of Joel Feinberg*, edited by Jules L. Coleman and Allen Buchanan, 42–75. Cambridge: Cambridge University Press, 1994.
- . "Luck Egalitarianism and Prioritarianism." *Ethics* 110, no. 2 (January 2000): 339–49.
- Broome, John. *Weighing Goods: Equality, Uncertainty and Time*. Oxford: Blackwell, 1991.
- Buchak, Lara. *Risk and Rationality*. Oxford: Oxford University Press, 2013.
- Dorsey, Dale. "The Hedonist's Dilemma." *Journal of Moral Philosophy* 8, no. 2 (2011): 173–96.
- Dworkin, Gerald. *The Theory and Practice of Autonomy*. Cambridge: Cambridge University Press, 1988.
- Fletcher, Guy. "A Fresh Start for the Objective-List Theory of Well-Being." *Utilitas* 25, no. 2 (June 2013): 206–20.
- Frankfurt, Harry G. *The Importance of What We Care About*. Cambridge: Cambridge University Press, 1988.
- . *Necessity, Volition, and Love*. Cambridge: Cambridge University Press, 1999.
- Hansson, Bengt. "Risk Aversion as a Problem of Conjoint Measurement." In *Decision, Probability, and Utility*, edited by Peter Gärdenfors and Nihl-Eric Sahlin, 136–58. Cambridge: Cambridge University Press, 1988.

⁴⁵ We would like to thank Campbell Brown, Susanne Burri, David Enoch, Ittay Nissan-Rozen, Mike Otsuka, Shlomi Segall, and Alex Voorhoeve for helpful comments on earlier versions of this article.

- Jaworska, Agnieszka. "Caring and Full Moral Standing." *Ethics* 117, no. 3 (April 2007): 460–97.
- Kagan, Shelly. "Well-Being as Enjoying the Good." *Philosophical Perspectives* 23 (2009): 253–72.
- Kymlicka, Will. *Contemporary Political Philosophy: An Introduction*. 2nd ed. Oxford: Oxford University Press, 2002.
- McCarthy, David. "Utilitarianism and Prioritarianism II." *Economics and Philosophy* 24, no. 1 (March 2008): 1–33.
- McKerlie, Dennis. "Dimensions of Equality." *Utilitas* 13, no. 3 (November 2001): 263–88.
- Nagel, Thomas. *Mortal Questions*. Cambridge: Cambridge University Press, 2012.
- Okasha, Samir. "On the Interpretation of Decision Theory." *Economics and Philosophy* 32, no. 3 (November 2016): 409–33.
- O'Neill, Martin. "Priority, Preference and Value." *Utilitas* 24, no. 3 (September 2012): 332–48.
- Otsuka, Michael. "Prioritarianism and the Measure of Utility." *Journal of Political Philosophy* 23, no. 1 (March 2015): 1–22.
- . "Prioritarianism and the Separateness of Persons." *Utilitas* 24, no. 3 (September 2012): 365–80.
- Otsuka, Michael, and Alex Voorhoeve. "Why It Matters That Some Are Worse Off than Others: An Argument against the Priority View." *Philosophy and Public Affairs* 37, no. 2 (Spring 2009): 171–99.
- Parfit, Derek. "Another Defence of the Priority View." *Utilitas* 24, no. 3 (September 2012): 399–440.
- . *Reasons and Persons*. Oxford: Oxford University Press, 1984.
- Rabinowicz, Wlodek. "Prioritarianism and Uncertainty: On the Interpersonal Addition Theorem and the Priority View." In *Exploring Practical Philosophy: From Action to Values*, edited by Dan Egonsson, Jonas Josefsson, Bjorn Petersson, and Toni Rønnow-Rasmussen, 139–65. Burlington: Ashgate, 2001.
- . "Prioritarianism for Prospects." *Utilitas* 14, no. 1 (March 2002): 2–21.
- Raz, Joseph. "The Problem of Authority: Revisiting the Service Conception." *Minnesota Law Review* 90, no. 4 (2006): 1003–44.
- . "The Role of Well-Being." *Philosophical Perspectives* 18 (2004): 269–94.
- Segall, Shlomi. "In Defense of Priority (and Equality)." *Politics, Philosophy and Economics* 14, no. 4 (November 2015): 343–64.
- Sen, Amartya. "Rawls versus Bentham: An Axiomatic Examination of the Pure Distribution Problem." *Theory and Decision* 4, no. 3–4 (February 1974): 301–9.

- Shoemaker, David W. "Caring, Identification, and Agency." *Ethics* 114, no. 1 (October 2003): 88–118.
- Stefánsson, H. Orri, and Richard Bradley. "What Is Risk Aversion?" *British Journal for the Philosophy of Science* (forthcoming).
- Temkin, Larry S. "Egalitarianism Defended." *Ethics* 113, no. 4 (July 2003): 764–82.
- Von Neumann, John, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. 6th ed. Princeton: Princeton University Press, 2004.
- Watkins, John W. N. "Towards a Unified Decision Theory: A Non-Bayesian Approach." In *Foundational Problems in the Special Sciences: Part Two*, edited by Robert E. Butts and Jaakko Hintikka, 345–79. Dordrecht: D. Reidel, 1977.
- Watson, Gary. "Free Agency." In *Agency and Answerability: Selected Essays*, 13–33. Oxford: Clarendon Press, 2004.
- Woodard, Christopher. "Classifying Theories of Welfare." *Philosophical Studies* 165, no. 3 (September 2013): 787–803.