

Person-Affecting Paretian Egalitarianism with Variable Population Size*

Bertil Tungodden and Peter Vallentyne

*Norwegian School of Economics and Business Administration, Department of Economics, Bergen, Norway;
Department of Philosophy, University of Missouri-Columbia, Columbia, Missouri.*

in *Intergenerational Equity and Sustainability*, edited by John Roemer and Kotaro Suzumura (Palgrave Publishers Ltd., 2006), ch. 6.

1. Introduction

Where there is a fixed population (i.e., who exists does not depend on what choice an agent makes), the deontic version of *anonymous Paretian egalitarianism* holds that an option is just if and only if (1) it is anonymously Pareto optimal (i.e., no feasible alternative has a permutation that is Pareto superior), and (2) it is no less equal than any other anonymously Pareto optimal option. We shall develop and discuss a version of this approach for the variable population case (i.e., where who exists does depend on what choice an agent makes). More specifically, we shall develop and discuss it in the context of a *person-affecting* framework—in which an option is just if and only if it wrongs no one according to certain plausible conditions on wronging.

2. The General Framework

We shall assume that, for any given option, there is a finite number of possible people who exist in that option. Moreover, we shall restrict our attention to cases where there is no uncertainty concerning the outcomes of choices.

To fully specify an egalitarian theory, one must specify the type of benefits that it seeks to equalize. Throughout the paper, however, we will leave open the relevant conception of benefit (resources, primary goods, brute luck well-being, etc.). References to a person being worse off than another should be understood in terms of the relevant benefits.

We shall assume, for the sake of argument, that benefits are ratio scale measurable and fully interpersonally comparable. The exact informational requirements, however, will depend on which version of Paretian egalitarianism (which we here develop) is adopted. For some versions, it will be sufficient to have ordinal measurability and comparability, combined with a norm level stating whether a life is worthwhile living or not. For other versions, a fully comparable ratio scale of benefits is needed. Since we will not focus on any specific version of Paretian egalitarianism, for generality, we assume that benefits are ratio scale measurable with full interpersonal comparability.

We shall be concerned with the assessment of the justice of alternatives, where alternatives are possible objects of choice (e.g., actions or social policies). Alternatives may have all kinds of features: they generate a certain distribution of benefits, satisfy or violate various rights, involve various intentions, and so on. In what follows, we shall assume that the only relevant information for the assessment of justice is the benefit distribution that an alternative generates. More formally, we shall assume:

Benefitism: Alternatives can be identified with (and thus their justice assessed solely on the basis of) their benefit distributions.

Benefitism is a generalization of welfarism. Although it does not assume that welfare (understood narrowly as subjective wellbeing) is all that matters, it does assume that justice supervenes on individual benefits. If two alternatives generate the same distribution of benefits, then they have the same status with respect to justice. Given Benefitism, we can identify an alternative with the benefit distribution that it generates, and in what follows we shall do so for

simplicity.

We also assume that the set of distributions generated by the set of possible alternatives is *rich* in the following sense:

Domain Richness: For any logically possible benefit distribution X, there is an alternative that generates that distribution.

This condition rules out, for example, the possibility that, where there are just three people, the distribution $\langle 3, 7, 9 \rangle$ (3 to the first person, 7 to the second, 9 to the third) is not one of the alternatives. All logically possible benefit distributions are among the alternatives. This is not to say that all are part of any given *feasible* set (the alternatives that are open to an agent on a given occasion). Of course, there are lots of logically possible benefit distributions that are not feasible on a given occasion. The claim here is about the range of benefit distributions that can be assessed by justice. The condition holds that such judgements can be made for all logically possible distributions. We believe that this is a highly plausible condition. Benefit distributions here play the role of test cases for the theory of justice. All logically possible test cases—assuming, as we shall, a finite population—are admissible.

We also impose the following assumption on the set of feasible sets.

Existence of Individually Best Feasible Option: For any given feasible set, for each individual, there is a maximum feasible benefit.

This rules out feasible sets where one or more person's benefits are unbounded (i.e., can

be greater than any standard number) and where everyone's benefits are bounded but one or more person's benefits has no maximum value (e.g., $1/2$, $3/4$, $7/8$, ...). Making sense of rational and moral choice in such cases is very difficult and we shall not attempt to do so here.

Benefitism, Domain Richness and Existence of Individually Best Feasible Option will be assumed throughout the paper, and thus we will not state these conditions explicitly when reporting the results.

Because we shall be appealing to egalitarian considerations, we need to make explicit some uncontroversial assumptions about the nature of equality. We assume:

Perfect Equality: A distribution X is more equal than a distribution Y if there is perfect equality among the existents in X and not perfect equality among the existents in Y.

Equality Weak Anonymous Contracting Extremes: A distribution X is more equal than a distribution Y, if some permutation of distribution X can be obtained from Y by (1) transferring a fixed amount of benefits from the uniquely best off person to the uniquely worst off person—but still leaving each the uniquely best off and the uniquely worst off person, respectively, and (2) making no changes in benefits to anyone else.¹

Equality Acyclicity: If, for distributions X_1, \dots, X_n , X_1 is more equal than X_2 , X_2 is more equal than X_3 , ... and X_{n-1} is more equal than X_n , then X_n is not more equal than X_1 .

These are each quite uncontroversial. Perfect Equality says, for example, that $\langle 2, 2, 2 \rangle$ is more equal than $\langle 1, 2, 2 \rangle$. Equality Weak Anonymous Contracting Extremes (which is a

weakening of the anonymous version of the well known Pigou-Dalton condition) says, for example, that $\langle 2,5,8 \rangle$ is more equal than $\langle 1,5,9 \rangle$. Equality Acyclicity is a weakened version of transitivity for equality. If X is more equal than Y, and Y is more equal than Z, it allows (unlike transitivity) that Z may be equally good as X or that the two are incomparable.

Because we do not assume that the equality relation is complete, throughout “a most equal anonymously Pareto optimal option” should be understood as “is anonymously Pareto optimal and no such option is more equal”. Thus, if there is some incompleteness in the equality relation, an option could still be judged a most equal option, even if it is not at least as equal as all other options.

Justice can be understood in *axiological* terms—what is at least as just as what (i.e., in terms of a justice ranking relation)—or in *deontic* terms—what is just (permitted by justice) relative to a set of feasible alternatives (i.e., in terms of a justice choice function). This latter approach does not attempt to provide a global ranking of alternatives. Instead, it attempts simply to determine which of any given set of feasible alternatives are just. It is well known that if a certain kind of contraction and expansion consistency is required, then the deontic approach is equivalent to the axiological approach. We doubt, however, that contraction consistency is a requirement of justice,² and hence that these two approaches are equivalent.

Justice can be understood in different ways, but we here understand it as concerned with what is *owed* to individuals in the sense of what is required to avoid wronging them. We shall thus assume:

Person-Affecting: An option is just if and only if it wrongs no one.

Person-Affecting would be a controversial thesis if it were a thesis about moral permissibility generally. It would claim that there are no *impersonal wrongs* (wrongs that wrong no one). Although one of us (Vallentyne) is inclined to defend this view, we do not here presuppose it. Instead, we simply limit our attention to justice as what we owe each other (including ourselves). So understood, Person-Affecting is simply a definition of our topic. If there are impersonal wrongs, then any account of justice so understood is an incomplete account of morality. A full account of moral permissibility would then need to deal with the further question of what things are impersonally wrong and how they should be traded-off with personal wrongs.

