A quantitative model for a theory of justice. Part I: Derivation and implications of the quasi-maximin principle

Stijn Bruers

Abstract
In this article we start with a unified mathematical formulation of three major consequentialist theories of justice: strict egalitarianism, maximin and utilitarianism. These theories can be represented in a continuum of theories that look at the quality of life (generalized well-being) of a set of individuals. Using the Rawlsian argument of the veil of ignorance and assuming a high but not maximum level of risk aversion, we arrive at a form of prioritarianism we will call quasi-maximin theory (QMM), which lies between maximin and utilitarianism. QMM-theory is also shown to unify both needs for equity and efficiency. We briefly overview the application of QMM-theory to different economical and political issues. QMM-prioritarianism is distinguished from absolute prioritarianism, which is shown to lie between an extreme form of sufficiantarianism and utilitarianism. We demonstrate that QMM is a better principle than absolute prioritarianism, that it can incorporate sufficiantarianism, and that it avoids problems of intransitivity in decision theory. Further problems of intergenerational justice and population ethics (e.g. the repugnant conclusion) will be dealt with in more detail.

Introduction: why a quantitative model?

In this article we want to unify different theories of justice and equality, by putting them into a coherent framework. In order to do this, we will try to use mathematical modeling as much as possible. Economists and natural scientists are familiar with the use of mathematical models. In moral philosophy however, only a few theories of justice (e.g. utilitarianism) have some more or less explicit reference to quantitative objects (e.g. utility,…).
Using a mathematical framework will help us to see different theories of justice and their mutual relationships in more clarity. Mathematical modeling offers an efficient toolbox that helps us to work towards a more unified theory of justice. The mathematical equations in this article are therefore not to be taken too literally, but they should be used as ways to simplify expressions of complex ideas. What we will attempt to do, is to combine different theories (utilitarianism, maximin, sufficientarianism, prioritarianism, egalitarianism,…) into a single mathematical expression which contains some parameters. These parameters can take different values, and for specific values we get a specific theory of justice.
In summary, we derive a mathematical formulation of a principle which w will call quasi-maximin (QMM) or positional prioritarianism. We will give two arguments for QMM-prioritarianism, one is based on a Rawlsian argument of impartiality (the veil of ignorance), whereby we assume that the person in the original position has a high but not maximum risk aversion. The second argument is based on compassion for the worst-off individuals, combined with a low but non-zero need for efficiency. Hence, efficiency is inversely related to risk aversion.
After giving a few political implications of QMM-theory, we will show that QMM-prioritarianism should be distinguished from absolute prioritarianism, which is given almost all attention in the literature (made popular by Parfit, 1991). We will demonstrate that our QMM-prioritarianism (unlike absolute prioritarianism, maximin or utilitarianism), is able to deal with problems of intransitivity (Temkin, 1987), non-identity (Parfot, 1984), substitutability, population ethics, the repugnant conclusion (Parfit, 1984), asymmetry in procreational duties (Narveson, 1967) and Allais paradox (Allais, 1953).
This article will focus on the study of consequentialist theories. In a follow-up article, more ethical principles (deontological principles and ethics of care) and ideas (the existence of moral intuitions and emotions) arising from studies in moral psychology will be incorporated.

The quality of life

The starting point of our theory of justice, is the notion of quality of life. We follow the psychological approach that equate quality of life with the satisfaction of universal needs (Maslow, 1943; Rosenberg, 2003). Universal needs not only contain physical well-being (food, water, movement, rest, health, safety,…) but also social needs (connection, compassion, acceptance, warmth, contribution,…), play (joy, humor,…), autonomy (freedom, space, independence, spontaneity,…), honesty (authenticity, integrity, trust…), peace (equality, harmony, order, beauty,…) and meaning (learning, growth, challenge, efficiency, clarity, creativity, purpose,…) (see Rosenberg 2003 for a more extensive list of universal needs). As Rosenberg points out, these universal needs should be distinguished from particular strategies (e.g. wanting a piece of chocolate as a strategy to satisfy the universal needs for food, pleasure,…).
Needs can have different intensities (e.g. a little hunger vs. being very hungry) and satisfactions also can have different levels (e.g. having access to a little bit vs. a lot of food). The higher the level of satisfaction and the higher the intensity of the corresponding need, the higher the quality of life. Feelings are indicators to see when needs are met or unmet.
Some beings can also have more needs than others (e.g. extra social needs for social beings). This means that it is possible that the potential total quality of life can be higher, the more needs someone has. In other words, a being with more needs can reach a higher overall quality of life (compared to a being with only a few needs) if all of his needs are satisfied. Also the potential total quality of life can be lower when a being has more needs. As a simplified example, suppose we have a being with only one need. His quality of life arising from that need can vary from e.g. -1 (needs far from being satisfied, so this being rather prefers to die), to 0 (needs satisfied to some extent so that for this being it doesn't matter if he lives or dies), to +1 (needs highly satisfied). A being with two needs (both of the same intensity), however, can have a quality of life ranging from -2 to +2, if the qualities of life of the individual needs can be added. This latter addition property is not a necessary condition for our theory we will discuss, and is only meant for didactical purposes. Here we want to address the possibility that different beings can have different potential total qualities of life.
Furthermore, as all sentient beings have subjective experiences of their feelings and needs, all sentient beings have a well-being or a quality of life for themselves. The model we are about to discuss therefore applies to all sentient beings. All reference to "society" and "persons" or "individuals" should therefore also include mentally disabled humans and animals. We should not restrict this theory of justice to only rational, self-conscious beings.
Let us now look into more detail what quality of life involves.
First, quality of life is more than the hedonist position that only looks at pain and pleasure. Philosophical problems like the 'experience machine' (Nozick, 1974) can be avoided. Suppose we have an experience machine that can give you feelings of pleasure for the rest of your life, by plugging your brains into this machine. However, the experiences in this machine are related to a world that is not real, and you might have a strong need for authenticity that will not be satisfied by this machine, and that is why you will be reluctant to step into this machine.
Second, quality of life is more than the welfarist (utilitarian) position that looks at the total well-being (utility, good) over the course of one's life. Take for example the argument of the long living oyster (Robert Crisp, 2008). Which life would you prefer: the life of a normal human being with life expectancy 80 years, or the life of an oyster with a life expectancy you may choose, but with a very small but positive and constant well-being? In short, the human being has a high well-being for a short period of time, the oyster has a low well-being, but integrating this low well-being over the very long course of its life, the total (integrated) well-being of the oyster can be higher than that of the human. Yet, a lot of people would prefer being born as the human, no matter how long the life expectancy of the oyster. This means

that these people value the quality of life of the human higher than that of the oyster. Why is that? Perhaps because they expect that leading a human life is less boring, and they have a need for variation or growth. These needs cannot be satisfied in the life of the oyster. Perhaps the oyster does not have those needs, but that still means that a human who has these extra needs and who has satisfied those needs, has a higher quality of life. (In this article, for simplicity we will often use the terms quality of life and well-being as synonyms.)

Quality of life is also more than the resourcist position (Dworkin, 1981) that looks at economical goods that can be distributed. Quality of life is more than the compensationist position of desert-principles of justice (Dick, 1975; Lamont 1997; Milne 1986; Sadurski 1985), which focuses at the compensations of virtuous work. And quality of life is more than the sufficiantarianist position of the capabilities approach (Nussbaum 1992, 2000; Sen 1992), which looks at basic functionings that one is free to choose to improve one's flourishing. Values like economic resources, income, wealth, jobs, compensations, capabilities, desire satisfaction,… all contribute to the quality of life, but cannot be reduced or set equal to the quality of life.

The continuum of consequentialist theories of justice

The consequentialist theories that will be discussed here, all focus at the qualities of life of a set of individuals. Suppose there are N individuals, and we have to make a choice between two options X and Y. Each individual has a corresponding quality of life (a lifetime expectation of well-being) that will be influenced by the choice X or Y. Therefore, situations X and Y can be represented as

$$X = (x_1; x_2, \dots x_N); Y = (y_1; y_2; \dots y_N),$$

Where x_i is the quality of life of person i in situation X.[1]

As a starting point, we assume that the values x_i and y_i for the same person i are ordinal numbers, which means that they can be ordered in a complete well-ordered set. In other words, it is meaningful to say that e.g. $x_i > y_i$, even though these values cannot be quantified. The order relation is complete if for all x_i and y_i we have $x_i > y_i$, $x_i < y_i$ or $x_i \approx y_i$. This assumption is not a strong assumption: in nearly all our choices we can compare our different needs and feelings affecting our quality of life. We might prefer visiting a friend over reading a book, we might prefer short term satisfaction of one need over long term satisfaction of another need,… So we are able to compare the qualities of life of different choices.

