

Diversity and the Social Contract

Ryan Muldoon

April 2009

Introduction

Motivation

Economic and political globalization has encouraged waves of immigration that have changed the social composition of countries, leaving once-homogenous cultures to deal with their newfound diversity. Unfortunately, Western countries have not handled this ongoing transition with much success. As Western Europe has seen immigration rise, it has seen support for liberal policies decline, and civil unrest increase. Liberal attitudes in Amsterdam are giving way to more restrictive policies. Paris has seen large-scale riots, first in the ethnic-minority banlieue, then among the more privileged students who resisted policies aimed at resolving the original tensions. In the United States, even with its longer history of ethnic diversity, immigration and affirmative action polarize political discourse.

I contend that these poor responses to new demographic realities are understandable, and even entirely unsurprising consequences of the classic liberal accounts of social contracts. Thanks to Rawls, social contract theory has become the dominant framework for considering questions of justice, but as it has been practiced so far, it assumes fairly homogeneous societies. This was, for much of Western history, a very reasonable assumption. When Rawls' *A Theory of Justice* was published, Western societies had manufacturing economies, and were relatively

homogenous, at least culturally if not always ethnically. As we go back to previous accounts of social contract theory, ethnic homogeneity only increases.

Though the crux of social contract theory has always been about how we are to reconcile our competing political interests, I submit that from Locke forward, the implicit assumption has been that the depth of political disagreement was not that great.¹ Likewise, accounts of tolerance, like Locke's *Essay on Toleration*, presented arguments for second-best solutions to religious disagreement: namely, that we cannot impose our views on others through the power of the state because they might one day be able to impose their views on us. And since we do not want to fight religious wars, we ought to tolerate the existence of other views.² This view assumes two things that I think are relevant to our understanding of current problems. The first thing to note is that diversity is a problem to be managed. There are no benefits seen. The second thing to recognize is that managing this diversity imposes a burden on individuals in society. Each person has to put up with those with whom they disagree.

Today, especially in Europe, the demographic story is much different. Not only has the nature of the economy changed drastically, but the composition of society has also changed. These countries, most with low birth rates, sought to

¹ In Hobbes, disagreement was essentially absolute: without a contract, we would be in a state of war. Locke dramatically weakened this assumption, and his account of the state of nature included individuals with recognized property rights, and people living successfully together. Rawls explicitly states that he is not going to argue with non-liberals.

² Though not all views. Locke claimed that atheists could not be tolerated.

bring in workers from other countries to maintain their welfare states. Those workers have stayed, and have now become sizable portions of the population, who expect rights like other citizens, including the freedom to continue their own cultural practices. And this is where we see the Lockean account of toleration fail. If we are to understand diversity as a burden, and disagreement becomes larger than the difference between Catholics and Protestants, then eventually people simply stop thinking that the burden is worth it. When this happens, we can expect to see policies enacted that aim to prevent new immigration inflows, and policies that either isolate minority groups or force them to take on majority culture. This is exactly what we see happening throughout the West.

This path is not sustainable. Economic logic dictates that more immigration will happen so long as richer countries have low birth rates, or economic globalization is maintained. But current political thinking has dictated that the resultant diversity creates political tension and eventually social unrest. And so, with the status quo, we reach an impasse: either wealthy countries drastically change their social contracts to contend with an aging, increasingly poor population, or they face civil unrest. I want to resist this false choice.

The aim of this dissertation is to provide an alternative account of social contract theory that explicitly aims to address the unique problems and opportunities that are found in diverse societies. In particular, it establishes three key ideas that differ from traditional social contract theory. First is the idea of *perspective* in political reasoning. Rather than assume that all citizens view moral

and political problems in the same way, it articulates a way to productively aggregate disparate views to find potential areas of agreement. The second key idea is my development of a *bargaining* model that requires virtually no agreement amongst its participants. And finally it presents a model of how agents can be *motivated* to seek out more diverse societies, even in the face of distributive disputes. These components combine to create a theory of social contracts that embraces Mill's concept of "experiments in living," by presenting a system of constant social experimentation and refinement. Rather than articulate an ideal end state, my focus is on the trajectory of change. Let us now briefly consider each of these ideas in turn.