Nonetheless, person-affecting justice is in itself an important moral topic. A common view is that it is permissible for the state (or private citizens) to forcibly restrict the liberty of citizens only when it is necessary to prevent them from wronging others. Prevention of impersonal wrongs is deemed an insufficient justification for forcibly restricting freedom. Justice in our sense thus provides the basis for assessing the legitimacy of state restrictions of liberty. Of course, if it is not legitimate for the state to restrict a person's liberty to prevent him from wronging *herself* (e.g., suicide), then our account of justice would need to be modified by excluding wrongs to oneself. Such a modification is straightforward once one identifies who the agent is in a given choice situation. For simplicity, however, we leave this modification aside.³

We shall also make the following two assumptions, which have been insightfully developed and defended by Roberts (1998, 2002) in the context of a person-affecting framework:

Non-Existence: A person is not wronged by an option if she does not exist under that option.

Best Feasible: A person is not wronged by an option if it is a best feasible option for her.⁴

Non-Existence states that individuals are not wronged by an option if they don't exist under that option. Possible individuals, that is, have no claims to come into existence. It's worth noting here that throughout we understand existence, relative to an option, in an atemporal way. Anyone who existed in the past, exists at the time of choice, or exists in the future, if a given option is adopted, is deemed to exist under that option.

Best Feasible states that a person is not wronged by an option if it is the best feasible option for her. Our assumption of Beneficentism (the view that justice is solely concerned with the benefits people get) ensures that whether a person is wronged is determined by the distribution of benefits (as opposed to non-benefits considerations). It leaves open, however, whether a person could be wronged even by the best feasible option for her. Best Feasible rules this out. One might object that in some such cases a person might still be wronged because her best feasible option is still not good enough (e.g., not enough for a decent life, or not enough to give her what she deserves). This objection makes sense if one is concerned with *ideal* justice, that is, with what justice requires ideally, independently of practical constraints of what is possible at the time of choice. We shall, however, limit our attention to *practical* justice which takes feasibility constraints as given, and asks what should be done in that situation. So understood Best Feasible is clearly plausible.⁵

We shall also assume:

No Prohibition Dilemmas: In any choice situation, at least one option is just.

This condition would be controversial if we were concerned with ideal justice, which does not take feasibility constraints into account. We are, however, considering practical justice, which takes such constraints as given, and ask what should be done. Even from this perspective, one could argue that sometimes nothing is just because nothing is good enough. We shall here, however, limit our focus to *comparative* practical justice, according to which justice is purely a matter of comparing favorably in the relevant respects with the feasible alternatives (e.g., being at least as favorable in the relevant respect as all (or 90%) of the feasible alternatives [which is always possible], as opposed to giving everyone an adequate level of benefits [which is not always possible]). Comparative practical justice always satisfies No Prohibition Dilemmas.

Call a framework *basic person-affecting* if it imposes Person-Affecting, Non-Existence, Best Feasible, and No Prohibition Dilemmas. Our task in this paper is to develop and defend a version of Paretian egalitarianism in the context of a basic person-affecting framework. It is worth noting that Roberts (1998) also invokes a principle that gives priority to benefits to those who exist in both of two alternatives over benefits to those who exist in only one. We will address this issue later in the paper. To start with, however, we do not invoke any such assumption.

3. Fixed Population

We shall here introduce a Paretian egalitarian theory that seems promising and is fully consistent with the basic person-affecting framework in the fixed population case (where the same people exist no matter what choice is made). In the next section, we show that this theory is inconsistent with the basic person-affecting framework where there is a variable population (where who exists depends on what option is chosen), and we show how the theory can be revised so as to be

fully consistent with the framework without altering its original judgements in the fixed population case.

Before stating the egalitarian theory that we shall develop, we need to introduce some definitions. An option is *Pareto superior* to another if and only if it makes someone better off and everyone else at least as well off. An option is *Pareto optimal* if and only if no feasible option is Pareto superior. An option is a *permutation* of another option if and only if it has the same distribution of benefits except perhaps with people occupying different positions in the distribution (e.g., $\langle 2, 1 \rangle$ is a permutation of $\langle 1, 2 \rangle$). An option is *anonymously Pareto superior* to another just in case it is Pareto superior to the other or to some permutation of the other. An option is *anonymously Pareto optimal* just in case no feasible option is anonymously Pareto superior to it. Thus, for example, $\langle 3, 1 \rangle$ is anonymously Pareto superior to $\langle 1, 2 \rangle$, and if these are the only two feasible alternatives, then $\langle 3, 1 \rangle$, but not $\langle 1, 2 \rangle$ is anonymously Pareto optimal—even though $\langle 2, 1 \rangle$ is not feasible. Anonymous Pareto optimality entails Pareto optimality but not vice-versa.

Where there is a fixed population, the following theory seems fairly plausible:

Fixed Population Anonymous Paretian Egalitarianism (FP-APE): An option is just if and only if it is a most equal anonymously Pareto optimal option.

This theory holds that a certain kind of efficiency—anonymous Pareto optimality—is prior to egalitarian considerations. An outcome is just only if it is efficient in this sense. If there are several options that are efficient, then only those that are the most equal among them are just. Of course, the theory is controversial. Many would reject the relevance of equality to justice.

Some might accept its relevance, but hold that it is more limited (e.g., limit the role of equality to breaking ties in total benefits). We shall not attempt to defend this condition here. Our task is to extend this theory to the variable population case in the context of a basic person-affecting framework. (See Tungodden and Vallentyne (2005) for some general results on Paretian egalitarianism in the fixed population case).

The rest of this section records some observations that are rather straightforwardly true in the fixed population case, but which will turn out to fail in the variable population case. To start, it will be instructive to note that FP-APE can be characterized in terms of the following two conditions:

Anonymous Strong Pareto: An option X is unjust if it is not anonymously Pareto optimal.

Weak Egalitarianism Injustice: An option X is unjust if there is a feasible alternative Y that is anonymously Pareto optimal and more equal than X.

Anonymous Strong Pareto strengthens the standard Pareto efficiency requirement by further requiring that even Pareto optimal options be judged unjust if one of their permutations is not Pareto optimal. The strengthening introduces a rather uncontroversial way of solving some of the cases where there is a conflict of interest in the population. Weak Egalitarian Injustice imposes an egalitarian requirement on how to solve the remaining cases of conflicts.

We now note some observations. For brevity, let us say that a theory is *the most permissive theory consistent with a given set of conditions* just in case the theory judges just every option judged just by any other theory that is consistent with the conditions.⁶ Consider

then:

Observation 1: In the fixed population case, FP-APE is the most permissive theory of justice consistent with the conjunction of Anonymous Strong Pareto and Weak Egalitarianism Injustice.

The proof is straightforward, and hence omitted. We now note that, in the fixed population case, FP-APE is fully consistent with the basic person-affecting framework.

Observation 2: In the fixed population case, given Person-Affecting, FP-APE is consistent with the conjunction of Best Feasible, Non-Existence, and No Prohibition Dilemmas.

Observation 2 can be established as follows. Given that the result only covers a fixed population, it is trivially true that FP-APE is consistent with Non-Existence. Moreover, consider any option X that is the best feasible option for someone that exists. If X is judged just by FP-APE, then Person-Affecting implies that no one is wronged in this alternative, which is consistent with Best Feasible. If X is not judged just by FP-APE, then it is not the most equal anonymously Pareto optimal option. Hence, there is someone who is worse off in this option than in the most equal anonymously Pareto optimal option. FP-APE is consistent with a theory of wronging that states that the person who is worse off in X than in the most equal anonymously Pareto optimal option is wronged in X and the person for which X is the best feasible option is not wronged in X. This theory of wronging is consistent with Best Feasible. No Prohibition Dilemmas is satisfied because (1) there is always at least one anonymously Pareto optimal option, and (2) given that (a) (as indicated above) “most equal” is stipulatively understood as “no

option is more equal”, and (b) Equality Acyclicity holds, there is always at least one most equal anonymously Pareto optimal option.

