

Forthcoming in *Utilitas*

Utilitarianism, Multiplicity, and Liberalism

Jeff Sebo
New York University

Abstract

In this paper I argue that utilitarianism requires us to tolerate intrapersonal disagreement for the same reasons that it requires us to tolerate interpersonal disagreement. I begin by arguing that multiplicity has many of the same costs and benefits as multiculturalism: it causes conflict, but it also allows us to perform experiments in living, adopt a division of labor, compartmentalize harm, and learn from ourselves. I then argue that, in light of these costs and benefits, utilitarianism requires us to adopt a “liberal system of individual self-government,” according to which we should not try to impose a single, unified set of beliefs and values on ourselves. Finally, I argue that we should apply this policy of liberal toleration to intrapersonal disagreement about utilitarianism too: if we want to maximize utility, then we should tolerate inner conflict not only about how to maximize utility but also about whether we should be trying to maximize utility in the first place.

1. Introduction

Many of us have the capacity to control how psychologically unified we are, at least to a degree. That is, we have the capacity to make our beliefs, desires, plans, and character traits a bit *more* coherent and stable than they presently are, and we also have the capacity to make them a bit *less* coherent and stable than they presently are. This raises an important question for morality: how psychologically unified should we try to be, and why?

My aim in this paper is to provide a utilitarian framework for answering this question, as well as to explore some of the implications of this framework. Specifically, I will argue that, if we want to maximize utility, then we should tolerate intrapersonal conflict for the same reasons that we should tolerate interpersonal conflict. Compare: at least since John Stuart Mill published *On Liberty*, philosophers have taken it for granted that, while multiculturalism has many costs, it also has many benefits: for example, it causes a lot of conflict, but it also allows us to perform experiments in living, adopt a division of labor, compartmentalize harm, and learn from each other. Thus, many philosophers now believe, utilitarianism requires us to adopt a liberal system of collective self-government, according to which the state should not try to impose a particular conception of the good life on its citizens but should rather allow each citizen to pursue their own conception of the good life, compatibly with other citizens being able to do the same.ⁱ

I will argue that the same considerations apply at the intrapersonal level as well. That is, I will argue that multiplicity has many of the same costs and benefits as multiculturalism, and therefore utilitarianism requires us to adopt a liberal system of individual self-government, according to which we should not try to impose a particular conception of the good life on ourselves across situations but should rather allow ourselves in each situation to pursue our

conception of the good life in that situation, compatibly with our being able to do the same in other situations. (To be clear, as I will discuss below, liberalism does not imply that we should try to *achieve* multiplicity, but rather only implies that we should *tolerate* multiplicity, within certain limits, if and when we do achieve it.) I will also argue that we should apply this principle of liberal toleration to intrapersonal disagreement about utilitarianism itself: if we want to maximize utility, then we should tolerate inner conflict not only about how to maximize utility but also about whether or not we should be trying to maximize utility in the first place.

I will proceed as follows. In section 2, I will explain what I mean by psychological unity and multiplicity. In section 3, I will review the costs of multiplicity, including indecision, compulsion, half-heartedness, and false belief. In section 4, I will argue that multiplicity has many of the same benefits as multiculturalism: it allows us to perform experiments in living, adopt a division of labor, compartmentalize harm, and learn from ourselves. Finally, in section 5, I will argue that, in light of these considerations, utilitarianism requires us to adopt a liberal system of individual self-government rather than an illiberal system of individual self-government, and I will explore some of the moral implications of this idea.

2. Unity and multiplicity

I begin by explaining what I mean by psychological unity and multiplicity.

First, what do I mean by unity? Throughout this paper I will use ‘unity’ in the following way: a set of beliefs, desires, plans, character traits, and so on are unified to the degree that they are coherent and mutually supporting. (In what follows I will use ‘psychological states’ to refer to all these states and dispositions.) So a set of beliefs is unified to the degree that all of the beliefs

in this set are coherent and mutually justifying. A set of desires is unified to the degree that all of the desires in this set are coherent and mutually satisfiable. A set of plans is unified to the degree that all of the plans in this set are coherent and mutually achievable. A set of character traits is unified to the degree that all of the character traits in this set cause similar kinds of behavior. And more generally, a set of psychological states is unified to the degree that they all support each other in these kinds of ways, e.g. to the degree that your beliefs and desires justify your plans and your beliefs, desires, and character traits cause you to act on your plans.

A set of psychological states is *fragmented*, then, to the degree that it falls short of this ideal. This can happen in a couple of ways. First, a set of psychological states is *synchronically* fragmented if you have at least some different, and conflicting, psychological states at one and the same time, e.g. if you believe that having an affair will ruin your marriage and you want to avoid ruining your marriage, yet you have no plans to stop having an affair. Second, a set of psychological states is *diachronically* fragmented if you have at least some different, and conflicting, psychological states across time, e.g. if you want to have an affair while around your friends but not while around your family. Stated in terms of this distinction, my interest in this paper is in assessing the value of what we might call *situational diachronic fragmentation*. That is, I am interested in exploring the costs and benefits of having at least some different, and conflicting, psychological states in different situations, such as at home and at work.

Some people find it useful to talk about situational diachronic fragmentation by describing themselves as having *multiple selves*. That is, instead of describing themselves as having different psychological states in different situations, they find it useful to describe themselves as having different *selves* in different situations, such as a family self and a work self. It will be useful for us to use similar language here, but we need to be a bit more precise than

people often are, since terms like ‘family self’ and ‘work self’ are multiply ambiguous.

Specifically, we need to make a distinction between the *roles* that we play in different situations, the *personas* that we adopt while playing those roles, and the *personalities* that we develop while adopting those personas.