Perspectives

Political disagreement is often cast as two or more parties having different preferences about which outcome out of a range of possible alternatives they are most interested in. Different people often have different political interests, and so this is an important idea to capture. But it does not capture even deeper sources of disagreement: the cases in which the parties do not agree on how to characterize the situation in the first place. Before we can even discuss what our competing preferences might be, we must first establish whether we agree with how to represent the problem. This issue of representation can be understood in terms of our perspectives. Perspectives do two things for us: they provide us with a

categorization of the world, and they provide us with a method of navigating that categorization. As societies become more diverse, we should expect that the range of perspectives found in the population will only grow.

On the face of it, this would seem like yet another kind of burden that diversity places on society. But I argue that different perspectives have epistemic advantages in different kinds of situations. There is no single perspective that is always best. Having a range of perspectives in the population can help a society solve more problems more effectively. In particular, we can take advantage of the fact of our diversity of perspectives to better assess our moral beliefs to help create the initial foundation of agreement in a social contract.

Bargaining

The notion of perspective is the concept around which the rest of the dissertation is organized. If we are to take perspectives seriously, there are several areas of social contract theory that must be adjusted. In particular, a focus on perspectives sheds light on why a public reason model of political deliberation cannot be used in a diverse society. Public reason requires the existence and adoption of a “neutral” perspective. I contend that such a perspective does not exist. Even worse, the public reason model limits what counts as acceptable in the public sphere. This model has good intentions, as it aims to allow only “shared” or “neutral” arguments, but I argue that it unfairly constrains individuals who have perspectives

that are not in the mainstream. Their views must be translated somehow into another framework of ideas before they can even be expressed. This almost assuredly reduces their political voice.

In place of public reason, I argue that a bargaining model of rights allocation in a social contract is most appropriate. Bargaining has several advantages. Most important among those is that bargains do not require a universally agreed-upon perspective. In fact, I show that a bargain can take place even where individuals in the bargain do not agree on what it is they are bargaining about. All each agent needs to concern herself with is her perception on whether she is gaining or losing. On this model, all perspectives are treated equally, as an agent's perspective is irrelevant to how other agents engage with her. But each agent is able to express her interests by making demands and concessions over the course of the bargain, even if agents with different perspectives might interpret those interests differently. In this way, we can arrive at an allocation of rights that leaves everyone with the same level of franchise, regardless of their perspective.

No-Subsidy Principle

The final component of the view of social contracts developed in this dissertation illustrates the conditions under which diversity can be actively embraced in society. Just as perspective diversity is epistemically beneficial, it is also economically beneficial. In particular, I argue that perspective diversity

encourages increased division of labor, and the resultant gains to trade create a larger social surplus. This fact ought to motivate increased diversity in society. In particular, agents have incentives to seek out people who are different from them to get larger gains to trade.

This positive story about the benefits of diversity to economic production is constrained by the very real disagreement over the distribution of the social surplus created by the increased production. Here, we ought to expect that as diversity increases, disagreement will also increase. I argue that the way around this source of tension is to introduce the “no-subsidy principle,” which is a distributive rule that assures that each coalition in society benefits from the social surplus created in proportion to their contributions to production. The goal of this principle is to assure social stability: when it is satisfied, no coalition is able to claim that they could be made better off if they left to form their own society. It assures that everyone shares from the system of social cooperation that is made possible by living in society. Though no one gets to dictate how the entire social surplus is distributed, each is assured that they get more of what they want than they would have in a less diverse society. It is this guarantee that assures that agents are rationally motivated to embrace diversity.

Putting it Together: Justice as a Trajectory

The account of social contract theory that I have developed explicitly avoids making claims about what the end result of the procedure I outline would look like. This is because ultimately my view is that we do not yet know what an optimal social contract would be. Not only that, but even if we found a social contract that was optimal for a given set of economic conditions and social demographics, that contract may become less optimal as the underlying social conditions change. Rather than endow a given end state with the status of a regulative ideal, I think it is more important to focus on the procedures by which a society can change its contract.

The three positive chapters of the dissertation should be seen as outlining an iterative procedure. We first attempt to assess our most robustly supported moral beliefs. Once we agree on what we view as simply unquestionable, we move on to what we can negotiate about. This negotiation, constrained by our previous step, determines an allocation of rights that we all can endorse. Finally, we find the implications of these rights by actually living with them and attempting small-scale experiments in living. As we gather more information about how accurate our assessments were of our support of our moral beliefs, and how well we did in determining a particular rights allocation, we can decide to go back and refine our social contract with our new information. Each stage of the process is designed to

ensure that no one is knowingly made worse off than they were before. If some change is found to not work, we simply try again. Every change that we as a society make is a refinement of how we think we ought to live with each other. Rather than stick with a social contract that was designed for a different era, or plunge in an entirely new direction, we should experiment. We can find what works, and abandon what does not.