In sum, the above results show that, for a fixed population, FP-APE is characterized by Anonymous Strong Pareto and Weak Egalitarian Injustice and is fully consistent with the basic person-affecting framework. As we shall now see, the latter is not the case when we move to the variable population case.⁷

4. Variable Population

In the variable population case, the people who exist under one option need not be the same as those who exist under another. We shall use “*” to denote non-existence. Thus, in the feasible set $\{ \langle 3, *, 2 \rangle, \langle 2, 4, * \rangle \}$, the first person exists in both options, the second person exists only in the second option, and the third person exists only in the first option.

There are several issues that need to be clarified if anonymous Paretian egalitarianism is to be applied in the variable population case. First, how do we understand equality? Second, how do we define an anonymously Pareto optimal option?

We will not impose any controversial assumptions about how to understand equality when the population size is variable. We only assume that equality is measured only among those who exist. Thus, for example, we assume that $\langle 2, 2, * \rangle$ is perfectly equal, whereas $\langle 2, 2, 0 \rangle$ is not.

With respect to the notion of anonymous Pareto optimality in the variable population case, we first need to make clear how to compare existence with non-existence. We assume that, for a given individual, (1) for any world in which she does not exist, there is some world (not necessarily accessible in a given choice situation) in which she exists that is equally good for her,

and (2) any two worlds in which she does not exist are equally good for her. The first assumption is not uncontroversial, but we believe it to be plausible. A world in which an individual receives sufficiently large benefits (e.g., a world that is full of happiness for her) is better for her than any world in which she does not exist, and that any world in which she does not exist is better for her than a world in which she receives sufficiently low negative benefits (e.g., a world full of pain and suffering for her). It is thus plausible to assume that there is some intermediate level of benefits that is equally good for her as non-existence. For greater defense, see Holtug (2001, 2005). The second assumption is plausible, since the only feature of worlds in which a person does not exist that is relevant for how good that world is for her is her non-existence. Given these two assumptions, we scale benefits so that the zero point is the level of benefits for which it is equally good to exist with those benefits than to not exist at all. Thus, we assume that $\langle 2, 1 \rangle$ is better for the second person than $\langle 2, * \rangle$, and that $\langle 2, * \rangle$ is better for her than $\langle 2, -1 \rangle$.

Given this understanding of when an individual is better off, the most natural understanding of Pareto optimality holds that $\langle *, 3 \rangle$ is not Pareto optimal when $\langle 1, 3 \rangle$ is feasible. This is because $\langle 1, 3 \rangle$ makes the first person better off and the second person no worse off. We shall understand Pareto optimality (and superiority) in this sense. We use, that is, the usual definition of Pareto optimality but combine it with the assumption that non-existence is equally valuable with existence with no benefits.

Next, how is a permutation to be understood for the definition of *anonymous* Pareto optimality? The most natural conception, which we shall adopt, simply treats * (non-existence) as one more value. Thus, $\langle 2, 3, * \rangle$ is a permutation of $\langle *, 3, 2 \rangle$, but $\langle 0, 3, 2 \rangle$ is not.

With these understandings, we can now show that, in the variable population case, FP-APE is not consistent with the basic person-affecting framework.

Observation 3: In the variable population case, given Person-Affecting, FP-APE does not satisfy the conjunction of Best Feasible and Non-Existence.

To prove the result, consider the feasible set $\langle *,3,1 \rangle$ and $\langle 2,2,* \rangle$. Both are anonymously Pareto optimal and $\langle 2,2,* \rangle$ is more equal. Hence, FP-APE judges $\langle 2,2,* \rangle$ as just and $\langle *,3,1 \rangle$ as unjust. Given Person-Affecting, this implies that someone is wronged in $\langle *,3,1 \rangle$. This, however, entails that the conjunction of Non-Existence and Best Feasible is violated. This is because Non-Existence entails that person 1 is not wronged in $\langle *,3,1 \rangle$ and Best Feasible entails that persons 2 and 3 are not wronged in $\langle *,3,1 \rangle$.

The problem, however, is not merely with FP-APE. We now note:

Observation 4: In the variable population case, Person-Affecting, Best Feasible and Non-Existence are jointly incompatible with each of Anonymous Strong Pareto and Weak Egalitarian Injustice.

The incompatibility with Weak Egalitarian Injustice is illustrated by the example given above. The incompatibility with Anonymous Strong Pareto can be seen by considering the feasible set consisting of $\langle *,1,5 \rangle$, and $\langle 5,1,5 \rangle$. Best Feasible and Non-Existence entail that no one is wronged in $\langle *,1,5 \rangle$ and Person-Affecting then entails that this option is just, which violates Anonymous Strong Pareto (since $\langle 5,1,5 \rangle$ is Pareto superior).

Thus, we need to weaken our Paretian and egalitarian conditions in order to make them compatible with the basic person-affecting framework. Call an option, X , *person-affecting*

anonymously Pareto optimal just in case there is no feasible option Y that (1) is anonymously Pareto superior to X and (2) *makes someone existing in X better off*. In the feasible set consisting of $\langle *, 1, 5 \rangle$ and $\langle 5, 1, 5 \rangle$, only the second is anonymously Pareto optimal, but both are person-affecting anonymously Pareto optimal (since no anonymously Pareto superior option makes anyone existing in $\langle *, 1, 5 \rangle$ better off). Consider, then:

Person-Affecting Anonymous Strong Pareto: An option X is unjust if it is not person-affecting anonymously Pareto optimal.

Person-Affecting Weak Egalitarianism Injustice: An option X is unjust if some person-affecting anonymously Pareto optimal option is more equal and makes someone existing in X better off.

In the fixed population case, these two conditions are equivalent to their original counterparts. In the variable population case, however, they are strictly weaker. Neither is violated in our above examples. Person-Affecting Weak Egalitarian Injustice is silent for the feasible set consisting of $\langle *, 3, 1 \rangle$ and $\langle 2, 2, * \rangle$. Although both are anonymously Pareto optimal—and hence person-affecting anonymously Pareto optimal—and $\langle 2, 2, * \rangle$ is more equal than $\langle *, 3, 1 \rangle$, the former does not make anyone existing in the latter better off. Likewise, Person-Affecting Anonymous Strong Pareto is silent for the feasible set consisting of $\langle *, 1, 5 \rangle$, and $\langle 5, 1, 5 \rangle$. Although the latter is anonymously Pareto superior to the former, it does not make anyone existing in the former better off.

Consider, then:

Person-Affecting Anonymous Paretian Egalitarianism-Version 1 (PA-APE1): An option, X, is just if and only if (1) X is a person-affecting anonymously Pareto optimal option, and (2) no other such option is more equal and makes someone existing in X better off.

This theory holds, for example, that all three options are just in the feasible set consisting of $\langle *, 3, 1, * \rangle$, $\langle 2, *, *, 3 \rangle$, and $\langle 2, 2, *, * \rangle$. Only the second is anonymously Pareto optimal, but all three are person-affecting anonymously optimal (since no other feasible option is both anonymously Pareto superior and makes someone existing in the former better off). Moreover, although the third is more equal than the other two, it does not make anyone existing in the other two better off. Hence, all three are judged just.

We now note:

Observation 5: In the variable population case, given Person-Affecting, PA-APE1 is consistent with the conjunction of Non-Existence, Best Feasible, No Prohibition Dilemmas, Person-Affecting Anonymous Strong Pareto, and Person-Affecting Weak Egalitarian Injustice.

The proof of this observation is as follows:

(1) To see that PA-APE1 is consistent with the conjunction of Non-Existence and Best Feasible, it suffices to note that PA-APE1 is compatible with a theory of wrongdoing that holds that an option X wrongs a person if and only if (a) she *exists in X*, (b) X is *not* a person-affecting anonymously Pareto optimal option, and (c) some person-affecting anonymously Pareto optimal option is more equal and makes her *better off* than X does. This theory of wrongdoing entails that

no person is wronged by an option if she does not exist under that option (Non-Existence) and that a person is not wronged by an option that is the best feasible option for her (Best Feasible).