First, we play multiple *roles* if we have different sets of responsibilities in different situations. We all play many roles in everyday life: for example, you might be a spouse, parent, friend, teacher, writer, activist, and more. Sometimes, and in some respects, your roles will complement each other: for instance, the roles ‘teacher’ and ‘writer’ are mutually supportive, and they also support roles like ‘spouse’ and ‘parent’ insofar as you need money to support your family. Likewise, the roles ‘spouse’ and ‘parent’ are mutually supportive, and they also support roles like ‘teacher’ and ‘writer’ insofar as you need love and support from your family in order to do good work. But other times, and in other respects, your roles will conflict with each other: for instance, if you have a high teaching load, then you might not be able to spend as much time and energy writing as you might like, and vice versa. And of course, similar conflicts will arise within your personal roles and between your personal and professional roles – especially when particular roles place a disproportionate demand on your time and energy.

Second, we have multiple *personas*, in my sense of the term, if we *present ourselves* as having at least some different, and conflicting, psychological states in different situations. For example, if you present yourself as compassionate while around your family and as dispassionate while around your colleagues, then we can say that you have a compassionate family persona and a dispassionate work persona.ⁱⁱ As this example shows, our roles and personas are related: we pick roles that fit our personas, and we also shape our personas to fit our roles. But the relationship between roles and personas is subtle. Sometimes the difference between them will

be thin: for instance, if your role is to play Hamlet, then one of the central requirements of this role will be that you adopt the persona of Hamlet. But other times the difference between them will be a bit thicker: for instance, if your role is to be a computer programmer, then your playing this role well will be compatible with your having many different personas while at the office.

Third, we have multiple *personalities*, in my sense of the term, if we actually *do* have at least some different, and conflicting, psychological states in different situations. For example, if you actually *are* a bit more compassionate while around your family and a bit more dispassionate while around your colleagues, then we can say that you have a compassionate family personality and a dispassionate work personality.ⁱⁱⁱ As with our roles and personas, our personas and personalities are related: we pick personas that fit our personalities, and we also shape our personalities to fit our personas. But also as with our roles and personas, the relationship between our personas and personalities is subtle. Sometimes the difference between them will be thin: for instance, if you want to present yourself as you actually are while in therapy. But other times the difference between them will be a bit thicker: for instance, if you want to present yourself as a loyal American while at work so that you can steal state secrets for Russia.

Stated in terms of the distinction between roles, personas, and personalities, my interest in this paper is in exploring the ethics of having multiple personalities in the sense that I have in mind here. I will be using ‘multiplicity,’ then, to refer to the tendency to have at least some different, and conflicting, psychological states in different situations. But it will be important for us to consider the ethics of playing multiple roles and adopting multiple personas as well, since, as I have said, our roles, personas, and personalities are closely related. If we play multiple roles, then we tend to adopt multiple personas as a result; and, if we adopt multiple personas, then we tend to develop multiple personalities as a result. Thus, in what follows we will have to question

the common, if tacit, assumption that many of us make in everyday life, which is that we should try to play multiple roles while having a single, unified personality. In some cases it might be possible and desirable for us to do this, but in other cases it might not be possible or desirable for us to do this. And in these latter cases, we will face a choice: is it more important for us to play multiple roles, or is it more important for us to have a single, unified personality?

3. Costs of multiplicity

I will now review the costs of having at least some different, and conflicting, psychological states in different situations. I will start with the individual costs and then turn, briefly, to the social costs. Fortunately, these costs are all well-known, so we can move relatively quickly here.

First, many philosophers have pointed out that, if we have different and conflicting psychological states in different situations, then we will be more likely to suffer from *indecision*. That is, we will be more likely to change our minds about what we have most reason to do from situation to situation.^{iv} For example, suppose that you tend to think, while at work, that you should take a promotion that will allow you to advance in your career, and you tend to think, while at home, that you should take a demotion that will allow you to spend more time with your family. In this case, you might not be able to commit to a decision about what to do: for example, you might decide, while at work, that you should take the promotion, and then you might decide, while at home, that you should take the demotion instead. Similarly, suppose that you tend to think, during the day, that you should come home by about midnight tonight, and you tend to think, during the night, that you should stay out till about 2am instead. In this case, you might not be able to commit to a decision about what to do either: for example, you might

decide, at home, to come home at midnight tonight, and then you might decide, at the bar, to stay out till 2am instead. And of course, insofar as you experience indecision, it will be difficult for you to commit to the kinds of projects and relationships that make life worthwhile.

Second, many philosophers have also pointed out that, if we have different and conflicting psychological states in different situations, then we will be more likely to suffer from *compulsion*. That is, we will be more likely to behave contrary to what we think, at the time of action, we have most reason to do.^v For example, suppose that you decide that you should keep your current job, and you maintain this judgment at work as well as at home. Nevertheless, you find yourself taking the promotion while at work anyway – not because you think that you *should* but rather because your desire to take the promotion is so powerful that it causes your behavior directly. Similarly, suppose that you decide that you have most reason to get home around 1am, and you maintain this judgment while at home as well as while at the bar. Nevertheless, you find yourself ordering another drink while at the bar anyway – again, not because you think that you *should* but rather because your desire to keep drinking is so powerful that it causes your behavior directly. And of course, insofar as you compulsively behave rather than autonomously act, it will be difficult for you to follow through on your commitments and sustain the kinds of projects and relationships that make life worthwhile.

Third, many philosophers have also pointed out that, if we have different and conflicting psychological states in different situations, then we will be more likely to suffer from *half-heartedness*. That is, we will be more likely to feel conflicted about, and only partly motivated to pursue, our chosen course of action.^{vi} For example, suppose that you decide to keep your current job, and this time you actually do it. Nevertheless, you wish that you could have taken the promotion while at work, and you wish that you could have taken the demotion while at home.