A much more difficult assumption, is the following step: there is an ordinality relationship between different individuals. I.e.: we are able to compare x_1 with x_2. This is the central most difficult (or vulnerable) point in our theory of justice: how to compare the qualities of life of different individuals? Is my satisfaction of visiting a friend higher that your satisfaction of reading a book? There is no clear method to solve these kind of questions. All we have are two heuristic methods: empathy and the Rawlsian thought experiment of the veil of ignorance (Rawls, 1971; 2001). We have to try to put ourselves into the position of the other, by using

---

[1] More general, in case of uncertainty, x_i is the expectation value of the quality of life. There is some tricky point, however. Consider two situations X and Y, with two persons A and B. In both situations, there is a probability distribution. Suppose in situation X, persons A and B always get outcome 1. But in situation Y, there is a probability ½ that A gets well-being 4 and B gets 0, and a probability ½ that A gets 0 and B gets 4. If we take the expectation values for persons A and B, we get expected well-being 1 in situation X and 2 (=½.0+½.4) in situation Y, for both persons A and B. So situation Y looks better, although we are sure that in Y the outcomes will always be unequal. A better alternative that we propose, is not to take the expectation value for a person, but take the expectation value for the lowest qualities of life, and similarly for the second to lowest etc. So, in situation Y we don't calculate x_A =½.0+½.4 for the well being of person A, but x_1 = ½.0+½.0=0 for the lowest quality of life, and x_2 =½.4+½.4=4 for the second to lowest quality of life. See also Rabinowicz (2002) and McCarthy (2003).

empathy, or by imagining that we could be the other person, with all his needs and feelings. The 'emotional' method of empathy and the 'rational' method of the veil of ignorance are rules of thumb to make educated guesses about the order of the qualities of life of different individuals.

A third assumption we have to make, seems also to be a big leap into superficiality: the values x_i and y_j are assumed to be cardinal numbers, i.e. quantitative numbers that can be multiplied, added, subtracted,… Although this step might seem superficial, it is in fact only for didactical purposes that we assume cardinality, because now we can use clear mathematical expressions. Therefore, we will speak of a quantitative "model" for a theory of justice.

Finally, we only take those individuals into account that have different qualities of life in situations X and Y. If a person always has the same quality of life in all possible situations, then we should not take him up in the calculations.

With the above assumptions, we will now discuss three consequentialist theories of justice: strict egalitarianism, maximin and utilitarianism. To simplify things, we first suppose there are only two individuals involved, and we have to choose between two situations X and Y. Second, we make an important assumption that all individuals are equal and that we are impartial with respect to the outcomes of individuals. This means that there is a symmetry in the quality of life of the different individuals. In other words: the situation $X = (x_1; x_2)$ is morally equal to the situation $X' = (x_2; x_1)$. So it doesn't matter if person 1 has quality of life x_1 and person 2 has quality of life x_2, or vice versa (person 1 has quality of life x_2, person 2 has x_1). This equality (impartiality or symmetry) is an important assumption. It means that we are always allowed to write X and Y as ordered sets, whereby $x_1 \leq x_2$ and $y_1 \leq y_2$. Note that with this symmetry, the index 1 in situation Y might refer to another person than the index 1 in situation X.

The three consequentialist theories of justice can now be presented in a simple set of mathematical inequalities:

- Strict egalitarianism: X is better than Y if and only if
$$x_2 - x_1 \leq y_2 - y_1$$
i.e. the difference between the qualities of life should be minimized.
- Maximin: X is better than Y if and only if
$$x_1 \geq y_1$$
i.e. the quality of life of the person in the worst position should be maximized.
- Utilitarianism: X is better than Y if and only if
$$x_1 + x_2 \geq y_1 + y_2$$
i.e. the total quality of life (total utility) should be maximized.

Interestingly, these mathematical expressions can be written in one inequality
$$x_1 + Qx_2 \geq y_1 + Qy_2$$
where the parameter Q takes the values:

- $Q = -1$: strict egalitarianism,
- $Q = 0$: maximin,
- $Q = +1$: utilitarianism

We can write $W(X) = x_1 + Qx_2$ as the 'moral weight' of situation X, and then we have:

X is better than Y if and only if W(X) > W(Y).

One can easily generalize this formulation to situations with N number of individuals, by using a recursive relation. The moral weight of situation X then reads:
$$W\big(X = (x_1; \dots x_N)\big) = x_1 + Q\big(x_2 + Q(x_3 + \cdots)\big) = \sum_{i=1}^{N} Q^{i-1}x_i.$$
In this formulation, we only have to take up the individuals whose qualities of life can change. Suppose that only the well-beings of persons i and j can change. Then the calculation to

optimize the distribution of well-being between persons i and j will contain a factor Q^{i-j}. So this factor strongly depends on the number of individuals between i and j. But these individuals have fixed well-beings that should not influence the results. Therefore it is better not to include them in the mathematical formulation, as it will strongly influence the distribution of well-being between i and j. In other words: if there are individuals who will have the same well-beings in situations X and Y, then these individuals should not be taken up in the expression of the moral weight (we refer to appendix 4 to elaborate more on this).
We now arrived at a continuum of theories of justice, with at the two extremes: strict egalitarianism ($Q = -1$) and utilitarianism ($Q = +1$), and maximin in the middle. Formulating the consequentialist theories of justice in this way gives us a clear view of the one unsolved problem: what value of Q to choose? There is always some arbitrariness by choosing Q. Why should the utilitarians be right by picking $Q = +1$ and not 0 or -1? And our moral intuitions make it even more arbitrary, by not having to decide between only the values $+1$, 0 or -1, but by a value somewhere between $+1$ and 0. This is what we will show next.

The quasi-maximin principle as a prioritarian theory of justice

The question we now have to ask, is whether there is a preferred value for the parameter Q, that best fits our moral intuitions. We follow the argument of the original position (the veil of ignorance) of John Rawls. Suppose that you don't know which one of the N persons you will be. However, you do know the laws of nature and psychology, and hence you do know the qualities of life of the N persons in situation X and Y. Given these two ordered sets $X = (x_1; \dots x_N)$ and $Y = (y_1; \dots y_N)$, and given the fact that you don't know who you will be, which situation would you prefer?
Rawls himself concluded that the maximin principle is the most suitable theory of justice. For Rawls, Q should be equal to 0. The utilitarians like Bentham, on the other hand, preferred $Q = 1$. Performing the thought experiment of the veil of ignorance, both the maximin and the utilitarian principles are valid. The crucial difference is the amount of risk aversion we have. In appendix 1, we show that the parameter Q is related to the amount of risk aversion. $Q = 0$ implies maximum risk aversion, $Q = 1$ implies risk neutrality. From behind the veil of ignorance, we don't know which one of the N individuals we will be. It's a game of chance whereby we try to maximize the quality of life of the person we will be. By symmetry/impartiality, we can assume that the probability to be person 1 equals the probability to be any other person. So we have a probability 1/N to become person 1. (However, note that in the original argument of Rawls, someone from behind the veil of ignorance should also have no knowledge of the probabilities to be a specific person. There is not a situation of risk, but of complete uncertainty when even probabilities are not known. See also Angner, 2004. We suppose here that the equal probability distribution is known from behind the veil of ignorance, as it is the most impartial probability distribution.)
Someone with a risk neutral attitude should prefer utilitarianism ($Q = 1$), as in utilitarianism, the total quality of life, and hence the expectation value of the quality of life, is maximized. To give an example: suppose $X = (10; 10)$ and $Y = (0; 100)$. So in situation X, both persons will have an equal quality of life, and the expectation value is 10. In situation Y, the first person has no quality of life, and the second is very happy, with a quality of 100. The expectation value in this situation equals ½.0+½.100=50.
However, a lot of people would not risk having a very low (zero) quality of life as in situation Y. They would prefer situation X, even if it has a lower expectation value, because in this situation they at least have quality 10. Only risk neutral people would prefer situation Y. The Rawlsian maximin strategy, on the other hand, implies a maximal risk aversion. People with

maximal risk aversion are real pessimists and always think as if they will become the person in the worst position. They ask the question: what if I would be the worst-off person?
In reality, not many people have maximum risk aversion. People with a high but not maximal risk aversion, would prefer a positive but small value of Q. This is in line with the often heard critique on the maximin principle, that it is not good to decrease the quality of life of the second person (the better off person), in order to increase the quality of life of the first person (the person who is worst off) by a negligible amount. This is the efficiency argument. To give an example, suppose situation $X = (1; 100)$, and situation $Y = (1{,}1; 2)$. In situation Y, the richer person (person 2) gives most of his wealth and resources to person 1 who is in extreme poverty, and this poor person uses these resources very inefficiently. So he still remains in extreme poverty. In other words, in situation Y, person 2 is himself driven into extreme poverty, in order to increase the well-being of person 1 with a tiny 10%. Someone with maximum risk aversion would prefer situation Y. However, it is expected that most people's intuitions would still prefer situation X (even though in X there is a high inequality). If situation X is preferred over Y, than we know that $Q \geq \frac{1{,}1-1}{100-2} = 0{,}001$.
In more general terms, the efficiency understood in this context, is the benefit/cost ratio. From the inequality $x_1 + Qx_2 \geq y_1 + Qy_2$, with $y_2 - x_2 \leq 0$ the cost (loss of well-being) of person 2 and $y_1 - x_1 \geq 0$ the benefit of person 1, we get

$$\frac{y_1 - x_1}{x_2 - y_2} \leq Q.$$

This is the efficiency constraint. In other words, when the efficiency (the left hand side of the above inequality) is too low, i.e. less then Q, we should *not* opt for the maximin distribution, because the efficiency is deemed too low. And the smaller Q, the less importance is given to efficiency (because the stronger the efficiency constraint will be). When the efficiency is higher than Q, we should implement the action. Hence the parameter Q serves as a lower bound for the efficiency.
So we finally arrive at the quasi-maximin (QMM) principle of justice, which corresponds with a positive but small value for Q, and which can be derived from an efficiency argument (assuming we have a non-zero need for efficiency, i.e. we don't think efficiency is completely unimportant), and also from the argument of the veil of ignorance (assuming one has a high but not maximum risk aversion). The importance of efficiency inversely relates to the level of risk aversion: the higher the risk aversion, the lower the efficiency is deemed important.
The quasi-maximin principle says that we should maximize the quality of life of the person in the worst position, except if the increase of his quality of life is negligible compared to the decrease of the quality of life of the other wealthier persons. QMM-theory is compatible with compassion for the worst-off people, assuming we have a non-zero need for efficiency, and also compatible with impartiality, assuming we have a non-maximum risk aversion.