Chapter 1:

Diversity and the Public Sphere in the Liberal Tradition

Introduction

A triumph of the Western Liberal political tradition is the enshrinement of the division of the public and private spheres. This distinction, aided by the development of individual rights, specifies the limits of the state, and the extent of the freedom of individuals. That is, it determines the areas in which the state has a legitimate claim to an interest in controlling or regulating the behavior of individuals or associations for the purposes of promoting the public interest. These areas are considered to be the public sphere – the realm of the political. Everything else is meant to be largely outside of government meddling. Here the interests of private citizens are supposed to freely manifest themselves: people can associate with whomever they please, say what they want, and for the most part pursue whatever personal life projects that they might have. The private sphere is thus the realm of individual interests.

In the private sphere, individual diversity is tolerated, and in a number of respects, considered to be a good thing. We are free to have our own musical tastes, our own preference for group association, and we are free to choose our profession and our hobbies. The only limits on our behavior are generally prohibitions on those actions that harm others. As the common example goes, my right to freely

punch stops at someone else's face. But this is not a particularly onerous restriction. As Rawls puts it, we are free to develop our own conception of the good, and pursue it without interference.

Microeconomics can be seen as a framework for studying our private-sphere interactions. As rational agents, we are free to make contracts and conduct trades with others to everyone's benefit. Individual differences in preferences, skills and knowledge create opportunities for markets to emerge and wealth to be created. Rational agents do not have to agree on much in a market: it does not matter what race or religion one is, or what one's political views are for there to be a successful transaction. The only relevant feature in a market interaction is whether or not one can agree to the price of a good or to the terms of a trade.

In the public sphere, however, there is a much greater need for wide agreement. Since the public sphere encapsulates those things that the state has an interest in regulating, a theory of participatory government may have difficulty with the kind of individual diversity that is encouraged in the private sphere. In the Western Liberal political tradition, however, thinkers like Locke, Rousseau and Rawls have contended that rational agents will come to agree on public sphere issues by means of reason alone. There is an apparent tension here - after all, both Locke and Rawls promote individual freedom in the private sphere, and so one would expect that these private sphere differences would have an effect on the public sphere. Instead of arriving at this conclusion, however, they argue that one's private preferences should not manifest themselves in the public sphere. One's role

as public citizen is distinct from one's role as private agent. It is further contended that rationality will lead citizens to agree on the shape of the public sphere. So, the ideal public sphere, which they claim is achievable, is one that all agree upon. This ideal of universal agreement can be seen in Rousseau's distinction between the General Will and the Will of All. The General Will is the agreement of individuals thinking of themselves as citizens on matters of the public sphere, whereas the Will of All is simply the union of the population's private sphere preferences.³ The Will of All might have a great deal of inconsistent preferences, reflecting a great deal of disagreement, but the General Will is explicitly defined to be that which we all agree upon as citizens. Rawls performs a similar kind of exercise in the Original Position, where one is hypothetically behind a thick veil of ignorance. One is meant to be able to characterize one's preferences for the basic structure of society without having any knowledge of one's private-sphere self. Because of this setup, Rawls argues that the reasoning behind the veil of ignorance will be the same for all individuals – therefore, regardless of private-sphere diversity, rationality leads to universal agreement in the public sphere.

This sharp boundary between public and private interests, and the lack of influence of private interests on public interests seems to be descriptively mistaken and theoretically worrisome. For instance, our private religious beliefs can have a

³ Rousseau's social contract, unlike Locke's or Rawls', does not protect private-sphere interests. Rather, the private sphere is overwhelmed by the public sphere as guided by the General Will. I raise this as an even more extreme example of the general phenomenon of establishing the public sphere with little regard to the private sphere. Mill is perhaps the archetype of private-sphere liberty, but I save that discussion for the next chapter.

direct influence on our public attitudes, as can be seen by Amish and Mormon rejections of some American political authority. Less extreme cases include nature-lovers fighting for national parks, or pacifists arguing against a large military. Not only are we in fact influenced by our private sphere versions of ourselves, but I also contend that this influence can be an avenue for social improvement, one that the rigid public-private split fails to recognize.