(2) To see that PA-APE1 satisfies No Prohibition Dilemmas, it suffices to note that (given Equality Acyclicity) there is always at least one option that is person-affecting anonymously Pareto optimal option and for which no other such option is more equal and makes someone existing in X better off.

(3) Finally, PA-APE1's satisfaction of Person-Affecting Anonymous Strong Pareto and Person-Affecting Weak Egalitarian Injustice follows trivially from its definition.

Although PA-APE1 is consistent with the basic person-affecting framework, we believe that it fails to capture some of the spirit of a person-affecting approach. Consider the feasible set consisting of $\langle 3, 2, * \rangle$, $\langle *, 2, 1 \rangle$, and $\langle 1, 3, * \rangle$. Only the first is person-affecting anonymously Pareto optimal. The second option is ruled out because the third option is anonymously Pareto superior and makes the second person better off. The third option is ruled out because the first is anonymously Pareto superior and makes the first person better off. Thus, PA-APE1 judges only the first option just. Why, however, should we think that $\langle *, 2, 1 \rangle$ is unjust? Assuming that $\langle 1, 3, * \rangle$ is unjust, everyone existing in $\langle *, 2, 1 \rangle$ is at least as well off as under every *just* option (since $\langle 3, 2, * \rangle$ is the only other possibly just option). More generally, we believe that the following condition is plausible in the context of the person-affecting approach:

No Just Improvements: An option does not wrong an individual if all feasible alternatives that make her better off are unjust.

In the above example, the feasible options are $\langle 3, 2, * \rangle$, $\langle *, 2, 1 \rangle$, and $\langle 1, 3, * \rangle$. No Just Improvements says that, if $\langle 1, 3, * \rangle$ is judged—by other conditions—unjust, then $\langle *, 2, 1 \rangle$ wrongs no one. Option $\langle 1, 3, * \rangle$ is the only option that makes someone in $\langle *, 2, 1 \rangle$ better off. Thus, if the former is unjust, it is not possible to make anyone existing in $\langle *, 2, 1 \rangle$ better off except by choosing an unjust option. No Just Improvements requires that, in this case, no one is wronged by $\langle *, 2, 1 \rangle$.

No Just Improvements is similar to Best Feasible. Both say that an individual is not wronged if no “admissible” option makes her better off. Best Feasible takes all feasible options to be admissible. No Just Improvements, on the other hand, takes options to be admissible only if they are just (on the basis of other conditions). Because it takes a more restrictive view of what is admissible, No Just Improvements entails Best Feasible, but not vice-versa. In what follows, then, we shall replace Best Feasible by the stronger No Just Improvements.

We can formally note that the above example establishes:

Observation 6: In the variable population case, given Person-Affecting, PA-APE1 violates No Just Improvements.

Indeed, the problem is more general:

Observation 7: In the variable population case, given Person-Affecting, Person-Affecting Anonymous Strong Pareto and Person-Affecting Weak Egalitarian Injustice are each incompatible with No Just Improvements.

The conflict with Person-Affecting Anonymous Strong Pareto is established by the above feasible set consisting of $\langle 3, 2, * \rangle$, $\langle *, 2, 1 \rangle$, and $\langle 1, 3, * \rangle$. The conflict with Person-Affecting Weak Egalitarian Injustice can be seen by considering the feasible set consisting of $\langle 3, 2, * \rangle$, $\langle *, 1, 4 \rangle$, and $\langle *, 5, 0 \rangle$. All three are person-affecting anonymously Pareto optimal. Person-Affecting Weak Egalitarian Injustice judges the third unjust (because, by Equality Weak Anonymous Contracting Extremes, the second is more equal and makes the third person better off) and also judges the second unjust (because the first is more equal and makes the second person better off). Given Person-Affecting, however, this violates No Just Improvements, since everyone who exists in $\langle *, 5, 0 \rangle$ is at least as well off as under $\langle 3, 2, * \rangle$, which is the only other possibly just alternative.

We believe that No Just Improvements is a plausible condition on justice and we shall therefore assume it in what follows. Call a framework *expanded person-affecting* just in case it satisfies No Just Improvements (and not merely Best Feasible), as well as Person-Affecting, Non-Existence, and No Prohibition Dilemmas. Thus, we must weaken our Pareto and equality conditions even further so as to make them compatible with this expanded framework. Let us say that two options are *anonymously Pareto incomparable* just in case neither is anonymously Pareto superior to the other and neither is a permutation of the other. Consider:

Conditional Person-Affecting Anonymous Strong Pareto: If, for any options X and Y in a given feasible set, (1) option X is just, and (2) X is anonymously Pareto superior to Y and makes someone existing in Y better off, then Y is not just.

Conditional Person-Affecting Weak Egalitarian Injustice: If, for any options X and Y in a given feasible set, (1) option X is just, and (2) X is anonymously Pareto incomparable to Y, more equal than Y, and makes someone existing in Y better off, then Y is not just.

We shall show that these two conditions are jointly compatible with the expanded person-affecting framework by appealing to the following theory, which we believe to be eminently plausible. To formulate this theory concisely, we introduce the term *recursively person-affecting most equal Pareto optimal option*, which is defined as follows, where the *unresolved set* is initially the feasible set and then sequentially modified as follows:

(1a) Determine which options are most equal anonymously Pareto optimal options relative to the unresolved set. These options are judged recursively person-affecting most equal Pareto optimal options and are removed from the unresolved set.

(1b) Determine which options have at least one existing person who is worse off than under some option judged to be a recursively person-affecting most equal Pareto optimal option by the previous step. These options are judged *not* to be recursively person-affecting most equal Pareto optimal options and are removed from the unresolved set.

(2) Repeat steps (1a) and (1b) in order until the unresolved set is empty.

(3) An option is a recursively person-affecting most equal Pareto optimal option if and only if so judged by this procedure.

We propose, then:

Person-Affecting Anonymous Paretian Egalitarianism-Version 2 (PA-APE2): An option is just if and only if it is a recursively person-affecting most equal Pareto optimal option.

We shall illustrate the above definition and the resulting theory with reference to the feasible set consisting of $\langle 5,7,* \rangle$, $\langle 9,3,* \rangle$, $\langle *,9,3 \rangle$, $\langle 9,* ,2 \rangle$, and $\langle *,8,4 \rangle$. In the first round, $\langle 5,7,* \rangle$ is judged just because, given Equality Weak Anonymous Contracting Extremes, it is the most equal anonymously Pareto optimal option and $\langle 9,3,* \rangle$ is judged unjust because it makes the second person worse off than $\langle 5,7,* \rangle$. The unresolved set at this point consists of $\langle *,9,3 \rangle$, $\langle 9,* ,2 \rangle$, and $\langle *,8,4 \rangle$. In the second round, by Equality Weak Anonymous Contracting Extremes, $\langle *,8,4 \rangle$ is judged just because it is the most equal anonymously Pareto optimal option relative to the unresolved set, and $\langle *,9,3 \rangle$ and $\langle 9,* ,2 \rangle$ are judged unjust because they each make the third person worse off than under $\langle *,8,4 \rangle$. Thus, PA-APE2 judges only $\langle 5,7,* \rangle$ and $\langle *,8,4 \rangle$ just. This satisfies No Just Improvements (given Person-Affecting), since each of the other three options makes at least one person worse off than under at least one of these two just options.

We now note that PA-APE2 is consistent with the expanded person-affecting framework:

Observation 8: In the variable population case, given Person-Affecting, PA-APE2 is consistent with the conjunction of Non-Existence, No Just Improvements (and Best Feasible), No Prohibition Dilemmas, Conditional Person-Affecting Strong Pareto, and Conditional Person-Affecting Weak Egalitarian Injustice.