And this not only takes away from your enjoyment of your work and family roles but also takes away some of the motivation that you need in order to play these roles well. Similarly, suppose that you decide to get home by 1am tonight, and this time you actually do it. Nevertheless, you wish that you could stay out till 2am while at the bar, and you wish that you had come home by midnight while at work the next day. And as before, this kind of ambivalence makes it hard for you to fully enjoy yourself or be fully effective in either of your roles.

Fourth, if we have different, and conflicting, beliefs in different situations, then we will have at least some false beliefs. Of course, this might not always be a matter of fully believing a false proposition or fully denying a true proposition. Instead, it might merely be a matter of having a bit less credence in a true proposition than we should, or of having a bit more credence in a false proposition than we should. Still, insofar as we have false beliefs even in this weak sense, we will be more likely to pursue ends that we would not pursue if we were fully informed, or to take means to our ends that we would not take if we were fully informed.

Moreover, these individual costs double as social costs, since, insofar as we suffer from indecision, compulsion, half-heartedness, and false belief, we will not be able to reliably follow through on our commitments to the other people in our lives, nor will we be able to reliably follow through on our aspirations to make the world a better place (insofar as we have those aspirations at all). Granted, this might not be true for all of our projects and relationships: for example, you might be able to reliably participate in a weekly card game even if you suffer from indecision, compulsion, half-heartedness, and false belief. But it will be true for the vast majority of projects and relationships that allow us to do good in the world: all else equal, the more indecision, compulsion, half-heartedness, and false belief we experience, the harder it will be for

us to have a career, have a family, or do any number of other things that would be good not only for us but for everybody else as well.

As a result of these individual and social costs of multiplicity, many philosophers have claimed (at least in passing; not many have felt the need to devote entire articles or books to this thesis) that we should try to be as psychologically unified as possible. For example, Derek Parfit, in the course of arguing that psychological connectedness and continuity with any cause is what matters in survival, morality, and rationality, writes:

I want my life to have certain kinds of overall unity. I do not want it to be very episodic, with continual fluctuations of my desires and concerns. Such fluctuations are compatible with full psychological continuity, but they would reduce psychological connectedness. This is [a] kind of change that most of us would regret.^{vii}

In fact, some philosophers even claim that we need to be psychologically unified in order to be autonomous at all.^{viii} For example, Harry Frankfurt writes that one of the aims and functions of agency is to

integrate the person both dynamically and statically. Dynamically insofar as it provides ... for coherence and unity of purpose over time; statically insofar as it establishes ... a reflexive or hierarchical structure by which the person's identity may in part be constituted.^{ix}

But while I agree that agency requires self-government, I do not agree that self-government requires unity. Consider the collective case: the ideal of collective self-government is clearly compatible with our having different beliefs and values in society. As long as we jointly commit to a fair set of laws, we can participate in collective self-government whether or not we have

“coherence and unity of purpose” in society – and, indeed, there are respects in which our having different beliefs and values in society helps, rather than hinders, our collective autonomy.

Similarly, I will argue that the ideal of individual self-government is fully compatible with our having different beliefs and values in different situations. As long as we commit to a fair set of laws over time, we can participate in individual self-government whether or not we have

“coherence and unity of purpose over time” – and, indeed, there are respects in which our having different beliefs and values over time helps, rather than hinders, our individual autonomy.

4. Benefits of multiplicity

I will now argue that multiplicity has many of the same benefits as multiculturalism: it allows us to perform experiments in living, adopt a division of labor, compartmentalize harm, and learn from ourselves. As in the last section, I will start with the individual benefits of multiplicity and then turn, briefly, to the social benefits. Of course, these benefits are a bit more familiar in the interpersonal case than in the intrapersonal case. Still, at least some psychologists are starting to appreciate them in the intrapersonal case, so we will be able to draw on recent psychology as well as common sense in order to make these points.

The first benefit of multiplicity is that it allows us to perform experiments in living.^x That is, it allows us to try out different ways of life and see which works best for us. This means trying out different roles, personas, and personalities and seeing how they fit. Many of us perform these experiments naturally, starting early in life. For instance, we act playful around our families, extroverted around the cool kids, and introverted around the outcasts. And as a result of performing these experiments, many of us develop multiple personalities as well, at

least to a degree. For example, we actually *become* a bit more playful around our families, a bit more extroverted around the cool kids, and a bit more introverted around the outcasts, as a result of drawing on these character traits while adopting these personas. And insofar as we have this experience, this is good, because it teaches us about who we are and who we are capable of becoming. Granted, we might experience a bit of inner conflict as a result of performing our experiments in living, e.g. we might think, while around the cool kids, that we should spend more time around the cool kids and then think, while around the outcasts, that we should spend more time around the outcasts. But we will also, more generally, be able to gravitate towards ways of life that, we agree across situations, are good for us, with the result that we will be able to make much more progress in the long run if we take an experimental approach to life than if we take a non-experimental approach to life.