Two examples of theories that require a form of state intervention in the private sphere, and thus support from the public sphere, are perfectionist or communitarian comprehensive moral doctrines.⁴ Such theories cannot be circumscribed to the private sphere: while they may shape individual values, they have public sphere requirements. For example, a communitarian doctrine requires that social institutions be designed in such a way as to support community connectedness, and a subordination of the individual to the community. Community history and cultural standards must be maintained, potentially to the restriction of options of the individuals in the community. These are necessary for maintaining private beliefs in communitarianism, but require cooperation from the public sphere.

The possibility that “private” moral beliefs could have public consequences is ignored by the attempts to maintain the sharp boundary between spheres. Yet, we

⁴ Here I take Perfectionism to be a doctrine that makes claims as to what the objective human good is, along with a framework for determining social and political mechanisms for optimally promoting the good. Communitarianism I take to be a rejection of liberalism’s emphasis on the individual in place of an emphasis on community structures.

cannot meaningfully discuss the value of such theories without reference to features of the public sphere. This divide between public and private spheres represents not only a restriction of the public sphere, but implicitly, it also represents a restriction on what is allowed in the private sphere. This move goes against one of the avowed goals of the public-private split: that beliefs and activities in the private sphere are unrestricted, so long as they follow the harm principle. As doctrines such as perfectionism or communitarianism are not obviously harmful or otherwise illegitimate, they can easily be accepted in the private sphere. But we know that they would have significant effects on the public sphere. This very fact violates the separation of the public sphere from the private sphere so central to the liberal Western tradition.

One may choose to argue against the legitimacy of certain moral doctrines that have public sphere requirements, and thus try to save the boundary. However, even this move is problematic. Indeed, shifting moral doctrines that cast doubt on the public – private cleavage to the public sphere is a very dangerous move – it deprives us of a whole set of values that we may want to guide our lives. The value of the private sphere is precisely its being like a laboratory where we can test values and their viability. Unless some of these values are harmful to others, why should we create artificial boundaries that exclude entire realms of possibilities?

For example, why should the value of social connectedness, so central to communitarian doctrine, be left out of the private sphere just because it has political consequences? Our non-intervention then becomes a kind of intervention by

prohibiting certain kinds of social arrangements, while leaving others open. It should be noted that a major consequence of creating sharp boundaries between spheres and transferring all those private values that involve a public commitment to the public sphere is that we severely curtail our chances of succeeding in building satisfactory theories of public institutions. Without having an opportunity for private-sphere experimentation, we cannot determine what institutions most effectively serve to help people achieve their life projects.

One very different, and radical approach to solving the problems that the public – private split creates is to eliminate the notion of the public sphere altogether, and attempt to scale up private sphere interactions to deal with large-scale problems. One example of this strategy is to use a purely market-based solution to social needs and problems. This move collapses the public into the private, but, as I shall discuss further, this collapse cannot possibly be completely successful. There are elements of this kind of solution in Libertarian approaches to social organization. In this case too, the problem of providing public goods and dealing with externalities⁵ raises its ugly head—having some means of providing national defense or a system of justice is a basic challenge for full Libertarian or

⁵ To make my point clear, let us briefly define “public good” and “externality.” A public good is a good that is non-rival and non-excludable. That is, once the good is provided, no one can be kept from enjoying it, nor does anyone’s consumption of the good reduce the amount left for everyone else. National defense is a classic example of a public good. An externality is a cost or benefit related to some good that is not captured in the price of that good.

anarcho-capitalist theories. Nozick, for example, relies on a notion of a minimal government in order to satisfy the provisioning of these basic public goods.⁶

It is not strictly the fault of Libertarian or anarcho-capitalist thinkers that they have not been able to develop means of public goods provisioning without some government intervention – it is a basic problem in microeconomics. Mancur Olson in *The Logic of Collective Action* and Russell Hardin in *Collective Action* outline the significant challenges to public goods provisioning in large groups. One of the main difficulties is that markets do not necessarily ensure that all the costs and benefits associated with production of some good are incorporated into its price – that is, not all costs and benefits can be *internalized*. For example, factories tend not to pay for the social costs of polluting rivers or the air without some kind of additional government regulation. This is due to the fact that the costs are spread widely, while the benefits of polluting are kept by the factory in the form of lower manufacturing costs. So even though the individual benefits are greatly outweighed by the collective costs, the polluter's share of the costs is overwhelmed by the benefits. This problem of incentive-compatibility arises for all problems of collective action. In fact, it is so severe that even a population of altruists will run into the same problems, as any individual's contribution is not sufficient to provide the good, and so the altruist will not give without some assurance that others will do the same.