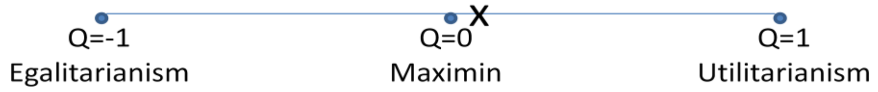

Figuur 1: the continuum of theories of justice, with the position of the quasi-maximin theory indicated with an X

The QMM-principle is a kind of prioritarianism, which gives priority to helping the worst off people (Scheffler 1982, Weirich 1983, Parfit 1997). The moral importance of a benefit for an individual is greater if the increase in quality of life for that individual is greater and if the quality of life of that individual (when she does not receive the benefit) is (relatively) lower. Quality of life, weighted by a priority function, is called "weighted well-being" (or weighted quality of life). The priority function depends on the level of well-being an individual has. In our model formulation, we can look at the benefit for individual i. Suppose in situation X she has the benefit, whereas in situation Y she does not have the benefit. Looking at the above moral weight functions $W(X)$ and $W(Y)$, we can derive that the weighted well-being B_i of person at level i reads:

$$B_i = Q^{i-1}(x_i - y_i).$$

The factor Q^{i-1} reflects the priority function, and $(x_i - y_i)$ is the increase in well-being. Prioritarianism now states that the sum of all weighted well-beings of all individuals should be maximized. So it combines in a single principle the value of efficiency (in terms of maximization of well-being) and the value of equality (in terms of giving priority to the worst off person). As is clear, for each value of Q, we have another kind of prioritarianism. Maximin and utilitarianism are the two extremes in the continuum of prioritarian ethics. Quasi-maximin, with a low positive value of Q, gives more importance to the value of equality than to efficiency.

One final remark, before we move to some applications. In the above description, the factor Q was interpreted as a constant. But the level of risk aversion might change when the difference between x_1 and x_2 changes. When the difference is very high, i.e. when x_1 is relatively very low, then the risk aversion might become higher. Also, it might be more important to give priority to the worst-off person, when the difference between the worst-off and better-off persons is large. So the factor Q might actually be a function inversely related to the difference between the well-beings of two consecutive individuals. For example: $Q(x_1; x_2) = \frac{q}{q+(x_2-x_1)}$, which is decreasing and equals 1 when $x_2 = x_1$. The moral weight of situation X could then be written again as a recursive relation:

$$W(X = (x_1; \dots x_N)) = x_1 + Q(x_1; x_2)\big(x_2 + Q(x_2; x_3)(x_3 + \cdots)\big).$$

Here, the function $Q(x_1; x_2)$ might decrease when x_1 decreases or $x_2 - x_1$ increases. There are of course other possibilities for the moral weight whereby the factor Q becomes a series of functions. Furthermore, Q might also depend on the number of individuals N involved. We will not try to determine possible expressions for Q that are in agreement with our intuitions about just distributions, risk aversions, efficiency, etc…

Applications of the quasi-maximin theory

Although the comparison between qualities of life of different individuals in different situations is very difficult, we can derive a set of approximate rules that can move us closer to the QMM-distribution of quality of life.

In his theory of justice, John Rawls derived three principles (Rawls, 1971, 2001):

1) equality of basic liberties and rights

2) equality of fair opportunity

3) equality of economic goods in terms of the difference principle: the distribution of economic goods should be according to maximin. That means that economic inequalities should be in the greatest benefit of the least advantaged persons.

First note that this latter Rawlsian difference principle refers to economic goods and not to the total quality of life. Economic goods (income, resources, wealth,…) only constitute of a subset of factors that contribute to quality of life. The QMM-theory as described in this article is more in line with the welfare based principles (like utilitarianism), and hence also incorporates the distribution of liberties, opportunities, capabilities and all other factors that contribute to the quality of life.

Hence, the above three Rawlsian principles can be easily be reframed in the QMM-theory. Let’s first look at the basic liberties and rights. In a follow-up article, we will encounter a Kantian right that cannot be derived from the QMM-theory or the original position, because it not necessarily affect the quality of life. So as for the first Rawlsian principle, we only look at rights and liberties that clearly affect the quality of life. Take for example the right to free speech. If I have a need for sharing ideas, I will feel frustrated when I do not have the right to free speech, and this obstruction will lower my quality of life. However, there are some speech acts (hate speech, insults,….) that can lower the quality of life of other people (the receivers). In most cases, allowing these disdainful speech acts will violate the QMM-principle. First, as Rosenberg notes (Rosenberg, 2003), someone uttering disdainful speech acts often implies that this person has unmet needs. Insults are a tragic expression of a person with an unmet need. If your boss insults you by saying that you are lazy, most likely that means that your boss feels frustrated and has an unmet need for e.g. efficiency, and that he only found a tragic way to express himself. Also hate speech and scapegoats indicate some unmet need (e.g. for social security, respect,…).

Let's try to apply our QMM-model to this problem. As a starting point, we have two persons. In situation X, there is no free speech. By lack of further details, and by the symmetry between the persons, we have to assume that a priori (all else equal) both persons have equal quality of life. Say $X = (1; 1)$. This is an important assumption in dealing with these kind of problems. In situation Y, person 2 insults person 1, by use of free speech. The quality of life of person 2 increases, but for person 1 it decreases. So $Y = (0{,}9; 1{,}1)$. Situation Y violates the QMM-principle. To summarize: not all speech acts satisfy the QMM-principle.

Moving to the second Rawlsian principle, when there is a scarcity of social, economic or political positions (education, jobs, elections,…), the equality of fair opportunity (and participation) can be derived from the original position. Only someone who is more talented, motivated, trustworthy or experienced to do a job that is beneficial to the least advantaged persons (or more generally a socially beneficial job that helps approaching the QMM-distribution), should have a higher prospect to get that job.

The third Rawlsian principle (the difference principle) can also easily be restated in the QMM-framework, whereby now the economic goods should be distributed according to quasi-maximin. This also means that for example disabled persons should get relatively more economic goods in order to compensate for their loss of well-being, except if the transfer of economic goods to these disabled persons cannot be done in a sufficiently efficient way. In other words, when we are only capable of increasing the well-being of the disabled person by a negligible amount by transferring huge amounts of resources to these disabled persons; we should not opt for the transfer.

So far for Rawls' difference principle. Let us also take a look at the resource-based and the desert-based principles.

In the resource-based principles of justice (Dworkin, 1981), one is concerned about the importance of personal responsibility. According to the QMM-theory, society should not keep on pouring resources down the drain, when worst-off people act very irresponsibly with these given resources (when they negligently squander them) or are highly inefficient in transforming these resources into quality of life (see Cohen 1989, Arneson 1989, Roemer 1996).

In the desert-based principles, one wants to emphasize someone's contribution (Miller 1976, Riley 1989), effort (Sadurski 1985, Milne 1986) or costs incurred in work (Dick 1975, Lamont 1997). According to the desert principle, we should distribute economic goods corresponding to the virtue or deservingness of a person (see e.g. Kagan 1999). If we now interpret virtuous work as work that contributes to the society, and more specifically that promotes the QMM-distribution of well-being, than things become clear. The more someone contributes to QMM, the more she should be rewarded. And the more her quality of life decreases by doing that important work (e.g. by doing hard, long, boring, dangerous,… work), the more she should be compensated for her loss of well-being. So the more her quality of life decreases and the more her work contributes to QMM, the more virtuous and deserving she is.