The second benefit of multiplicity is that it allows us to adopt a personal division of labor. That is, it allows us to play different roles in different situations, to adopt different personas for our roles, and to develop different personalities for our personas. Many of us will discover through our experiments in living that we need to play a complex set of roles in order to live well. For instance, we may decide that we need to focus primarily on work in order to live well, but that we also need to have a family and friends so that we can do well in our career. Alternatively, we may decide that we need to focus primarily on family in order to live well, but that we also need to have a career so that we can support our family. Either way, then, many of us will need to play different and partly conflicting roles in life in order to live well. And as with our experiments in living, many of us will find that, in the course of playing these roles, we will naturally develop multiple personas for playing them, and we will naturally develop multiple personalities for adopting these personas. If so, then, once again, this is good, because it will

make us much more capable social performers overall. Granted, as with our experiments in living, we might experience a bit of inner conflict as a result of adopting this division of labor, e.g. we might think, while at work, that we should prioritize work and we might think, while at home, that we should prioritize family. But we will also, more generally, be able to immerse ourselves in a much wider range of projects and relationships if we allow ourselves to develop different personalities for different roles than if we try to maintain a one-size-fits-all personality for all of them.

The third benefit of multiplicity is that it allows us to compartmentalize harm. One of the main causes of dissociative identity disorder is that, when we have a traumatic experience, we cope with this experience by creating multiple personalities – one who remembers the traumatic experience, and another (or perhaps several others) who are protected from having to do this.^{xi} Of course, dissociative identity disorder is an extreme, and pathological, response to psychological trauma. Yet we can also use our capacity for dissociation to compartmentalize harm in less extreme and more precautionary ways: specifically, we can develop different personalities with different talents and interests so that, if one of our personalities has a traumatic experience and/or becomes less useful for us, then our other personalities will be relatively protected from that experience and/or might be relatively more useful for us.^{xii} For example, if one of your friends or family members passes away, then it will help for you to have other projects and relationships to throw yourself into, thereby bringing out personalities who feel the loss of your friend or family member less acutely. Similarly, if you lose your job, then it will help to have other talents and interests to fall back on. Thus, while having multiple personalities might cause us to develop more talents and interests than we can fully pursue, it will also, more

generally, allow us to survive and flourish as a person when bad things happen to us much more effectively than a single, coherent and stable personality would allow us to.

The final benefit of multiplicity that I will talk about here is that we can learn from ourselves and become more tolerant and open-minded as a result. The reason is the same as in the interpersonal case.^{xiii} If we always had the same perspective, then it would be easy to dismiss other perspectives without fully engaging with them. As a result, we would eventually start to accept our own perspective dogmatically, without fully understanding why we accept it or how we might improve upon it. But if we have different perspectives in different situations, then it will be much harder for us to dismiss other perspectives, and it will be much easier for us to engage with them and understand and improve upon our own current perspective as a result. For example, if you feel happy-go-lucky around your family and cynical and world-weary around your colleagues, then the interaction between these perspectives will make it difficult for you to accept either view unquestioningly, and will make your personalities much more likely to gravitate towards a balanced, nuanced perspective on the value of life over time. Thus, while having different perspectives might cause us to have at least some false beliefs, it will also, more generally, motivate us to challenge, improve upon, and justify our beliefs much more than a single, coherent and stable perspective would motivate us to.

As with the costs of multiplicity, these individual benefits double as social benefits. This is true not only because of the general connection between happiness and productivity but also because performing experiments in living, adopting a division of labor, compartmentalizing harm, and learning from ourselves are useful means to the end of making a difference in the world. For instance, if we try out different ways of life, then we will be better able to discover which ways of life will allow us to do the most good. If we play different roles and develop at

least some different psychological states for each role, then we will be better able to do good through our roles – either because we will be better able to do good through *all* of them or because we will be better able to do good through *some* of them and to support this effort through others. If we have different talents and interests, then we will be better able to cope with setbacks that we experience with respect to some of our projects and relationships. And if we have at least some different, and conflicting, beliefs, then we will be more open to new possibilities for making a difference as time passes and circumstances change.

5. A liberal system of individual self-government

I will now argue that, in light of these considerations, utilitarianism requires us to adopt a liberal system of individual self-government, according to which we should not impose a particular set of beliefs and values on ourselves across situations but should rather allow ourselves in each situation to pursue our beliefs and values in that situation, compatibly with our being able to do the same in other situations as well. (To be clear, as I will discuss below, liberalism does not imply that we should try to *achieve* multiplicity, but rather only implies that we should *tolerate* multiplicity, within certain limits, if and when we do achieve it.) I will also argue that we should apply this policy of liberal toleration to inner conflict about utilitarianism as well: if we want to maximize utility, then we should tolerate inner conflict not only about how to maximize utility but also about whether or not we should be trying to maximize utility in the first place.

Before I begin, a few clarifications about my terminology and strategy in what follows. First, a note about what I mean by ‘utility’. I will be focusing here on what system of individual self-government we should adopt according to classical utilitarianism, where classical

utilitarianism is the view that an action is right if and only if it maximizes pleasure in the world. I am restricting my focus to this theory for the sake of simplicity and specificity, as well as because philosophers are already familiar with the idea that utilitarianism implies liberalism in the collective case. With that said, I want to be clear from the outset that very little in my argument will actually turn on the details of classical utilitarianism. As I will discuss below, I think that the general form of my argument applies to any consequentialist theory of morality or rationality (with a few caveats and exceptions), and so I hope that my discussion here will be of interest to consequentialists in general as well as to utilitarians in particular.

Second, a note about what I mean by a ‘system of individual self-government.’ I will be using this term to refer to our procedures for creating, interpreting, and enforcing the laws that govern our treatment of ourselves. In using this term (and defining it this way), I am intentionally drawing an analogy with systems of collective self-government, i.e., procedures for creating, interpreting, and enforcing the laws that govern our treatment of each other. My reason for drawing this analogy is simple: our thinking about the value of liberty and toleration is clearer and more fully developed in the collective case than in the individual case. Thus, I think that it will be useful for us to start by considering the collective case, and then to observe that the same considerations apply in the individual case as well – even if the details are a bit different.