⁶ See Nozick's *Anarchy, State, and Utopia* (1974) for an elaboration of his account of the minimal state in a Libertarian framework.

A strategy that supports collapsing the public sphere into the private sphere is to “internalize” the externalities by creating new property rights, along with markets in which we can buy and sell these new rights. The Coase Theorem details how this incentive compatibility issue could be resolved if we were willing to devise a system of “complete” individual rights, such that individuals owned, say, vouchers for breathing fresh air. If the transaction costs were near-negligible, then the problem of collective action would go away, as there would be ample opportunities for the factory to simply purchase the fresh air rights of those affected.⁷ However, we do not have such fine-grained allotments of individual rights. So we find ourselves with no means to make individual incentives compatible with public goods provisioning. Because of this problem, the standard economic conclusion is that individual efforts to provide public goods will fail. Individuals are always motivated to “free ride” on the contributions of others, and not contribute to public goods themselves.

So we find ourselves in a position where taking a purely private sphere approach to organizing society has fundamental difficulties, and allowing the public sphere to ignore the features of the private sphere is unmanageable. Instead of either approach, we can aim for a novel path between these extremes. We retain a public-private distinction, but instead of allowing the public sphere to be determined in isolation from the private sphere, we show that it can be actively

⁷ For a more detailed discussion of the Coase Theorem and its implications for allocations of rights, see Coase's “The Problem of Social Cost”, *Journal of Law and Economics* (October 1960) or *The Firm, The Market, and The Law* (1990)

shaped by it. The most important, and indeed motivating, consequence of this is that the diversity that is fostered at the private level is retained at the public level as well. Rather than being a problem to be mitigated, this diversity can be used to better refine a system of social organization and allow it to adapt to the changing needs of a society.

The goal of this chapter is to examine the view of the public sphere as found in social contract theory and compare it to the analysis of diversity that is found in economic thought. I contend that the approach taken by social contract theorists is able to capture only some negative aspects, and not any benefits, of diversity, as it is analyzed in economics. I first clarify what I mean when I use the word “diversity.” Upon establishing the definitions I require, I begin the substantive analysis. I use Rousseau’s notion of the General Will, and Rawls’ notion of the Original Position, as major examples of the style of reasoning common to the Western Liberal tradition with regards to the public sphere. I then examine several ideas from economic theory to argue that there are a number of private sphere benefits to diversity. I leave to later chapters the work of fully accounting for how these benefits manifest themselves. Finally I conclude by arguing that the economic analysis presents problems for the approach taken by the Western Liberal tradition.

Diversity of What?

A term like “diversity” is in some sense ill-suited for the task ahead. By itself, it does not identify a precise concept, as it is a relational term without any particular

restrictions on its domain. As such, we are left with a term whose sense is an amorphous mixture of related concepts. In my view this fuzzy meaning led to a use of “diversity” in public discourse germane to a notion of cultural or identity diversity. I contend that identity diversity is not a basic concept, and must be broken down into its constituent parts, to better understand cultural diversity, but also to see how the basic parts interact with each other.

To do this, we must consider the ways in which individuals in a society can be diverse. Once these distinctions have been made, we can then see the multi-faceted ways in which these kinds of diversity can interact with each other. It is through an understanding of these interactions that many of the problems that the social contract faces can be solved.