Someone who contributes more to the well-being of the worst-off persons, should get prior access to more economic goods.
In a desert-based theory of justice, one often adds the "greater gap principle". The greater the gap between what someone deserves and what someone has, the more priority should be given to decreasing that gap. In a sense, this is a generalization of prioritarianism as defined above, where priority should now be given to the most deserving person. The more deserving person is not always the worst-off person, but can also be the more virtuous person. So not only the weighted well-being matters (as in simple prioritarianism), but also someone's contribution to society (to approach the QMM-distribution, i.e. to contribute to the total weighted well-being of all people in society) should be rewarded. And this latter reward is only possible when it is not in conflict with the QMM-theory itself.
Let's give an example. Suppose we have an ill person, a doctor and a very rich and wealthy person. The quality of life distribution in situation X equals (1;10;100). Now, the rich person can give money to the doctor so that the doctor is motivated to heal the ill person. We now get situation $Y = (10; 30; 30)$. Note that the increase in well-being of the doctor (30 – 10 = 20) can be larger than the increase of well-being of the ill person (which equals 9 in the example. We might compare this desert principle with a negative feedback mechanism which acts as an attractor mechanism when there are disturbances away from the QMM-equilibrium position. This feedback mechanism actively pulls towards the state of QMM, by stimulating (rewarding) people who contribute to the total weighted well being of society.
To summarize, we see that the QMM-theory combines and encompasses a lot of different ideas: prioritarianism (keeping the balance between Rawlsian maximin and Benthamian utilitarianism) and desert-based, welfare-based and resource-based theories. There is however one important set of theories of justice, called sufficientarianism, that seems to be a competitor with our QMM-theory. This will be discussed in the next section.

Absolute versus positional prioritarianism
The above quasi-maximin (QMM) prioritarianism is an intermediate between maximin and utilitarianism. This QMM-prioritarianism can also be termed positional prioritarianism. With this we mean that the moral weight of the well-being of person i equals $Q^{i-1}x_i$, so it depends on the position of person i relative to the other individuals. This also means that a benefit that falls at a person can have different moral values, depending whether or not there are other individuals with lower qualities of life. Take for example a situation with two persons, with well-beings x_1 and x_2 ($> x_1$). A benefit of size b for person 2 will result in an increase in moral weight of Qb. But if person 1 was not present and x_2 remains the same, the same benefit bestowed to the same person will now give a higher increase in moral weight, namely b.
So the lower the position of a person relative to others, the higher the moral weight of a benefit, and hence the higher the priority. But in the literature, prioritarianism is mostly understood as 'absolute prioritarianism' (Parfit, 1991), and this should not be confused with our positional prioritarianism. In absolute prioritarianism, only absolute levels of well-being are important, and it doesn't matter whether or not there are other individuals present. The moral weight of a benefit only depends on the absolute level of well-being of the person.
This absolute prioritarianism has a mathematical expression for the moral weight that differs from our QMM-model.

$$W_{AP}\big(X = (x_1; \dots x_N)\big) = \sum_{i=1}^{N} f(x_i),$$

where $f(x)$ is an increasing concave function, called the priority weighted utility function (Broome 1991; Brown, 2007; Holtug 2006; Rabinowicz 2002; McCarthy 2003, 2008). An example is a root function $f(x) = \sqrt[n]{x}$. The steepness of the curve decreases, which means

that for a given benefit b the moral weight W_{AP} is increased the most if b falls at the person with the lowest well-being. (In other words: The increment $f(x+b)-f(x)$ is higher, the lower x.)
However, there is a problem with this formulation, as some scale invariance might be broken (see Brown, 2007). Especially for negative values of well-being, the formulation of absolute prioritarianism becomes more tricky. Brown (2007) therefore suggested what he called a 'positive prioritarian' weighted utility function, which has a stretched s-shape. For example:

$$f_{PP}(x) = \begin{cases} \sqrt[n]{x} & \text{for } x \geq 0, \\ -\sqrt[n]{-x} & \text{for } x < 0. \end{cases}$$

Note that this function is inversely prioritarian for negative values of well-being (i.e. a convex function), and positively prioritarian for positive values off well-being (i.e. a concave function).
If the parameter n in the above weighted utility function equals 1, then we get utilitarianism. Another interesting limit is when n moves to infinity. In this limit, the positive prioritarian weighted utility function becomes a threshold function: it flips from -1 to +1 at $x=0$. This is nothing but a special form of sharp sufficientarianism, where $x=0$ is the sufficiency threshold level of well-being (the word "sharp" refers to the sharpness of the threshold: there is no smooth function at the threshold, but a discontinuous jump).
Sufficientarianism states that when the well-being of a person is above some threshold level, a minimally decent life for that person is fulfilled. Differences in qualities of life between different persons are not significant when all their levels are above this threshold. Lifting the qualities of life of persons above this threshold, is therefore considered sufficient for this theory of justice. So the sufficientarian objective is that as many individuals as possible should have a quality of life above some threshold.
In summary, there are at least three extremal options for consequentialist theories of justice (i.e. theories that wants to improve the qualities of life of individuals):
1) The average quality of life of all persons should be as high as possible. This is the Benthamian/Millsian utilitarianist position
2) The quality of life of the lowest person should be as high as possible. This is the Rawlsian maximin position.
3) The number of people with qualities of life above a threshold value should be as high as possible. This is the sufficientarian position.
Now we have two forms of prioritarianism. 1) a QMM-prioritarianism that lies in between a Rawlsian maximin and a Benthamian utilitarianism, and 2) an absolute prioritarianism that lies in between a "sharp zero-threshold" sufficientarianism and utilitarianism. So QMM-theory unifies the first two choices, absolute prioritarianism unifies the first and third options. The following figure gives a schematic representation of the different theories.

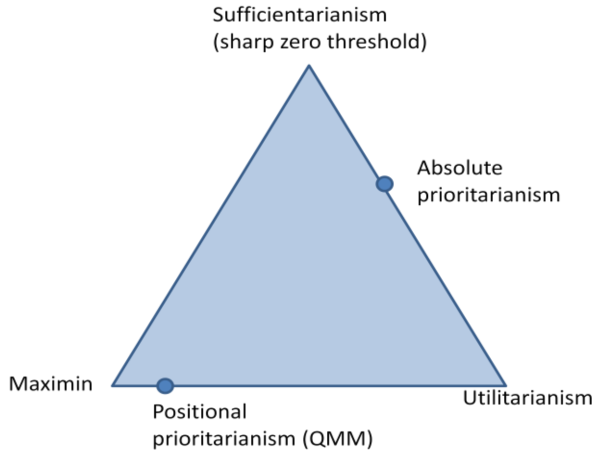

There are an infinite ways of how to unify QMM and absolute prioritarianism, so to move closer to the center of the above triangle. One obvious possibility is taking the new moral weight to be:

$$W_{QMM+AP}(X) = \sum_{i=1}^{N} Q^{i-1} f(x_i).$$

We will not elaborate on what this unification implies, as it may not be necessary to unify those two prioritarian theories. Instead, we will discuss some further properties, strengths and weaknesses of these two types of prioritarianism, and we will argue that QMM-theory has some advantages over absolute prioritarianism.

1) Parameter freedom

In QMM-theory, there is only one unknown parameter Q, whereas absolute prioritarianism has a whole function $f(x)$ that needs to be specified. The latter has the advantage that you can more freely play around and look for a utility function that corresponds well with our intuitions, but it also has the disadvantage of indeterminacy.

2) Thresholds

Note that the QMM-theory with moral weight $W_{QMM}(X) = \sum_{i=1}^{N} Q^{i-1} x_i$ is scale and translation invariant (i.e. it is invariant with respect to affine transformations). This means that we can rescale all $x_i \to s x_i$ or shift all $x_i \to x_i + t$, without changing the QMM-theory. The advantage of scale invariance is that we are allowed to be less specific in the determination of the qualities of life. The absolute prioritarian weight function W_{AP} on the other hand is not scale invariant. This also means that in QMM there is no preferred reference value for x_i that can act as a threshold like in sufficientarianism. If we want to include a threshold in the QMM formulation, we would need to break this scaling and translation invariance, for example by turning the parameter Q into a function of the qualities of life. This is not that farfetched, as it is likely that the amount of risk aversion may depend on the stakes, i.e. on the values x_i. When the stakes are higher (when the qualities of life of the worst-off persons are lower), the risk aversion might turn out to be higher, and $Q(X)$ will be lower. We will not elaborate further on this approach of breaking invariances here, and conclude that the absolute prioritarian view has a threshold built-in more naturally.

In the above formulation, the priority weighted utility function f_{PP} has a built-in threshold at $x = 0$. In the limit of $n \to \infty$, this threshold becomes very sharp, and we get a strict (sharp) sufficientarianism. But this strict sufficientarianism has some counter-intuitive problems. For example this sharp threshold allows us to decrease the qualities of life of persons below the threshold, in order to increase the quality of life of one person just a tiny amount, from e.g.

$x = -0{,}001$ to $x = 0{,}001$. Further, inequalities are not so important when everyone has a quality of life higher than the threshold value. And below this threshold it doesn't really matter how bad one's situation is. These things are not in line with what a rational person from behind the veil of ignorance would prefer.

So we would like to avoid such a sharp threshold, and a high but finite value of the root n serves exactly as a smoothening of the utility function. Absolute prioritarianism therefore avoids a counter-intuitive sharpness problem of strict sufficientarianism, just as positional prioritarianism (QMM) avoided a counter-intuitive inefficiency problem of strict maximin.