With that in mind, consider the collective case. The central question for a utilitarian political philosopher is: what system of creating, interpreting, and enforcing the laws that govern our treatment of each other will maximize utility? Of course, there are many options to choose from, but for our purposes here it will be enough to identify two ends on a spectrum. At one end of the spectrum, the state can adopt a *liberal* system of collective self-government. On this view, the state should not try to impose a particular “comprehensive conception of the good,” i.e., a

particular comprehensive conception of how we ought to live, on its citizens. Instead, the state should allow each citizen to pursue their own comprehensive conception of the good, compatibly with other citizens being able to do the same.^{xiv} If we take this approach, then we can expect that a relatively wide range of comprehensive conceptions of the good will emerge in political society, and our central question will be: how can we resolve conflicts among people with different beliefs and values in a fair way?

At the other end of the spectrum, we can adopt an *illiberal* system of collective self-government. On this view, the state *should* impose a particular comprehensive conception of the good on its citizens. This imposition can take more or less extreme forms. In a less extreme case, the state can *promote* a particular set of beliefs and values but not force the population to live according to it. In a more extreme case, the state can not only promote a particular set of beliefs and values but also force the population to live according to it, for example by manipulating, coercing, or physically restraining people. If we take either of these approaches, then we can expect that a relatively narrow range of comprehensive conceptions of the good will emerge in political society, and our central question will be: how can we effectively unite as many people as possible behind the beliefs and values that we sincerely think are best?

One might think that utilitarians would support illiberalism, since the central aim of illiberalism is to unify people behind the set of beliefs and values that will cause them to live well. But as it happens, most utilitarians support liberalism instead.^{xv} Not because they value liberty and autonomy for their own sake, but rather because they believe that a free, open society is more likely to maximize utility in the long run than an authoritarian, repressive society is. The most famous expression of this idea comes from John Stuart Mill in *On Liberty*, but many other utilitarians have argued for this conclusion as well.^{xvi} For our purposes here we can summarize

these arguments as follows: even though liberty allows for a lot of conflict (since it allows for the kind of multiculturalism that causes a lot of conflict), it also allows us to make much more progress in the long run than illiberalism does, since it allows us to discover, and gravitate towards, the social and political systems that work best for us – or at least, it allows us to come closer to achieving that aim than if we simply assumed, without evidence, that our current social and political systems work best for us. Moreover, a liberal system of collective self-government allows us to resolve interpersonal conflicts in a relatively fair way, with the result that we can enjoy many of the benefits of multiculturalism while minimizing many of the costs.

Importantly, many utilitarians also think that these considerations extend to disagreement about utilitarianism itself. That is, they think that the state should tolerate disagreement not only about *how* to maximize utility but also about *whether* to maximize utility. Of course, this is not to say that state officials cannot believe in, or act on, the principle of utility themselves. Rather, the idea is that, insofar as state officials are creating, interpreting, and enforcing the laws that govern our treatment of each other, they have an indirect utilitarian obligation to not impose the principle of utility – or indeed any comprehensive conception of the good – on everyone else.

Now consider the individual case. Some of us have multiple personalities, and some of us do not. Still, most if not all of us have the capacity to *develop* multiple personalities, and the mere possibility of multiplicity raises the question: should we allow ourselves to have different, and conflicting, beliefs and values in different situations, or should we try to impose a single, unified set of beliefs and values on ourselves across situations? If we take the former approach, then we will have adopted what I will call a *liberal system of individual self-government*, and our central question will be: how can we resolve the conflicts that we experience across situations in a fair way? In contrast, if we take the latter approach, then we will have accepted what I will call

an *illiberal system of individual self-government*, and our central question will be: how can we effectively unite ourselves behind the beliefs and values that we currently think are best?

Most philosophers assume that we should adopt what I am calling an illiberal system of individual self-government: they assume that we should try to have a single, unified set of beliefs and values, even if we have to manipulate, coerce, and physically restrain ourselves across situations in order to accomplish this aim.^{xvii} In my view, however, the costs and benefits of multiplicity make it the case that utilitarians, at least, should adopt liberalism instead, for the same general reason as in the collective case: a free, open approach to our own psychology is more likely to maximize utility in the long run than an authoritarian, repressive approach to our own psychology is.

To see why we should accept this claim, consider how the utilitarian argument for liberalism applies at the individual level. First, even though liberalism allows for inner conflict (since it allows for the kind of multiplicity that causes inner conflict), it also allows us to make much more progress in the long run than illiberalism does, since it allows us to discover, and gravitate towards, the ways of life that work best for us – or at least, it allows us to come closer to achieving that aim than if we simply assumed, without evidence, that our current set of roles, personas, and personalities is best for us. To be clear, I am not saying here that, in light of the costs and benefits of multiplicity, we know for a fact that multiplicity is best for us, and therefore we should adopt liberalism so as to achieve multiplicity. Instead, what I am saying is that, in light of the costs and benefits of multiplicity, we do not know for a fact that multiplicity is *not* best for us. It is an open, empirical question what kind of psychology will allow us to maximize utility, and therefore we should adopt a policy of experimentation and liberal toleration so that we can answer this question for ourselves. Specifically, we should perform experiments in living

so that we can learn what roles will allow us to maximize utility, what personas will allow us to play those roles well, and what personalities will allow us to adopt those personas well – and we should tolerate inner conflict if and when we experience it so that we can support these experiments as well as the ways of life that they guide us towards.