The proof is as follows, where Person-Affecting is assumed throughout:

- (1) To see that PA-APE2 is consistent with Non-Existence, it suffices to note that PA-APE2 is compatible with a theory of wrongdoing that holds that an option X wrongs a person if and only if
- (a) she *exists in X*, (b) X is *not* a recursively person-affecting most equal Pareto optimal option,

and (c) some recursively person-affecting anonymously Pareto optimal option is more equal and makes her *better off* than X does.

(2) To see that PA-APE2 is consistent with No Just Improvements, it suffices to note that an option is only judged unjust by PA-APE2 on the basis of a comparison with an option that is judged just.

(3) To see that PA-APE2 satisfies No Prohibition Dilemmas, it suffices to note that (given Equality Acyclicity) there is always at least one most equal person-affecting anonymously Pareto optimal option.

(4) To see that PA-APE2 satisfies Conditional Person-Affecting Strong Pareto consider any feasible set and any two options X and Y thereof, where (a) option X is *just*, and (b) X is anonymously Pareto superior to Y and makes someone existing in Y better off. Since option X is considered just, by step (1a) of the recursive procedure, X must have been the most equal anonymously Pareto optimal options relative to the unresolved set at some step of the procedure. If option Y was part of this unresolved set, then it would be ruled unjust according to step (1b). If Y was not part of this unresolved set, then it would have been judged just or unjust by a previous step in the procedure in relation to a larger unresolved set. However, option X would also have had to be part of this larger unresolved set (since it was judged just only at a later step), and thus option Y could not be an anonymously Pareto optimal option in this larger unresolved set. In sum, option Y has to be judged unjust in relation to one of the two unresolved sets, which is consistent with Conditional Person-Affecting Strong Pareto.

(5) To see that PA-APE2 satisfies Conditional Person-Affecting Weak Egalitarian Injustice, consider any feasible set and any two options X and Y, where (a) option X is *just*, and (b) X is anonymously Pareto incomparable to Y, more equal than Y, and makes someone existing in Y

better off. By exactly the same line of reasoning as above, we can show that the recursive procedure has to judge Y as unjust, which is consistent with Conditional Person-Affecting Weak Egalitarian Injustice.⁸

One might wonder whether PA-APE2 is the most permissive theory consistent with the above conditions. The following example shows that this is not so. Consider a theory, PA-APE2* that makes exactly the same judgements as PA-APE2 for all feasible sets except the one consisting of $\langle 3,3,* \rangle$, $\langle *,2,4 \rangle$, $\langle 2,*,4 \rangle$, and $\langle 4,2,* \rangle$. Here PA-APE2 judges only the first just, whereas PA-APE2*, we stipulate, judges only the last three just. Given Observation 8, and the stipulation that PA-APE2* makes the same judgements as PA-APE2 for all other feasible sets, it follows that PA-APE2* is consistent with all the conditions of the observation for all other feasible sets. Now note that, for the above feasible set, PA-APE2* is consistent with the view that the first option is unjust because it wrongs only the first person. For this set, then, the theory is consistent with No Existence, No Just Improvements, No Prohibition Dilemmas, and the two conditional conditions. So, for this set, PA-APE2* judges just some options that PA-APE2 does not. This establishes that PA-APE2 is not *the* most permissive theory consistent with those conditions.

Although PA-APE2 is not *the* most permissive theory consistent with the conditions of the above observation, it might nonetheless be *a maximally permissive* theory consistent with these conditions, where this means that no other theory that consistent with these conditions (1) judges just every option that it judges just, and (2) also judges just some option that it judges unjust. For illustration of this notion, suppose that theories T1, T2, and T3 satisfy a given set of conditions, and, relative to the feasible set $\{X,Y,Z\}$, T1 judges only X just, T2 judges only X and Y just, and T3 judges only Y and Z just. In this case, T1 is not *a maximally permissive* theory

consistent with the conditions (since, for this feasible set, the set of options judged just by T2 is a strict superset of those judged just by T1). T2 and T3 may, however, each be *a maximally permissive* theory consistent with those conditions (if they make suitable judgements for other feasible sets). Neither, however, is *the* most permissive theory consistent with the conditions, since each judges some option just that the other judges unjust in the same feasible set.

We now note:

Observation 9: In the variable population case, PA-APE2 is a maximally permissive theory of justice consistent with the conjunction of Conditional Person-Affecting Anonymous Strong Pareto and Conditional Person-Affecting Weak Egalitarianism Injustice.

Above we proved that PA-APE2 is consistent with these two conditions. Here we prove that it is a maximally permissive theory consistent with these conditions. Consider any option, X, judged to be unjust in a given step of the recursive procedure. It is judged unjust in that step if and only if (1) it is not a most equal anonymously Pareto optimal option relative to the unresolved set for that step, and (2) some such option, Y, makes someone in X better off. It follows that Y must be either anonymously Pareto superior to X (which would imply that X is not an anonymously Pareto optimal option) or anonymously Pareto incomparable and more equal (which would imply that it is not a most equal anonymously Pareto optimal option). Given that Y is judged just and makes someone in X better off, in the former case Conditional Person-Affecting Anonymous Strong Pareto requires that X be judged unjust, and in the latter case, Conditional Person-Affecting Weak Egalitarian Injustice requires that X be judged unjust. Thus, given the options judged just by PA-APE2, all theories consistent with Conditional Person-

Affecting Anonymous Strong Pareto and Conditional Person-Affecting Weak Egalitarian

Injustice must judge all the remaining options unjust. Hence, no other theory consistent with the conditions of the observation (1) judges just every option judged just by PA-APE2, and (2) also judges just some option judged unjust by PA-APE2.

One problem remains that we need to address: Should benefits to individuals who will exist no matter what choice one makes have priority over benefits to those who will exist only if certain choices are made? We now turn to this issue and related issues.

5. Gratuitous Deprivation

Consider the feasible set consisting of just $\langle 9,9,9,9,* \rangle$ and $\langle 1,1,1,1,1 \rangle$. Is the second option just? According to PA-APE2, both are just, since both are most equal anonymously Pareto optimal options. Many, however, would argue that the second option is not just on the ground that benefits to those who will exist no matter what option is chosen (the first four people in this example) have a certain kind of priority over benefits to those who exist only if certain options are chosen (the fifth person in this example). If the choice is simply between giving those who will exist no matter what very good lives, or creating an extra person with the result that everyone will have a low quality life, it seems unjust to bring the extra person into the world—at least where (1) both are most equal anonymously Pareto options and (2) everyone that exists in the world without the extra person is as well off as is feasible.

The following condition captures a version of this intuition:

Ultra Weak Gratuitous Deprivation: An individual is wronged by an option X if (1) she exists in all feasible options, and (2) there is an option, Y, such that (a) Y is a most equal anonymously

Pareto optimal option, (b) Y makes her better off, and (c) everyone who exists in Y is as well off as is feasible.

This condition holds that $\langle 1,1,1,1,1 \rangle$ is unjust relative to the feasible set consisting of $\langle 9,9,9,9,* \rangle$ and $\langle 1,1,1,1,1 \rangle$. In this case, both are most equal anonymously Pareto optimal options, but the condition requires that a certain priority be given to the benefits of those who will exist no matter what choice is made. This priority, however, is very weak. First, the condition is silent when the option that is better for the definite existents is not a most equal anonymously Pareto optimal option. For example, it is silent for the feasible set consisting of $\langle 1,* ,3 \rangle$, $\langle 4,0,* \rangle$. Here, although the second option is better for the first person (the only definite existent), it is not a most equal anonymously Pareto optimal. Second, the condition is silent, when even one person is not as well off as possible. For example, it is silent for the feasible set consisting of $\langle 1,1,1,1,1 \rangle$, $\langle 9,9,9,9,* \rangle$, and $\langle 10,0,0,0,0 \rangle$. Here, although the second option is a most equal anonymously Pareto optimal option and better for all definite existents than $\langle 1,1,1,1,1 \rangle$, the latter is not judged unjust by this condition. This is because the first person is not as well off as possible in $\langle 9,9,9,9,* \rangle$. The condition only applies when some definitely existing people are better off and makes *everyone* (definitely existing or not) as well off as is feasible.