Yet, there is a serious problem with this threshold in absolute prioritarianism: it is at the level zero. This is considered way too low. A quality of life equal to say 0,01 means that this life is only just worth living. A quality of life equal to a negative value means that you'd be so depressed that you'd prefer not to be born. But sufficientarianism should do more than just lifting people above well-being 0. Because all persons – even the very poor – who are not willing to commit suicide are already above this threshold. So this zero-threshold sufficientarianism is already satisfied to a high degree, even in situations with extreme poverty.

So we can assume that the threshold value should be at some positive value T. It is one of the major difficulties of sufficientarianism to decide what counts as a sufficient decent life. Absolute prioritarianism doesn't avoid this problem, and it has matters even worse: Shifting the weight function to the right, i.e. $f(x) \to f(x - T)$ turns us in another contra-intuitive situation where we have inverse prioritarianism even for some positive values of well-being (i.e. when x is between 0 and T).

There might be another, better way in which the sufficientarian position leaks into our QMM theory of justice. In the above formulation, we took the function $f(x)$ as a function of the quality of life. But it might be that the quality of life x_i itself is a function of a set of resources or capabilities, and that the sufficiency requirement is at the level of resources.

The capabilities approach (Nussbaum 1992, 2000; Sen 1992) is probably the most influential sufficientarian theory of justice. The theory states that when all basic (important) capabilities of a person are above some threshold level, a minimally decent life for that person is fulfilled.

Also the resourcist perspective can be translated in a kind of sufficientarianism as well, if we apply the psychological theory of diminishing marginal utility at the level of resources. In other words, we can suppose that the well-being x_i is an increasing function of resources or capabilities. And this function is concave for high values of x_i. So let's take a set of resources that we will represent with a cardinal number r_i. Then it might be possible that e.g. :

$$x_i(r_i) = \begin{cases} b_i \sqrt[k]{r_i - T} & \text{for } r_i \geq T, \\ -b_i \sqrt[k]{T - r_i} & \text{for } r_i < T. \end{cases}$$

with T a threshold for the basic resources/capabilities, k a 'sharpness' parameter, and b_i can be a function that depends on other non-resourcist factors that contribute to the quality of life of person i. (I leave the question open whether the parameters T and k can just like b_i and r_i now also differ between different individuals.) This seems to be a more natural way of how sufficientarianism can be included in our QMM-theory with moral weight $W_{QMM}(X) = \sum_{i=1}^{N} Q^{i-1} x_i(r_i)$.

We therefore believe that we don't need the absolute prioritarianism formulation at all: on the level of qualities of life, we can use the QMM-formulation, and on the level of resources or capabilities we can use a sufficientarian formulation.

3) The intransitivity problem

Temkin (1987) gave an argument that seems to violate a most fundamental axiom in ethical decision theory, namely transitivity. This axiom states that if situation A is better than B, and

B is better than C, than A should be better than C. Temkins intransitivity argument goes as follows.

Temkin first refers to what he calls the First Standard View (FSV). Most people judge that a situation in which someone suffers a lot is better than a situation where two people suffer a lot, but a little bit less than the one person in the first situation. So take situation $A = (-100; 0; \dots 0)$ and situation $B = (-99; -99; 0; \dots 0)$, then A is better than B. Two extreme sufferers is better than one extreme sufferer. (In this example, the first person has a negative well-being and the other people have well-being 0. But the argument is translation invariant, so one could add a constant well-being to every person to lift every person above zero, without affecting the argument. So situation A could as well be $A = (1; 101; 101; \dots 101)$.)

We can again state that by induction, situation $C = (-98; -98; -98; 0; \dots 0)$ is even worse than situation B, situation D (with 4 people in the worst position) is worse than C, and so forth. However, most people also accept a Second Standard View (SSV), which says that situation A is far worse than a situation where any number of people suffer almost nothing. So take situation $Z = (-1; -1; \dots; -1)$, then this is better than situation A, no matter how many people are involved. This shows that transitivity is violated, because $A > B > C \dots > Z > A$.

Utilitarianism is compatible with the first standard view, but not with the second. If the number of people N in situation Z is high enough, then the sum of N times -1 can be lower than -100, which means that Z becomes worse than A. Even if a billion people had a well-being of -0,01, then this is worse than situation A.

However, maximin is compatible with the second standard view, but not with the first. Maximin says that B is better than A, because in B, the well-being of the worst-off individual is a little higher than in A.

But what about our quasi-maximin principle? As we can see in appendix 2, QMM-theory satisfies both FSV and SSV, and yet the intransitivity problem can be avoided! The appendix also demonstrates that absolute prioritarianism doesn't have this property, as it behaves like utilitarianism by only respecting FSV. Only in its extreme version of sharp sufficientarianism, can the transitivity problem be avoided. But this sufficientarianism had other problems, as we have seen.

Intergenerational justice and the repugnant conclusion

The QMM-theory not only involves current living sentient beings, but all future generations as well: From behind the veil of ignorance, you don't know in which time period you will be born. However, the number of people living in the future is not necessarily fixed. Until now, we have only looked at situations with a fixed number of people N. We can generalize the QMM-theory to situations where the number of people varies. As current generations can have influence on the number of people that will exist in the future, these situations are encountered in problems of population ethics and intergenerational justice.

Let us start with the problem of the Repugnant Conclusion (Parfit, 1984). The argument goes as follows. Start with a world with one very happy person. So, situation $A = (a_1)$, with the quality of life $a_1 = 100$. Now, what would happen if we increased this population, with another very happy (but slightly less happier) person, i.e.: situation $B = (100; 98)$. The utilitarianist would say that this situation is better: the total number of happy persons increased, and the total happiness increased. So B is better than A. This is the first step. The second step is just equalizing the qualities of life: situation $C = (99; 99)$ is considered to be even better than B, because now there is more equality. One can repeat step 1, by introducing a third person with quality of life say 97. After repeating step 1 and step 2, we move to the optimal situation Z which contains an almost infinite number of people with an almost zero quality of life. So Z is better than A, or in other words, we should boost population growth, even if all our qualities of live becomes very low. But this conclusion is repugnant.

How to solve this problem? We note that even working with the QMM-formulation with moral weights $W(X) = \sum_{i=1}^{N} Q^{i-1} x_i$, and Q positive, results in the same repugnant conclusion. There is however a simple way out, if we would take the moral weights

$$W(X) = \frac{\sum_{i=1}^{N} Q^{i-1} x_i}{\sum_{i=1}^{N} Q^{i-1}}$$

This is a weighted average of qualities of life, weighted with the polynomial powers of the parameter Q. Using this formulation, we see that situation $B = (100; 98)$ is in fact worse than situation $A = (100)$. Why is this so? Does this follow from the veil of ignorance? It does: In situation A you have probability 1 to be a person with quality of life 100. But in situation B, you have probability 1/2 to be a less happier person. Even if the difference in well-being is small, you would not (as a rational being striving for maximal quality of life) prefer situation B.

To summarize, situation X (with N persons) is better than situation Y (with M persons), if and only if

$$\frac{\sum_{i=1}^{N} Q^{i-1} x_i}{\sum_{i=1}^{N} Q^{i-1}} \geq \frac{\sum_{j=1}^{M} Q^{j-1} y_j}{\sum_{j=1}^{M} Q^{j-1}}.$$

In appendix 3 we describe another well-known problem of leveling down and welfare diffusion, which seems similar to the repugnant conclusion, but it has different implications.

Two remarks are in order. First, in this QMM-formulation, we didn't say anything about the identity of the different persons. It doesn't matter if person 1 in situation X equals person 1 in situation Y. So Parfit's Non-Identity Problem (Parfit, 1984) does not pose itself here. The Non-Identity Problem asks questions like: Can we harm a person if the other option we had would be that this person would not exist? Or can we harm a future person by bringing that person into existence? The reason why the non-identity problem is avoided, is because the QMM-theory is not about harming someone, but about just distributions of quality of life. With the above QMM-formulation, we simply can avoid the tricky questions raised by Parfit.

As a second remark, the numbers N and M refer to the number of individuals who are alive in the world-histories of situations X and Y respectively. So they also include the individuals that will exist in the future. But they do not include the individuals whose well-beings are never affected by a choice between X and Y. So if a person lives in both situation X and Y, and it has the same well-being in both situations, then we choose not to take him into account in the above inequality expression. However, see appendix 4 for further implications of this choice, especially in its relation with Allais paradox.

With this new formulation, let's look at two other problems.

1) The problem of replaceability.

One might object that in our QMM-theory (as in other consequentialist or utilitarian theories), persons are nothing but receptacles of well-being. It is not a moral problem if a person is simply replaced by another person with the same level of expected well-being. For example, you are allowed to kill someone as long as you let another child be born, which will have the same expectation of quality of life as the murdered person would have if he was not killed.

One can counter this objection by simply stating that persons have intrinsic value, which simply means here that these persons cannot be replaced without violating something deemed important. However, the problem of replaceability is already solved in the QMM-theory. Take situation X=(100), i.e. in situation X, there is a person with total quality of life equal to 100. In situation Y, you kill this person somewhere in the middle of his life, and you let a second person be born which will have the same quality of life equal to 100. So situation $Y = (50; 100)$. According to utilitarianism, the second situation is better, as it has a 50% increase in total well-being. But applying the QMM-theory, we see that if Q<1/2, then $W(X) = 100 > W(Y) = 50 + Q.100$. So the person should not be killed and replaced.