We can expect that this policy of experimentation and liberal toleration will have different results for different people, depending on our talents, interests, capacity for multiplicity, and other such individuating features. Some of us will learn from our experiments that we are most likely to maximize utility through a set of roles, personas, and personalities that are relatively fragmented: for example, you might decide that you need to be a CIA agent in order to maximize utility and that you need to be a spouse and a parent in order to support your cover – roles that place many conflicting demands on you socially and psychologically. Others of us will learn from our experiments that we are most likely to maximize utility through a set of roles, personas, and personalities that are relatively unified: for example, you might decide that you need to be a teacher and writer in order to maximize utility, and that you need to be a spouse and parent in order to find inspiration for your teaching and writing – roles that place *some* conflicting demands on you but which generally speaking benefit from your having similar thoughts, feelings, and habits across situations. Still, no matter how unified or fragmented we become as a result of our experiments in living, we will still have an indirect utilitarian obligation to tolerate inner conflict if and when we experience it, so that we can remain open to developing in new directions over time.

The second main reason we should accept a liberal system of individual self-government according to utilitarianism is that it allows us to resolve inner conflicts in a relatively fair way, with the result that we can enjoy many of the benefits of multiplicity while minimizing many of

the costs (insofar as we experience multiplicity at all). Of course, we need to say more about how liberalism works in the individual case before we can evaluate this claim. In particular, we need a substantive theory of what kinds of intrapersonal harms we should prohibit and what kinds of intrapersonal harms we should permit according to liberalism, in order to develop a concrete sense of how liberalism allows us to resolve intrapersonal conflict. I want to emphasize that I think that my general argument here is compatible with many different answers to this question. With that said, I will propose one possible answer to this question and show how this answer allows our personalities to work at cross purposes to a degree, while also allowing them to compromise, coordinate, and get along overall.

Here is my proposal: insofar as we have multiple personalities that disagree about how to live, our personalities have a duty to jointly commit to a fair set of laws (understood as a shared way of life that represents a fair compromise among them), and they have a right to bind each other, i.e., to manipulate, coerce, and physically restrain each other, only insofar as they need to do so in order to enforce these laws.

To see how this conception of individual liberalism works in practice, recall the examples that we considered in section 3. First, suppose that you tend to think, while at work, that you should take a promotion that will advance your career, and then you tend to think, while at home, that you should take a demotion that will allow you to spend more time with your family. In this case, a fair compromise between your work personality and your family personality might involve your keeping your current job. And if so, then we can say that your work personality is permitted to keep your current job (even though your family personality prefers the demotion) but is not permitted to take the promotion. Similarly, we can say that your family personality is permitted to use manipulation, coercion, or physical restraint to make your work personality

keep your current job (even though your work personality prefers the promotion) but is not permitted to use these methods to make your work personality take the demotion.

Similarly, recall the case where you tend to think, during the day, that you should stay out until around midnight tonight, and then you tend to think, during the night, that you should stay out until around 2am tonight. In this case, a fair compromise between your day personality and your night personality might involve your staying out until, say, 1am tonight. And if so, then we can say that your night personality is permitted to stay out until 1am (even though your day personality prefers that they come home by midnight) but is not permitted to stay out until 2am. Similarly, we can say that your day personality is permitted to use manipulation, coercion, and physical restraint to make your night personality come home by 1am (even though your night personality prefers to stay out until 2am) but is not permitted to use these methods to make your night personality come home by midnight.

As these examples indicate, this conception of individual liberalism is not opposed to self-binding, i.e. to our manipulating, coercing, or physically restraining ourselves, in all cases. Indeed, it endorses self-binding in most of the standard cases that philosophers and psychologists discuss in the literature. For example, if you have an addictive personality who can never have one drink without having ten, then you are justified in doing everything you can to keep them from having any drinks at all. Similarly (though this kind of case is less extreme), if you know that your work personality is likely to stay at work later than is fair to your other personalities, or that your family personality is likely to stay at home later than is fair to your other personalities, then you are justified in doing everything you can to ensure that they treat your other personalities fairly. So the point is not that we have to tolerate inner conflict in all cases, but rather that we have to tolerate inner conflict within certain limits, and that we are permitted to

engage in self-binding insofar as, and only insofar as, we need to in order to respect intrapersonal liberty and tolerate intrapersonal conflict overall.

Granted, this conception of intrapersonal liberalism will not always provide perfect solutions to our intrapersonal problems. For example, suppose that part of you thinks that you should get married and part of you thinks that you should stay single, or that part of you thinks that you should move to New York and part of you thinks that you should move to Texas. In these cases, you might be able to accommodate both personalities to a degree (for example, by having an open marriage or by moving to Texas but visiting New York), but the fact remains that you will ultimately have to favor one personality over another. But this is fine – indeed, it is inevitable for any political system. What matters is not that we identify a decision procedure that makes every personality equally happy but rather that we identify a decision procedure that considers the interests of each personality equally so that we can resolve conflicts in a fair and impartial way. And I believe that my conception of liberalism accomplishes that aim: it allows our personalities to work at cross purposes to a degree, yet it also ensures that they compromise, coordinate, and get along overall. And while our personalities might not always regard the plans that result from this decision procedure as *ideal*, they will nevertheless be strongly motivated to comply with them – and, insofar as this is true, they will be resilient in the face of indecision, compulsion, half-heartedness, and false belief – since they will recognize that, all else equal, they will be able to do more good in the long run if they reside in a “just soul” than if they reside in an intrapersonal “state of nature” defined as a war of each personality against all.