One of us (Vallentyne) is inclined to accept Ultra Weak Gratuitous Deprivation (indeed something much stronger), but one of us (Tungodden) is inclined to reject it as an added requirement for the expanded person-affecting framework. To see why one might reject it, consider the feasible set consisting of $\langle 100,* ,* ,* ,* ,* \rangle$ and $\langle 99,99,99,99,99 \rangle$. Ultra Weak Gratuitous Deprivation holds that the first person is wronged by the second option, and thus,

given Person-Affecting, this entails that the second option is unjust. More generally, Ultra Weak Gratuitous Deprivation holds that providing even a very small benefit to just one person who definitely will exist takes priority over large benefits to many more people who exist only if certain choices are made—as long as the former is a most equal anonymously Pareto option. Many people will find that implication difficult to accept.

We shall not attempt to resolve this issue. Below we shall propose a modification to PA-APE2 if some gratuitous deprivation condition is accepted. For the record, however, we shall briefly note several ways that Ultra Weak Gratuitous Deprivation can be strengthened. Each of these is endorsed by one of us (Vallentyne) and rejected by one of us (Tungodden).

To start consider:

Weak Gratuitous Deprivation: An individual is wronged by an option X if (1) she exists in X, and (2) there is an option, Y, such that (a) Y is a most equal anonymously Pareto optimal option, (b) Y makes her better off, and (c) everyone who exists in Y is as well off as is feasible.

This is like the original condition except that it merely requires that the individual exist *in the given option* rather than that she *exist in all feasible options*. The revised condition thus does not give priority to definite existents as such. Instead, it rules out (roughly) adding people to the world when it would have been possible to add only a proper subset of them and make the members of the subset better off in a certain way. Consider, for example, the feasible set consisting of $\langle 1, 3, * \rangle$, $\langle 3, *, 1 \rangle$, and $\langle *, 1, 1 \rangle$. The original condition is silent because there are no definite existents. The revised condition, however, judges the first option unjust (because the

second option is a most equal anonymously Pareto optimal option, makes the first person, who exists in both, better off, and makes everyone as well off as feasible).

One further strengthening is to drop the requirement in (2a) that the “dominating” option be a most equal anonymously Pareto optimal option and merely require that everyone existing in both of the options be at least as well off in the “dominating” option:

Moderate Gratuitous Deprivation: An individual is wronged by an option X if (1) she exists in X, and (2) there is an option, Y, such that (a) everyone who exists in both X and Y is at least as well off in Y as in X, (b) Y makes her better off, and (c) everyone who exists in Y is as well off as is feasible.

Unlike the Weak Gratuitous Deprivation, this judges the first option unjust in the feasible set consisting of $\langle 1, *, 3 \rangle$, $\langle 4, 0, * \rangle$ —even though $\langle 4, 0, * \rangle$ is not a most equal anonymously Pareto optimal option.

A final strengthening is to replace the requirement in (2c) that everyone in Y be as well off as is feasible with the requirement that this be so for those who exist in Y but not in X:

Strong Gratuitous Deprivation: An individual is wronged by an option X if (1) she exists in X, and (2) there is an option, Y, such that (a) everyone who exists in both X and Y is at least as well off in Y as in X, (b) Y makes her better off, and (c) everyone who exists in Y but not in X is as well off as is feasible.⁹

Unlike the above conditions, this judges the first option unjust in the feasible set consisting of $\langle 1, *, 3 \rangle$, $\langle 4, 0, * \rangle$, and $\langle 5, *, 1 \rangle$. The above conditions are silent about the first option because neither the second nor the third option makes all existents as well off as feasible. The revised condition, however, judges $\langle 1, *, 3 \rangle$ unjust, because $\langle 4, 0, * \rangle$ makes the first person (the only shared existent) better off and makes the second person (the only person who exists in the second but not the first) as well off as possible. The fact that the $\langle 4, 0, * \rangle$ does not make the first person (who exists in both) as well off as possible is not deemed relevant.

Unfortunately, we cannot here resolve the issue of whether any of these conditions should be accepted. The important point is that, if we accept at least Ultra Weak Gratuitous Deprivation, then we must modify PA-APE2. As it stands, that theory says that, relative to the feasible set consisting of $\langle 9, 9, 9, 9, * \rangle$ and $\langle 1, 1, 1, 1, 1 \rangle$, both options are just (since both are most equal anonymously Pareto optimal options). Ultra Weak Gratuitous Deprivations, on the other hand, requires that $\langle 1, 1, 1, 1, 1 \rangle$ be judged unjust.

We shall suppose that, if any gratuitous deprivation condition is imposed, it will be one of the above. For brevity, let us say that a condition C on gratuitous deprivation is *admissible* just in case it is either one of the above conditions or “the empty condition” that deems that no one is wronged by imposing no requirements. We shall say that, relative to a given feasible set, an option C-gratuitously deprives a person just in case she is wronged according to C.

The most natural revision—which is our final formulation—is the following:

Person-Affecting Anonymous Paretian Egalitarianism-with no C-gratuitous deprivation

(PA-APE-C): An option is just if and only if, *relative to those feasible options that C-gratuitously deprive no one*, it is a recursively person-affecting most equal Pareto optimal option.

This is just like PA-APE2, except that prior to beginning the recursive procedure, it first eliminates options that C-gratuitously deprive someone.

PA-APE-C is fully consistent with the expanded person-affecting framework combined with No Just Improvements and any gratuitous deprivation condition C:

Observation 10: In the variable population case, given Person-Affecting, PA-APE-C is consistent with the conjunction of any admissible condition C on gratuitous deprivation, Non-Existence, No Just Improvements (and hence Best Feasible), and No Prohibition Dilemmas.

The proof is as follows, where Person-Affecting is assumed throughout:

(1) It follows straightforwardly that PA-APE-C is consistent with any condition C on gratuitous deprivation.

(2) It follows from Observation 8 that PA-APE-C is consistent with the conjunction of Non-Existence, No Just Improvements (and hence Best Feasible) and No Prohibition Dilemmas relative to the set of options that gratuitously deprive no one. We now address whether this is so relative to entire feasible set.

(3) To see that PA-APE-C is consistent with No Prohibition Dilemmas for the entire feasible set, we have to show first that for any admissible C, there will always be a non-empty set of alternatives that C-gratuitously deprives no one. Consider any admissible version of C and suppose that someone is C-gratuitously deprived in X. This means that there is an option Y where (i) this person exists in X, and (ii) there is an option, Y, such that (a) everyone who exists in both X and Y is at least as well off in Y as in X, (b) Y makes her better off, and (c) everyone

who exists in Y but not in X is as well off as is feasible. If Y gratuitously deprives no one, then we have established that at least one option gratuitously deprives no one. If Y does gratuitously deprive someone, then it must gratuitously deprive someone who exists in both X and Y (since all other individuals are as well off as feasible). Hence, there must exist some W where (a) everyone who exists in both Y and W is at least as well off in W as in Y, (b) W makes some person (who exists in both X and Y) better off than in Y, and (c) everyone who exists in W but not in Y is as well off as is feasible. Given that we have assumed Existence of Individually Best Feasible Option, it follows that there exists such an alternative W that does not gratuitously deprive anyone. Hence, given Equality Acyclicity, it follows that there exists at least one most equal person-affecting anonymously Pareto optimal option—thereby satisfying No Prohibition Dilemmas.

(4) It is obvious that PA-APE-C also is consistent with Non Existence for the entire feasible set, given that any admissible condition C on gratuitous deprivation only states that someone is wronged in an option if they exist in that option.