2) The problem of asymmetry in procreational duties

This asymmetry (Narveson, 1967; Mulgan, 2006) says that we do not have an obligation to give birth to children (out of the interests of those children), but we do have an obligation *not* to give birth to children when we know that the lives of those children will be miserable. Think about the problem of parents knowing they have a genetic defect which means that their potential child will be seriously handicapped. Why is there this asymmetry? Where does it come from?

Suppose we have a set of individuals, with qualities of life $X = (x_1; x_2; \ldots x_N)$. Now, suppose that a new child is born, which will have quality of life x', and suppose that the presence of this child does not affect the qualities of life of the other persons. The question is: what is the minimal level for x', in order to decide that this new situation (with the added child) is better than the old situation (without the child)? I.e., when is $X' = (x_1; \ldots x_N; x')$ better than X? Solving $W(X') \geq W(X)$, we get:

$$x' \geq \frac{\sum_{i=1}^{N} Q^{N-i} x_i}{\sum_{i=1}^{N} Q^{N-i}} = x_{min}(x_1; \ldots x_N)$$

So, in other words, if the parents know that the handicapped child will have a well-being lower than x_{min}, than it is better not to procreate, because procreation would result in a less just world. The above inequality for x' also follows from the veil of ignorance with high risk aversion. You would not prefer the situation X' when $x' \leq x_{min}$, as in this situation you have a higher probability to a worse outcome. If you would have maximal risk aversion ($Q = 0$), then $x' \geq x_{min} = x_N$. So you would only prefer X' if the new person has a quality of life higher than the person with the highest quality of life. In summary: we have an obligation to restrict our procreation, to give birth only to children with an expected quality of life $x' \geq x_{min}$. This solves the asymmetry problem.

One might object this line of reasoning, with the following example. Suppose the parents know that their potential child will have to wear glasses. Suppose wearing glasses lowers the quality of life a little bit. So we have the situation X with two happy parents and no child: $X = (100; 100)$, and the situation with the less happy child: $Y = (99; 100; 100)$. The child has a quality of life equal to say 99. According to QMM-theory, the parents should not give birth to this child. They have to abstain from procreation, because $W(X) > W(Y)$. This is a hard conclusion. But this objection is most likely not valid. We should not forget that the parents can have a need for procreation, and they will have a lower quality of life when they are forbidden to procreate. In other words, situation X might actually be: $X = (80; 80)$, and then it becomes clear that $W(X) < W(Y)$.

A new inequality measure

QMM-theory allows us to derive a new inequality metric. In order to see this, we first note that QMM is a combination of a utilitarian and an egalitarian theory. Let's start with a two person situation. As we have seen, the moral weight can be expressed as a linear combination of a utilitarian and egalitarian term which can be expressed as: $W(X) = a(x_1 + x_2) + b(x_2 - x_1)$. It is clear that the parameter $Q = \frac{a-b}{a+b}$. Here, the second term represents the inequality, but when there are more than 2 individuals, matters get more complicated to measure inequality. But with our QMM-formulation including the normalization factor $\sum_{i=1}^{N} Q^{i-1}$ in the denominator of the moral weight, we can easily get to a new measure of inequality. We can write the moral weight now as $W_{QMM}(X) = U(X)(1 - I(X))$, with $U(X) = \sum x_i / N$ the utilitarian function of aggregated well-being, and

$$I(X) = \frac{\sum_i \left(\sum_k Q^{k-1} - NQ^{i-1}\right) x_i}{\sum_l Q^{l-1} \sum_m x_m}$$

an inequality metric. The reason why the latter is an appropriate inequality metric, is because $I(X) = 0$ if all x_i are equal. Without the denominator term in $W_{QMM}(X)$, the latter would not be the case. Furthermore, the metric satisfies anonymity (impartiality, i.e. symmetry between different persons), scale independence (if all well-beings are multiplied by a factor, then $I(X)$ remains the same), population independence (the metric does not depend on the size of the population: if the population N increases, $I(X)$ does not increase proportional to N, because it is a ratio of functions that are extensive in N) and Pigou-Dalton (if there is a transfer from a rich person to a poor person, whereupon the rich person still remains richer than the poor person, the inequality should decrease). Our prioritarian theory automatically satisfies Pigou-Dalton, and therefore so does the inequality measure.

This inequality measure could be used in economics, if for example x_i is interpreted as the income or wealth of person i. The moral weight W_{QMM} can then be interpreted as a social welfare function (Sen, 1982). Note that we cannot derive an inequality measure from absolute prioritarianism; only our positional prioritarianism is suitable. For example, using absolute prioritarianism, we do not get $I(X) = 0$ when all x_i are equal, and there is also no scale independance.

Summary

In this article, we derived the quasi-maximin principle as a unified model for consequentialist theories of justice. Quasi-maximin is a theory close to maximin, but with a very small tendency towards utilitarianism. The "quasi" in QMM-theory is derived in two different ways: first from the impartial veil of ignorance with a high but not maximal risk aversion, and second from the equality principle with a small but not zero preference for efficiency. QMM is a special form of prioritarianism that is compatible with the moral virtue of compassion.

Some implications of QMM were discusses, and an elegant mathematical expression of the moral weight $W(X)$ of a situation X was derived. With this moral weight, we can avoid the intransitivity problem, the repugnant conclusion, and many other problems in ethics.

In a follow-up article, we will take a look into moral psychology and encounter some moral intuitions that might overrule the QMM-principle and cannot be derived from the Rawlsian veil of ignorance. One of those moral intuitions relates to the Kantian deontological ethics of the basic right. This is the right not to be used as merely means to our ends. We will extend the mathematical expression of the moral weight $W(X)$ to encapsulate this basic right.

Appendix 1: Quasi-maximin and risk aversion

In order to understand the relation between quasi-maximin and risk aversion, start with the following situation. Suppose you will be born as one of two persons. In the first situation, the qualities of life of the two persons (the pay-offs) are given by $X = (x_1; x_2)$, with $x_1 < x_2$. In this situation, there is some risk involved, as you can be born as the first person who is in the worst position. In the second situation, we have $Y = (y_1; y_2)$, with $y_1 = y_2$. So in the second situation, there is no uncertainty in outcomes, there is no gambling. The question is: In which world (which situation) would you prefer to be born?

If the moral weight of both situations are the same, than both situations are equivalent, and you would not prefer one over the other. Using $W_{QMM}(X) = W_{QMM}(Y)$, we get

$$y_1 = \frac{x_1 + Qx_2}{1 + Q}.$$

When the qualities of life in situation Y take this value, then situation Y is as good as situation X.

In the theory of risk aversion (Arrow, 1965; Pratt, 1964), there is the notion of the utility function, which is a function of the possible outcomes. The utility function for QMM-theory can be derived using the condition that when the two situations are equivalent, the average utilities in both situations should be the same. So we get as a condition:

$$\frac{u(x_1)+u(x_2)}{2}=\frac{u(y_1)+u(y_1)}{2}=u\left(\frac{x_1+Qx_2}{1+Q}\right).$$

Taking x_1 fixed, we can show that the utility function for the value x_2 (relative to x_1) that satisfies the above equation reads:

$$u_{x_1}(x_2)=\sqrt[n]{x_2-x_1},$$

with

$$n=\frac{\ln\left(\frac{1+Q}{Q}\right)}{\ln(2)};\quad Q=\frac{1}{2^n-1}.$$

(Note that we have a utility function which is very similar to the priority weighted utility function of absolute prioritarianism. See main text, where the priority weighted utility function reads $f(x)=\sqrt[n]{x}$. The only important difference, is that our utility function $u_{x_1}(x)$ is relative to x_1, the well-being of the worst-off person, instead of the zero threshold. In absolute prioritarianism, it is easy to see that $u(x)=f(x)$.)

Finally, we can derive the so called Arrow-Pratt measure of relative risk aversion (Arrow, 1965; Pratt, 1964), for fixed x_1, which is here defined as:

$$R_{x_1}(x_2)\equiv-(x_2-x_1)\frac{u''_{x_1}(x_2)}{u'_{x_1}(x_2)}=1-\frac{\ln 2}{\ln\left(\frac{1+Q}{Q}\right)}.$$

What is important, is that the above equation clearly shows a negative relationship between our parameter Q and the measure of risk aversion: when Q increases, the risk aversion decreases. We also note that our QMM-theory with constant parameter Q has a constant relative risk aversion. The corresponding utility function is called iso-elastic in the theory of risk aversion.

Appendix 2: The intransitivity problem

In this appendix we demonstrate that the intransitivity problem (Temkin, 1987) is avoided in the quasi-maximin theory. We give a general description. Let us start with situation S_1 whereby a number of people n_1 all have well-being w lower than all the other people. Due to translation invariance in de QMM-theory, we can write that the n_1 people have well-being $-w$ and the others have 0. In situation S_2 there are $N_2=n_1+n_2$ people, all with well-being $-w+\delta_1$, where δ_1 is a small, positive value. The others remain at well-being 0. Situation S_3 has $N_3=n_1+n_2+n_3$ people at well-being $-w+\delta_1+\delta_2$. And so forth.