Finally, I think that, as in the collective case, we should apply this policy of liberal toleration to intrapersonal disagreement about utilitarianism itself. That is, we should tolerate inner conflict not only about how to maximize utility but also about whether or not we should be

trying to maximize utility in the first place. The reason is simple but important: it is possible that, in order for us to maximize utility through our roles, we will have to approach these roles with a kind of partiality and particularity that conflicts with the radical impartiality and universality inherent in the principle of utility. For example, you might justify being a spouse, parent, teacher, writer, and activist on the grounds that this set of roles allows you to maximize utility (as far as you can tell). But then in order for you to play each of these roles effectively, you might have to become a person who values them for their own sakes, rather than merely for the sake of overall utility. And once you value these roles for their own sakes, you will lose your commitment to utilitarianism, at least to a degree, in the course of playing these roles in everyday life.^{xviii}

To be clear, I am not saying here that we know for a fact that utilitarianism is self-effacing. That is, I am not saying that we know for a fact that all people, in all situations, will have to reject utilitarianism in order to comply with it. I am rather saying that we do not know for a fact that utilitarianism is *not* self-effacing for at least some people, in some situations. It is an open question what degree of belief in utilitarianism will allow us to maximize utility, in the same kind of way that it is an open question what kind of psychology will allow us to maximize utility more generally. And if this is right, then we should extend our policy of liberal toleration to inner conflict about utilitarianism too, so that we can create the social and psychological space necessary for performing our experiments in living, as well as for playing the roles, adopting the personas, and developing the personalities that these experiments guide us towards.

At least in theory, there is nothing wrong with this result. As Parfit says, the aim of a moral theory is to be correct, not to be believed.^{xix} But we do have to be careful in practice. After all, if we adopt liberalism on utilitarian grounds, then we might become so invested in our

projects and relationships that we lose sight of what justifies them in the first place. We might then let our attachment to these projects and relationships guide our thinking about how to live at least as much as, if not much more than, considerations of utility, with mixed results. As in the collective case, I think that this is a serious concern, but I also think that we should accept liberalism anyway, for two related reasons. First, I expect that, if we accept liberalism on utilitarian grounds, then the typical result will *not* be that we end up fully rejecting utilitarianism in everyday life. Rather, the typical result will be that we lose our commitment to utilitarianism to a degree in everyday life, and then we get it back to a degree in the cool hour of deliberation. Second, even if we do end up fully rejecting utilitarianism in everyday life, we are likely to remain committed to liberalism, since liberalism is compatible with many reasonable comprehensive conceptions of the good (which is, of course, part of its appeal); and our commitment to liberalism is likely to still push us in good, if not optimal, directions.

Can we guarantee that liberalism will be for the best? Of course not. As in the collective case, we have few if any guarantees about what will maximize utility. But we can say that, in light of the costs and benefits of multiplicity, we are more likely to maximize utility in the long run if we take a liberal approach to individual self-government than if we take an illiberal approach to it – even if a foreseeable consequence of this approach is that we will experience more inner conflict than we otherwise would, including inner conflict about utilitarianism.

6. Conclusion

I have argued that, according to utilitarianism, we should adopt a liberal system of individual self-government for the same reasons that we should adopt a liberal system of collective self-

government: a policy of tolerating inner conflict allows us to discover, and gravitate towards, the ways of life that allow us to maximize utility, and it also allows us to enjoy many of the benefits of multiplicity while minimizing many of the costs (insofar as we experience multiplicity at all). Moreover, I have argued that we should extend this policy of liberal toleration to inner conflict about utilitarianism itself: if we want to maximize utility, then we should tolerate inner conflict not only about *how* to maximize utility but also about *whether* to maximize utility in the first place.

If my argument here is correct, then it has surprising but ultimately plausible moral implications. For example, it places substantive limits on the morality of self-binding. I have developed these substantive limits here as follows: insofar as we feel conflicted about how to live, we have a duty to commit to a set of laws that represents a fair compromise between our different, and conflicting, perspectives, and we have a right to manipulate, coerce, and physically restrain ourselves insofar as, and only insofar as, we need to do so in order to enforce these laws. This constraint on self-binding is surprising because it shows that we have a moral obligation to treat ourselves with the same kind of toleration that we have a moral obligation to treat others with. But it is ultimately plausible because it is compatible with the idea that we should still manipulate, coerce, and physically restrain ourselves in order to, say, prevent an addictive or destructive personality from harming the rest of the intrapersonal community – and this is, of course, the kind of self-binding that philosophers most often discuss and endorse.

In closing, I would like to say a word about the scope of my argument in this paper. As I said in section 5, I have focused on utilitarianism here for the sake of simplicity and specificity, and because philosophers are already familiar with the idea that utilitarianism supports liberalism in the collective case. But while some of the details in my discussion have depended on my

choice of utilitarianism in particular, I think that the general form of my argument applies to many other consequentialist theories of morality and rationality as well. Stated in these more general terms, my thesis is that a liberal system of individual self-government allows us to discover, and gravitate towards, the ways of life that allows us to take the means towards our ends (whatever those ends happen to be), and it also allows us to enjoy many of the benefits of multiplicity while minimizing many of the costs (insofar as we experience multiplicity at all). Of course, I would have to say much more in order to establish this more general result. And there might be some theories of morality and rationality, such as deontological theories or consequentialist theories that value unity for its own sake, to which this general argument will not apply, or will only partially apply. (Though I think that deontologists should accept liberalism too, for reasons that I discuss elsewhere.^{xx}) But insofar as my argument does generalize, my conclusion about the morality of self-binding will generalize as well, and we will have to think in much more detail about how the ideal of intrapersonal liberty interacts with other ideals in morality and rationality including personal autonomy, responsibility, sovereignty, and more.

ⁱ Throughout this paper, I will use ‘comprehensive conception of the good’ to refer to a comprehensive conception of how to live. For more on this use of the term, see John Rawls, *Political Liberalism* (New York, 2005), pp. 19-20.