(5) To see that No Just Improvements is satisfied for the entire feasible set, we have to establish that for any alternative X that gratuitously deprives someone, there exists a just alternative that makes someone existing in X better off. If, for any admissible condition C on gratuitous deprivation, X C-gratuitously deprives someone, then, given (3), there is an alternative W that does not gratuitously deprive anyone and such that (i) everyone who exists in both X and W is better off in W and (ii) everyone who only exists in W is as well off as is feasible. Given PA-APE-C, if W is not just, then there exists some Z that is just and makes someone existing in W better off in Z than in W. However, those who are better off in Z than in W must be among those who exist in X, since the others in W are as well off as is feasible. This implies that there is

someone existing in X who is better off in Z (since everyone existing in X is better off in W and some of them are better off in Z) and hence that No Just Improvements is satisfied.

We now note:

Observation 11: In the variable population case, PA-APE-C is a maximally permissive theory of justice consistent with the conjunction of an admissible condition C on gratuitous deprivation, Conditional Person-Affecting Anonymous Strong Pareto, and Conditional Person-Affecting Weak Egalitarianism Injustice.

The proof is straightforward: Any theory satisfying an admissible condition C on gratuitous deprivation has to judge unjust —as does PA-APE-C—options that C-gratuitously deprive someone. Consider now the set of all remaining options. With respect to this set PA-APE-C makes the same judgements as PA-APE2. Observation 9 establishes that PA-APE2 is a maximally permissive theory consistent with Conditional Person-Affecting Anonymous Strong Pareto and Conditional Person-Affecting Weak Egalitarianism Injustice. It follows that PA-APE-C is a most permissive theory consistent with the three conditions.

We believe that PA-APE-C is the most plausible way of adapting anonymously Paretian egalitarianism to the expanded person-affecting framework. Obviously, PA-APE-C requires more scrutiny before it can be accepted with any confidence. Our task here, however, is simply to formulate and motivate a promising person-affecting version of anonymous Paretian egalitarianism. We believe that PA-APE-C is such a theory. Before concluding, we shall note how the recursive person-affecting approach used by PA-APE-C can be generalized to other theories (such as utilitarianism).

6. A Generalization: Recursively Person-Affecting Theories

A generalized form of the recursively person-affecting procedure that we invoked to define PA-APE2 can be used to make any theory consistent with the expanded person-affecting framework—as expanded to include No Just Improvements, and perhaps an admissible condition on gratuitous deprivation. Moreover, we suggest that this way of making a theory consistent with the expanded person-affecting framework is the most plausible way of doing so.

Consider any theory of justice, *T*, and any admissible condition, *C*, on gratuitous deprivation. Define *recursively person-affecting T-C* as follows, where the *unresolved set* is initially the feasible set and then sequentially modified as follows:

- (1) Judge unjust all options that *C*-gratuitously deprive someone and remove them from the unresolved set.
- (2a) Determine which options are judged *just* by *T* relative to the unresolved set. These options are judged just and are removed from the unresolved set.
- (2b) Determine which options have at least one existing person who is worse off than under some option judged to be recursively person-affecting *T*-just by the previous step. These options are judged *not* to be just and are removed from the unresolved set.
- (3) Repeat steps (2a) and (2b) in order until the unresolved set is empty.
- (4) Recursively person-affecting *T* judges an option just if and only if it is so judged by this procedure.

Consider, for example, (total) utilitarianism, saying that the minimal set of just alternatives consist of all alternatives with the greatest total utility. It is incompatible with the expanded person-affecting framework. We suggest that the most plausible modification of

utilitarianism consistent with the expanded person-affecting framework is recursively person-affecting utilitarianism. For illustration, consider the feasible set consisting of $\langle *,2,5 \rangle$, $\langle 3,3,* \rangle$, $\langle 3,*,2 \rangle$, and $\langle *,4,* \rangle$. Utilitarianism judges only the first just and, given Person-Affecting, this violates the conjunction of Non-Existence and No Just Improvements (which requires that $\langle 3,3,* \rangle$, and $\langle *,4,* \rangle$ each also be judged just, since, in each, each existing person is at least as well off as under the only just option, $\langle *,2,5 \rangle$). Recursively Person-Affecting Utilitarianism, however, satisfies both these conditions in this case. Assuming that no condition on gratuitous deprivation is imposed, then, in the first round, no judgements are made. In the second round, $\langle *,2,5 \rangle$ is judged just and $\langle 3,*,2 \rangle$ is judged unjust. In the third round, $\langle 3,3,* \rangle$ is judged just and nothing is judged unjust. In the fourth and final round, $\langle *,4,* \rangle$ is judged just. This is consistent with the expanded person-affecting framework.

Suppose now that Strong Gratuitous Deprivation is imposed. In this case, in the first round, $\langle *,2,5 \rangle$ and $\langle 3,3,* \rangle$ are judged unjust (because each gratuitously deprives the second person in comparison with $\langle *,4,* \rangle$). In the second round, $\langle 3,*,2 \rangle$ is judged just and nothing is judged unjust. In the third round, $\langle *,4,* \rangle$ is judged just and nothing is judged unjust. Thus, all only $\langle 3,*,2 \rangle$ and $\langle *,4,* \rangle$ are judged just, and inspection shows that this is consistent with the expanded person-affecting framework.

We now note the following general result:

Observation 12: For any theory of justice, T, that satisfies No Prohibition Dilemmas, and any admissible condition, C, on gratuitous deprivation, recursively person-affecting T-C is consistent with the conjunction of Person-Affecting, Non-Existence, No Just Improvements (and Best Feasible), No Prohibition Dilemmas, and condition C.

The proof follows straightforwardly from the proofs of Observation 8 and Observation 10.

The generalized recursive person-affecting procedure thus converts any theory of justice into one that satisfies the expanded person-affecting framework. We believe, moreover, that it does so in a particularly plausible manner. Thus, for example, if one is committed to utilitarianism in the fixed population case, and endorses the expanded person-affecting framework with admissible condition C on gratuitous deprivation, then, we suggest, one should endorse recursively person-affecting utilitarianism-C. We will not, however, attempt to argue for this claim here.

7. Conclusion

We have assumed the person-affecting framework, which is defined by Person-Affecting, Non-Existence, and Best Feasible. We further suggested that Best Feasible should be strengthened to No Just Improvements. Finally, we assumed that, in the fixed population case, FP-APE is correct. In the variable population case, however, FP-APE is ruled out by the person-affecting framework. More generally, Strong Pareto and Weak Egalitarianism Injustice are each ruled out. We suggested that each should be weakened in a certain way and showed that PA-APE2 is consistent with the conjunction of all these conditions.

We also discussed the issue of gratuitous deprivation, but came to no conclusion on this difficult issue. We suggested, however, that, if some admissible condition, C, of gratuitous deprivation is imposed, then PA-APE2 should simply be applied to the set of options that satisfy that condition. More exactly, we suggested that the following view is plausible (PA-APE-C): An option is just if and only if, *relative to those feasible options that C-gratuitously deprive no one,*

it is a recursively person-affecting most equal Pareto optimal option.

Finally, we suggested that, for any theory of justice, T, the most plausible way of modifying it to make it compatible with the expanded person-affecting framework—augmented by No Just Improvements, and an admissible gratuitous deprivation condition, C—is to apply a generalized version of the recursive procedure invoked by PA-APE-C to obtain recursively person-affecting T-C.

We close with a few comments on how the person-affecting framework—and thus any recursively person-affecting theory—deals with various versions of the (deontic) repugnant conclusion. Because the repugnant conclusion raises particular problems for utilitarianism (and similar aggregative theories), we shall focus on recursively person-affecting utilitarianism for illustration.