The First Standard View can now be expressed in terms of the moral weights $W_{QMM}(S_k)$ for all situation S_k. This gives a series of conditions as inequalities:

$$-w\sum_{i=1}^{N_1}Q^{i-1}\geq-(w-\delta_1)\sum_{j=1}^{N_2}Q^{j-1}\geq\cdots\geq-\left(w-\sum_{k=1}^{p-1}\delta_k\right)\sum_{l=1}^{N_p}Q^{l-1}\geq\cdots$$

Write $\sum_{k=1}^{p-1}\delta_k=\Delta_p$, and keep in mind that

$$\sum_{i=1}^{N}Q^{i-1}=\frac{1-Q^N}{1-Q},$$

then we can write the above inequality between situation S_1 and S_p as

$$\frac{1-Q^{N_1}}{1-Q^{N_p}} \leq 1-\frac{\Delta_p}{w}.$$

Take the limit where p becomes very large (Q^{N_p} becomes negligible), then we get the constraint $\Delta_p \leq wQ^{N_1} < w$. This means that there is an upper bound on the increments of well-being. In other words, we cannot get the well-being arbitrarily close to zero without violating the First Standard View (FSV). If the δ_k are high enough so that the well-being $w-\Delta_p$ can move to 0, then there comes a point where the inequality in the series flips. I.e. $W_{QMM}(S_1) > W_{QMM}(S_2) > \cdots > W_{QMM}(S_p) < W_{QMM}(S_{p+1}) < \cdots$. So FSV is not valid anymore from a certain point onwards, if the increments are too large.

Note that if $Q=1$, we get the constraint $\Delta_p \leq w$, so in this case the well-being $w-\Delta_p$ can always move arbitrarily close to zero, and we see that utilitarianism implies that FSV is always satisfied, and the Second Standard View (SSV) is not. On the other hand, for $Q=0$ (maximin), we get the constraint $\Delta_p \leq 0$, so no increments are possible that respect the FSV. In other words: maximin always violates FSV (but respects SSV).

Let us now move to absolute prioritarianism. We can write the same series of inequalities, this time with the moral weight $W_{AP}(S)$ as a sum of priority weighted utility functions $f(x_i)$:

$$N_1 f(-w) \geq N_2 f(-w+\delta_1) \geq \cdots \geq N_p f(-w+\Delta_p) \geq \cdots.$$

The constraint now reads:

$$f(-w+\Delta_p) \leq \frac{N_1}{N_p} f(-w).$$

For the limit where p becomes very large, the right hand side goes to zero, and the constraint becomes $f(-w+\Delta_p) \leq 0$ or $-w+\Delta_p \leq 0$. So it is always possible to find a large enough population N_p whereby the well-being can move arbitrarily close to zero. Hence, absolute prioritarianism always satisfies FSV.

There is one exception, and that is when the utility function $f(x)$ becomes the threshold function of sharp sufficientarianism. Take the situation where N_1 people are below the threshold, and the others are above. This situation is better than situation S_2 where N_2 people are below the threshold at well-being $-w+\delta_1$. Moving to S_3, S_4 and so forth, this level of well-being moves closer and closer to the threshold 0, and every time the situation gets worse, because more and more people fall below the threshold. But at one point, the well-being $-w+\Delta_p$ crossed the threshold, and all the N_p people are now suddenly above the threshold. So we get the series:

$$W(S_1) > W(S_2) > \cdots W(S_{p-1}) \ll W(S_p) = W(S_{p+1}) = \cdots$$

The major difference between this sharp sufficiantarianism and absolute prioritarianism is that the utility function in absolute prioritarianism is an always increasing function, whereas in sharp sufficiantarianism it is constant for all values above the threshold (and also all values below the threshold). That is the reason why sharp sufficientarianism can still satisfy both FSV and SSV, be it in a strange way.

Appendix 3: The leveling down and welfare diffusion argument

There is an argument slightly similar to the repugnant conclusion, which leads to a conclusion that seems counter-intuitive to some people: the leveling down and welfare diffusion argument (Persson, 2010). Suppose in situation X we have $N-1$ very poor individuals and one individual with a high well-being (e.g. a high income). Let's write for simplicity situation $X=(0;0;\ldots;x_N)$. In situation Y, the well-being (income) of the N-th person is equally divided amongst all persons: $Y=\left(\frac{x_N}{N};\frac{x_N}{N};\ldots\frac{x_N}{N}\right)$. So the well-being of the N-th person is leveled down by diffusing its welfare. One can show that $W_{QMM}(X)=\sum_{i=1}^{N} Q^{i-1}x_i <$

$W_{QMM}(Y) = \sum_{i=1}^{N} Q^{i-1} y_i$, which means that it is good to drive someone in nearly complete poverty, even if it only means a negligible increase of income of a vast group of very poor people. Yes, it is even possible that situation $Z = (a; a; \ldots a)$, with $a < \frac{x_N}{N}$ is even better than situation X. So even an overall loss of income and well-being is permissible. This argument is similar to the efficiency argument which resulted in the introduction of the factor Q in the quasi-maximin theory. However, the QMM-equation of the moral weight does not seem to correspond with our need for efficiency in this case.
But also a situation $V = (b; b; \ldots b)$, with b slightly higher than $\frac{x_N}{N}$ can still be preferable. So most consequentialist theories struggle with this argument: situation V is better than situation X according to strict egalitarianism, maximin, prioritarianism and utilitarianism. Hence, if we want to avoid this conclusion, we seriously have to rewrite the mathematical formulation of these consequentialist theories. Only sufficientarianism is able to deal with this diffusion problem, to some level.
But do we really want to avoid this conclusion? After all, making N people a little bit happier can be quite a lot if N is big. And if N is small, the share of the richest person ($\frac{x_N}{N}$) will still be sufficiently big.
I leave it up for further research to find a formulation that avoids the welfare diffusion argument, in case the reader really wants to avoid it. One tentative solution is taking a more coarse grained description of the well-beings. In other words: $x_i \approx x_i + \delta$ for δ very small. For example, if we would look at the income, and we would distribute 100$ among 1 million poor people, those people will gain 0,0001$. We can suppose that such small increments will not be noticed. This means that the welfare diffusion has some limits.

Appendix 4: Independence and Allais paradox

Allais paradox (Allais, 1953, Kahneman and Tversky, 1979) is a famous example where the important axiom of independence in decision theory gets violated. An example of the paradox goes as follows. Suppose we have two experiments, each consisting of two gambles. In the first experiment, you can choose between gamble 1, whereby you have probability 100% to gain \$100, and gamble 2, whereby you have probability 89% to gain \$100, 1% to gain nothing and 10% to gain \$500. Most people would choose to play gamble 1, especially when they have risk aversion, because in gamble 2 they risk not to gain anything. In gamble 1 they are always certain to receive some benefit.
In a second experiment, people can choose between gambles 3 and 4. Gamble 3 has the following probabilities: 89% to gain nothing and 11% to gain \$100. Gamble 4: 90% to gain nothing and 10% to gain \$500. In this second experiment, people would prefer to play gamble 4, because they gain \$500 at an almost equal probability as gaining \$100 in gamble 3.
The curious thing is that experiments one and two are in fact quite similar, as can be seen by writing it as in the following table.

Gamble 1	Gamble 2	Gamble 3	Gamble 4
89% 100$	89% 100$	89% 0$	89% 0$
1% 100$	1% 0$	1% 100$	1% 0$
10% 100$	10% 500$	10% 100$	10% 500$

In the first experiment, we can decouple a 89% probability to gain \$100. So we can write for gamble 1 that you have 89% to gain \$100, and a 1%+10% to gain \$100. But if we would now set the gains for those 89% in gambles 1 and 2 equal to zero, we get gambles 3 and 4. So it is strange why people prefer gamble 1 over 2, but 4 over 3, because gamble 4 is quite similar to gamble 2, and 3 is similar to 1. Gambles 1 and 2 are both changed in the very same way (setting a gain to zero), and yet the order of preference of the gambles suddenly changes.

Let's translate this Allais paradox to an example of distributional justice. Suppose we have four situations, each involving three individuals. In situation S_1, all three persons have equal well-being 1. I.e.: $S_1 = (1;1;1)$. For the other situations, we have $S_2 = (1;0;5)$, $S_3 = (0;1;1)$, $S_4 = (0;0;5)$.

You can see that these situations are similar to the above gambles, whereby the first person refers to the 89% probability, the second to the 1%, and the third to the 10% probability, as in the above table. If we have risk aversion, we'd prefer A over B. But in the above example of Allais paradox one would choose situation 4 over situation 3.

It is also clear to see that in the choice between situations 1 and 2, the first person has always well-being 1, so its presence should not matter according to the axiom of independence. The same goes for situations 3 and 4. But if we neglect the first person, then situation 1 becomes nothing but situation 3 (which becomes $(1;1)$), and situation 2 equals situation 4.