ⁱⁱ Erving Goffman develops this dramaturgical theory of social self-presentation in *The Presentation of Self In Everyday Life* (New York, 1959), R. Keith Sawyer, *Creating Conversations: Improvisation in Everyday Discourse* (Cresskill, NJ, 2001), Lee Strasberg, *A*

Dream of Passion: The Development of the Method (New York, 1988), and J. David Velleman, *How We Get Along* (Cambridge, 2009).

ⁱⁱⁱ To be clear, I am using ‘multiple personality’ in a wide sense that includes, but is not limited to, dissociative disorders such as dissociative identity disorder. I will not consider these much more extreme forms of multiplicity in this paper.

^{iv} See, for example, Richard Holton, ‘Intention and weakness of will’, *The Journal of Philosophy*, 96:5, 1999, pp. 241-262 as well as Frank Jackson, ‘Weakness of will’, *Mind*, 93:369, 1984, pp. 1-13.

^v See, for example, Donald Davidson, ‘How is weakness of the will possible?’ in *Essays on Actions and Events*. Oxford, 1969/2001 and Gary Watson, ‘Skepticism about weakness of will’, *The Philosophical Review*, 86:3, 316-39, 1977.

^{vi} See, for example, Harry Frankfurt, ‘Identity and Wholeheartedness,’ in Ferdinand David Schoeman (ed.), *Responsibility, Character, and the Emotions: New Essays in Moral Psychology* (Cambridge, 1987).

^{vii} Derek Parfit, *Reasons and Persons* (Oxford, 1984), p. 301.

^{viii} See, for example, Michael Bratman, ‘Practical Reasoning and Weakness of the Will’, *Noûs*, 13 (1979): pp. 131–51; Harry Frankfurt, ‘Freedom of the Will and the Concept of a Person’, in *The Importance of What We Care About* (Cambridge, 1988), pp. 11-25; David Shoemaker, ‘Caring, Identification, and Agency’, *Ethics*, 114(1) (2003): pp. 88–118; and Gary Watson, ‘Free Agency’, *Journal of Philosophy*, 72 (1975): pp. 205–20.

^{ix} Harry Frankfurt, ‘Identification and Wholeheartedness,’ *The Importance of What We Care About* (Cambridge, 1988), p. 175.

^x For more on experiments in living, see John Stuart Mill, ‘On Liberty’ (New York, 2004), p. 59.

^{xi} For more on dissociative identity disorder, see Deborah Haddock, *The Dissociative Identity Disorder Sourcebook* (New York, 2001).

^{xii} For evidence that multiplicity helps us to compartmentalize harm in this way, see A. G. Morgan and R. Janoff-Bulman, ‘Positive and Negative Self-complexity: Patterns of Adjustment Following Traumatic and Non-Traumatic Life Experiences,’ *Journal of Social and Clinical Psychology*, 1994, No. 13, pp. 63-85 and P. W. Linville, ‘Self-Complexity as a Cognitive Buffer Against Stress-Related Illness and Depression’, *Journal of Personality and Social Psychology*, 1987, Vol. 52, pp. 663-76. For general discussion of this issue, see Rita Carter, *Multiplicity* (New York, 2008), pp. 74-90.

^{xiii} John Stuart Mill, *On Liberty*, Chapter 2.

^{xiv} Of course, this is not to say that that the state should tolerate *all* comprehensive conceptions of the good. Rather, the idea is that the state should tolerate all *reasonable* comprehensive conceptions of the good, where a comprehensive conception of the good is reasonable in the relevant sense if and only if it supports a free and open exchange of ideas in society. See John Rawls, *Political Liberalism* (New York, 2005), pp. 48-54 for more on this idea.

^{xv} To be clear, this is not to say that utilitarians support the same kind of liberalism as deontologists. For example, as David Brink points out, one difference between Millian liberalism and Rawlsian liberalism is that Millian liberalism is perfectionist whereas Rawlsian liberalism is not. David Brink, “Mill’s Moral and Political Philosophy,” *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/mill-moral-political/>

^{xvi} For good contemporary discussions of this idea F. Berger, *Happiness, Justice and Freedom: The Moral and Political Philosophy of John Stuart Mill* (Berkeley, 1984) and Katarzyna de Lazari-Radek and Peter Singer, ‘Secrecy in Consequentialism: A Defence of Esoteric Morality,’ *Ratio* XXIII (2010), pp. 34-58.

^{xvii} How can we manipulate, coerce, and physically restrain ourselves? Ulysses arranged for people to bind him to the mast, but we have many other, less extreme options available to us as well, which philosophers and psychologists discuss in the literature on self-binding. For example, if you want to make yourself get up in the morning, you can place your alarm clock on the other side of the room. If you want to prevent yourself from gambling (or at least pressure yourself not to do so), you can pledge to donate money to a charity you oppose if you get caught gambling. And so on. As I will note below, liberalism and illiberalism agree that self-binding is often morally permissible; what they disagree about is whether or not we should use self-binding in order to impose a single, unified set of beliefs and values on ourselves across situations. For much more on the nature and rationality of self-binding, see Jon Elster, *Ulysses Unbound* (Cambridge, 2000).

^{xviii} For more on indirect utilitarianism, see David Brink, ‘Utilitarian Morality and the Personal Point of View,’ *Journal of Philosophy*, 83:8 (1986), pp. 417-48 and Peter Railton, ‘Alienation, Consequentialism, and the Demands of Morality,’ *Philosophy and Public Affairs*, 13:2 (1984), pp. 134-71.

^{xix} Derek Parfit, *Reasons and Persons* (Oxford, 1984), p. 24.

^{xx} Jeff Sebo, “The Just Soul,” *The Journal of Value Inquiry* (forthcoming).