Suppose that one has the choice between (1) an option where many people have good lives, and (2) an option where those people do not exist, many more other people exist with lives just barely worth living, and the total benefits are greater. Suppose, for example, that the choice is between $\langle 9, 9, *, * \dots [20 \text{ times}] \dots *, * \rangle$ and $\langle *, *, 1, 1 \dots [20 \text{ times}], \dots 1, 1 \rangle$. (For simplicity, we use small numbers of people for illustration, but the idea can be made more striking by supposing that each number represents a billion people.) The first has two people with a total of 18 and the second has 20 different people with a total of 20. Because all individuals existing in the first option are as well off as feasible, the person-affecting framework ensures that it is judged just. Thus, a strong form of the repugnant conclusion is avoided. Justice does not *require* one to choose the option producing a highly populated but fairly bleak world. Even recursively person-affecting utilitarianism agrees with this judgement: It judges both just.

The person-affecting framework, however, is subject to a weak version of the repugnant

conclusion in cases such as the above. The framework—that is, the conjunction of Person-Affecting, Best Feasible (or with the stronger No Just Improvements), and No Existence—requires that, in the above case, justice *allow* one to choose the option producing highly populated but fairly bleak world. This is because, in this particular kind of case, everyone in that world is as well off as is feasible. The judgement that it is just to choose such an option (even if it is also just not to do so) will strike many as bizarre. Within the person-affecting framework, however, it is inevitable and natural. Who is wronged by such a choice? Not the individuals who exist with the bleak, but worth living, lives. Their lives are better than non-existence, which is the only alternative. Nor are individuals who do not exist wronged. Hence, no one is wronged and the option is indeed just. Of course, it might be impersonally wrong to choose such an option, but we have set aside that issue in this paper.

Let us now consider a repugnant conclusion case in which some people exist under more than one option. Suppose, for example, that the choice is between $\langle 9, 9, *, * \dots [20 \text{ times}] \dots *, * \rangle$ and $\langle 1, 1, \dots [22 \text{ times}], \dots 1, 1 \rangle$. As above, the first option has two people with a total of 18. Because both are still as well off as feasible, the person-affecting framework rightly requires that the first option be judged just. This time, however, the second option has 22 people with a total of 22 *and two of these people also exist under the first option*. It is thus no longer true that everyone in the second option is as well off as feasible, and the person-affecting framework no longer requires that the second option be judged just. Nonetheless, both PA-APE2 and recursively person-affecting utilitarianism judge the second option just. Again, this avoids the strong version of the repugnant conclusion (since the second option is not required by justice), but it faces the weak version thereof (since the second option is permitted by justice). If, however, we further add at least Ultra Weak Gratuitous Deprivation (which requires that the

second option be judged unjust), then even the weak form of the repugnant conclusion is avoided in these kinds of cases. Although we have left open whether Ultra Weak Gratuitous Deprivation should be endorsed, it is clear that it provides an important way of avoiding certain versions of the repugnant conclusion.

In sum, the person-affecting framework, we believe, has the resources to avoid the main problematic versions of the repugnant conclusion. Because we find some version of anonymous Paretian egalitarianism attractive, we have focused on it. We believe that PA-APE3-C is the most plausible version thereof that is compatible with the expanded person-affecting framework. Our more general claim, however, is that the generic recursively person-affecting procedure is a plausible way of converting any theory into one that is consistent with the expanded person-affecting framework. Obviously, many of the judgements invoked in the paper are controversial. We hope nonetheless that we have at least established that the person-affecting framework should be taken seriously and that there are promising ways of developing anonymous Paretian egalitarianism—and other theories—within this framework.

References

- Arrhenius, Gustaf (2005) 'The Paradoxes of Future Generations and Normative Theory', in J. Ryberg and T. Tännsjö (eds.) *The Repugnant Conclusion* (Dordrecht: Springer), pp. 201-218.
- Blackorby, Charles, Walter Bossert, and David Donaldson (2005) *Population Issues in Social-Choice Theory, Welfare Economics and Ethics* (Cambridge, MA: Cambridge University Press).
- Broome, John (2004) *Weighing Lives* (Oxford, Oxford University Press).
- Holtug, Nils (2005) 'Person-affecting Moralities', in Jesper Ryberg and Torbjorn Tannsjo (eds.), *The Repugnant Conclusion* (Dordrecht: Springer), pp. 129-161.
- Holtug, Nils (2001) 'On the Value of Coming into Existence', *Journal of Ethics*, vol 5, pp. 361-384.
- Roberts, Melinda (1998) *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law* (New York: Rowman & Littlefield Publishers).
- Roberts, Melinda (2002) 'A New Way of Doing the Best That We Can: Person-based Consequentialism and the Equality Problem', *Ethics*, vol. 112, pp. 315-350.
- Sen, Amartya (1993) 'Internal Consistency of Choice', *Econometrica*, vol. 61, pp. 495-521.
Reprinted in Amartya Sen, *Rationality and Freedom* (Cambridge, MA: Harvard University Press, 2002).
- Tungodden, Bertil and Peter Vallentyne (2005) 'On the Possibility of Paretian Egalitarianism', *Journal of Philosophy* 102 (2005): 126-54.
- Peter Vallentyne, 'Equality, Efficiency, and Priority for the Worse Off', *Economics and Philosophy* 16 (2000a): 1-19.

Peter Vallentyne, 'Critical Notice of *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law* by Melinda Roberts', *Nous* 34 (2000b): 634-47.

Notes

* For helpful comments, we thank Walter Bossert, Nils Holtug, Wlodek Rabinowicz, Melinda Roberts, the participants at the International Economics Association Roundtable Meeting on Intergenerational Equity, Hakone, Japan (2005), and the participants at the Philosophical Aspects of Social Choice and Welfare conference, Caen, France (2005).

¹ This is an anonymous version of a condition introduced and discussed by Vallentyne (2000a).

² More specifically, we deny that Alpha is required: If an alternative is judged just relative to a given feasible set, then it is also judged just from any subset containing it. For criticism of this condition, see Tungodden and Vallentyne (2005) and Sen (1993).

³ The person-affecting idea can also be expressed in terms of axiological justice: A distribution is less just than another only if it is worse for someone. We are, however, skeptical that the axiological person-affecting approach is promising in the variable population case. Any such approach, we believe, will have to be radically incomplete so as to avoid generating cycles of betterness (i.e., where X_1 is better than X_2 , which is better than X_3 , ... which is better than X_n , which is better than X_1).

⁴ As stated, Best Feasible overlaps with Non-Existence when non-existence is a best feasible option for an individual (e.g., both say that person two is not wronged by $\langle 2, * \rangle$ when the only alternative is $\langle 2, -3 \rangle$). To avoid this overlap, we could have restricted Best Feasible to only cover options where a person exists, but we have not done so since this would require cumbersome expressions below.

⁵ Note that Best Feasible is compatible with holding that a person is wronged by being created

with a life not worth living, when non-existence is feasible.

⁶ We here assume that theories of justice are identical if, for all possible feasible sets, they judge the same options just. Thus, there is only one theory that is a maximally permissive theory consistent with given conditions.

⁷ The literature on variable population ethics is extensive. See, for example, Arrhenius (2005), Blackorby, Bossert and Donaldson (2005), Broome (2004), Holtug (2005), Roberts (1998, 2002), and Vallentyne (2000b).

⁸ For the record, we note that PA-APE2 violates the weak anonymity condition that requires that, if an option and a permutation thereof are each feasible, then either both are just or neither is. To see that violation, consider the feasible set consisting of $\langle 3, 3, * \rangle$, $\langle 3, 1, * \rangle$, and $\langle *, 3, 1 \rangle$. PA-APE2 judges only the first and third option just. The violation of anonymity is effectively unavoidable within the person-affecting framework.

⁹ This principle is tentatively endorsed by Roberts (1998) in n. 48 of ch. 2. It is a strengthening of her official principle D^* , which is the same except that clause (4) says that no one exists in Y but not in X.