Let us now finally see what happens with the moral weight W_{QMM}. In the main text, we mentioned that the sum should only include those persons whose well-being is different in the different choices. So in calculating whether S_1 is better than S_2, we do not include the first person, because the first person has the same well-being in both choices. Then we get $W(S_1) = 1 + Q > W(S_2) = 5Q^2$, when $Q < 1/4$. But if this is the case, then also automatically $W(S_3) > W(S_4)$. So by neglecting the person whose well-being is not affected, we respect the axiom of independence.

But what would happen if we did incorporate the independent person in the expression of the moral weight? Then we get $W(S_1) = 1 + Q + Q^2 > W(S_2) = Q + 5Q^2$, when $Q < 1/2$. But if $Q > 1/4$, we see that $W(S_3) = Q + Q^2 < W(S_4) = 5Q^2$. In other words, when Q is in the range $\frac{1}{4} < Q < \frac{1}{2}$, it is possible that situation 4 is better than situation 3, although situation 2 was worse than situation 1. This is Allais paradox.

What this means is that we have two choices: 1) we want to respect the axiom of independence, which can be done by not taking the independent person up in the calculations, or 2) we want to respect peoples intuitions that result in Allais paradox (and violate the axiom of independence), which can be done by taking the independent person up in the calculations. I'll leave it to the reader to choose what he prefers. What is important, is that our QMM-theory can be very easily adapted in order to incorporate Allais paradox: we simply have to take all persons up in the calculations, instead of only the persons whose well-beings are affected by the choices.

Appendix 5: Another moral weight function for prioritarianism

There is another possibility for the moral weight function that corresponds with what is known as the strong prioritarian totalism: "One distribution is at least as just as a second if and only if its total weighted benefits are at least as great, where the weights for a given benefit finitely decrease as the benefits prior to the increment finitely increase." (Vallentyne, 2009 p150)

The above expression can be interpreted with the following mathematical formulation. Suppose again that we have the ordered set of qualities of life $X = (x_1; x_2; \dots x_N)$. Now, the weighted well-being of person 1 equals x_1, the weighted well-being of person 2 does not equal Qx_2, but $x_1 + Q(x_2 - x_1)$. For the third person we have $x_1 + Q(x_2 - x_1) + Q^2(x_3 - x_2)$. We see that each difference in well-being is now weighted, instead of the total well-being.

In general, the weighted well-being of person i equals

$$B_i = \sum_{j=1}^{i} Q^{j-1}(x_j - x_{j-1}),$$

with $x_0 \equiv 0$.

With this, we get the new formulation of the moral weight function of prioritarian totalism:

$$W_P(X) = \sum_{i=1}^{N} B_i = \sum_{i=1}^{N} (N - i + 1)Q^{i-1}(x_i - x_{i-1}).$$

Note that if $Q = -\infty$, we get strict egalitarianism. For $Q = 0$ and $Q = 1$, we get again maximin and utilitarianism respectively.

Next, we also want to avoid the repugnant conclusion in this new formulation. This can be done by slightly modifying the above weight function into:

$$W_P(X) = \frac{\sum_{i=1}^{N}(N - i + 1)Q^{i-1}(x_i - x_{i-1})}{\sum_{i=1}^{N}(N - i + 1)Q^{i-1}}.$$

As with the QMM weight function, we can also derive the risk aversion equations for $W_P(X)$. The corresponding utility function equals:

$$u_{x_1}(x_2) = \sqrt[n]{x_2 - \frac{(1-Q)}{(2-Q)}x_1},$$

with

$$n = \frac{\ln\left(\frac{2}{Q}\right)}{\ln(2)} \text{ and } Q = \frac{1}{2^{n-1}}.$$

The Arrow-Pratt measure of relative risk aversion for fixed x_1 can be defined as:

$$R_{x_1}(x_2) \equiv -\left(x_2 - \frac{(1-Q)}{(2-Q)}x_1\right)\frac{u''_{x_1}(x_2)}{u'_{x_1}(x_2)} = 1 - \frac{\ln 2}{\ln\left(\frac{2}{Q}\right)}.$$

References

Allais, M., 1953, *Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine*, Econometrica 21, 503-546.

Angner, E., 2004, *Revisiting Rawls*: *A Theory of Justice in the Light of Levi's Theory of Decision*. Theoria 70 (1):3-21.

Arrow, K.J., 1965, *The theory of risk aversion*, in *Aspects of the Theory of Risk Bearing*, by Y.J. Saatio, Helsinki.

Arneson, R.J., 1989, *Equality and Equal Opportunity for Welfare*, Philosophical Studies 56, pp. 77-93.

Broome, I., 1991, *Weighing Goods*, Oxford; Basil Blackwell.

Brown, C., 2007, *Prioritarianism for variable populations*, Philosophical Studies, 134:325–361.

Cohen, G. A., 1989, *On the Currency of Egalitarian Justice*, Ethics 99, pp. 906-944.

Crisp, R, 2008, *Well-being*, Stanford Encyclopedia of Philosophy.

Dick, J., 1975, *How to Justify a Distribution of Earnings*, Philosophy and Public Affairs, 4: 248-72.

Dworkin, R., 1981, *What is Equality? Part 1: Equality of Resources*, Philosophy and Public Affairs, 10: 185-246

Holtug, N., 2006, *Prioritarianism*, in: Egalitarianism, new essays on the nature and value of equality, Holtug & Lippert-Rasmussen (eds.), Clarendon Press.

Kahneman, D. and Tversky, A., 1979, *Prospect Theory: An Analysis of Decision under Risk*, Econometrica 47, 263-291.

Kagan, S., 1999, *Equality and Desert*, in L.P. Pojman and O. McLeod (eds.), *What Do We Deserve? A Reader on Justice and Desert*, Oxford and New York: Oxford University Press, pp. 298-314.

Lamont, J., 1994, *The Concept of Desert in Distributive Justice*, The Philosophical Quarterly, 44: 45-64

Maslow, A., 1943, *A Theory of Human motivation*, Psychological Review, Vol. 50, nr.4, pp. 370–396.
McCarthy, D., 2003, Prospects for prioritarianism
McCarthy, D., 2008, *Utilitarianism and prioritarianism II*, Economics and Philosophy, 24 (2008) 1–33.
Miller, D., 1976, *Social Justice*, Oxford: Clarendon Press.
Milne, H., 1986, *Desert, effort and equality*, Journal of Applied Philosophy, 3: 235-243.
Mulgan, T., 2006, *Future People. A Moderate Consequentialist Account of our Obligations to Future Generations*, Oxford: Clarendon Press.
Narveson, J., 1967, *Utilitarianism and New Generations*, Mind, 76: 62-72.
Nussbaum, M., 1992, *Human Functioning and Social Justice: In Defense of Aristotelian Essentialism*, Political Theory 20, pp. 202-246.
Nussbaum, M., 2000, *Women and Human Development: The Capabilities Approach*, Cambridge: Cambridge University Press.
Parfit, D., 1984, *Reasons and Persons*, Oxford: Clarendon Press.
Parfit, D., 1991, *Equality or Priority, The Lindlev Lecture*, Lawrence: University of Kansas.
Parfit, D., 1997, *Equality and Priority*, Ratio 10, pp. 202-221.
Persson, I., 2010, *Prioritarianism, leveling down and welfare diffusion*, Ethical Theory and Moral Practice, DOI: 10.1007/s10677-010-9242-y.
Pratt, J. W., 1964, *Risk aversion in the small and in the large*, Econometrica 32, 122-136.
Rabinowicz, W., 2002, *Prioritarianism for Prospects*, Utilitas 14, 2–21.
Rawls, J., 1971, *A Theory of Justice*, Cambridge: Harvard University Press, rev. ed. 1999.
Rawls, J., 2001, *Justice as Fairness: A Restatement*, ed. by E. Kelly, Cambridge: Harvard University Press.
Riley, J., 1989, *Justice Under Capitalism*, Markets and Justice, ed. J.W. Chapman, New York: New York University Press, 122-162
Roemer, J.E., 1996, *Theories of Distributive Justice*, Cambridge: Harvard University Press.
Rosenberg, M., 2003, *Nonviolent Communication: A Language of Life.* 2nd Edition. Encinitas, CA: PuddleDancer Press.
Sadurski, W., 1985, *Giving Desert Its Due*, Dordrecht, Holland: D. Reidel.
Scheffler, S., 1982, *The Rejection of Consequentialism*, Oxford and New York: Oxford University Press.
Sen, A., 1982, *Choice, Welfare and Measurement*, MIT Press.
Sen, A., 1992, *Inequality Reexamined*, Cambridge: Harvard University Press.
Temkin, L., 1987, *Intransitivity and the Mere Addition Paradox*, Philosophy and Public Affairs 16, no. 2, pp. 138–187.
Vallentyne, P., 2009, *Sen on sufficiency, priority and equality*. In *Amartya Sen*, C.W. Morris et al., Cambridge University Press, p. 139.
Weirich, P., 1983, *Utility Tempered with Equality*, Nous 17, pp. 423-39.