Identity Diversity and Why Groups are not the Unit of Analysis

Most readily apparent to many of us is Identity Diversity. People identify with different racial, social, and gender groups, and these distinctions are frequently used to divide a population. However, identity diversity is also a shifting notion – over time, some identities fall away, and new ones appear. Race, ethnicity, and social groupings are not natural kinds – there are no principled, objective ways of defining borders between most groups. One might attempt to create distinctions between social groups by means of shared histories: individuals whose families are from the same geographical location, and have significant rates of inter-marriage

might be grouped together.⁸ Though a technique like this is not going to produce sharp boundaries, let us suppose that it could, at least, provide rough groupings. The fuzzy borders between groups may cause problems, but we can at least assume that for many groups, there are sharp boundaries between almost any two of them. So, let us suppose that for ethnic groupings we can find an objective standard by which we can differentiate between most pairs of groups. Even with this, we are not yet out of the woods: while this technique may be successful at differentiating Irish from Italians, it does not seem to work for separating Capitalists from Communists, Catholics from Protestants, Unionists from Loyalists, or Punks from Goths. Social groupings may often have some positive correlation with genealogy, but for many significant groups, it was the spread of ideas that created distinctions. As such, it seems unlikely that there exists an objective distinction that can be used to set boundaries between groups.

Without something like an objective method of distinguishing groups, we are unable to speak precisely about identity groups as a class and describe their properties and dynamics. For example, there does not seem to be much in common among the natures of the Red Sox Nation, Republicans, lesbians, academics, and the British. Each provides some cultural identity, but to a greater or lesser degree. Red

⁸ We already know from biology, however, that this is going to be a fruitless task. There is significant literature arguing against the concept of a species (or at least reducing a species to an individual). For a sample of this literature, see Mishler, B. and Donoghue, M., 1982, "Species Concepts: A Case for Pluralism", *Systematic Zoology*, 31: 491-503 and Ereshefsky, M., 1998, "Species Pluralism and Anti-Realism", *Philosophy of Science*, 65:103-120. If this cannot be done for species, identifying particular cultural groups, a much finer-grained distinction, has very little chance of success.

Sox fans are unlikely to push for political concessions, beyond perhaps protecting Fenway Park from destruction, whereas other groups are much more inclined to push for political concessions. On the other hand, Red Sox fans have a great deal of shared culture and demonstrate in-group behavior even with strangers. Red Sox fans probably publicly gather at greater frequencies than, say, Republicans or academics. But even so, being a Red Sox fan has only limited, if any, spillover into one's public sphere attitudes, whereas the political goals of lesbians, Republicans or academics are much more pronounced. This example might encourage us to think that we can at least make distinctions based on how much groups address the public sphere, but this also fails. Consider the case of the Sierra Club. The Sierra Club was originally an organization whose purpose was to organize outings for people who loved the outdoors. It provided a means to plan trips, meet other people interested in the outdoors, and other similar functions. However, it evolved into an organization whose primary aim is to lobby the government to protect the environment by means of regulation. One can see how this could have happened: the membership was interested in outdoor activity, and from that, they gained an interest in protecting the outdoors to ensure their continued enjoyment of it. So they were able to join together to lobby the government to supply a public good for which they had private motivations. It is not inconceivable to imagine that Red Sox fans could eventually turn to political activism to perhaps lift immigration restrictions on Cuban or Dominican nationals, given their tendency to be good baseball players.

Groups can do more than shift from a private focus to a public focus.

Groups can disappear altogether, like the Whigs, or they can merge, as in some ways the European countries are beginning to do. Groups can also splinter into subgroups, once the group's membership starts recognizing additional distinctions that they find increasingly salient. The various strains of Christianity (or even just Roman Catholicism) are examples of this kind of splintering over time. This inconstancy of groups over time, not just of membership, but of focus and distinguishing features further suggests that treating groups themselves, and thus cultural identity, as the relevant unit of analysis in a discussion of diversity in the public sphere would be a mistake. While groups surely exist, they are both difficult to differentiate at an abstract level, and inconstant through time.

Though we are able to dismiss the concept of groups being natural kinds, this objective distinction does not preclude a subjective perception of groups as natural kinds. That is, it is possible that Muslims or Italians could view themselves as natural kinds. They could hold this view regardless of the objective facts of the matter. It is likely that the members of at least some group hold this view, and so we need to consider this and see what consequences it has for the analysis I suggest.

Fortunately, there is a straightforward response to this worry. Recall that we are trying to determine what the appropriate unit of analysis is in a social and political context. Even if the individuals have the mistaken belief that their group is in some sense natural, the abstract analysis of the social dynamic still ought to focus on individuals, for the reasons elaborated above. However, there should be some

mechanism by which we can describe groups that recognizes this subjective stance. Some groups will exist for longer periods of time, or perhaps have heightened attitudes towards in-group individuals, and this should be recognized. That is, individuals who consider themselves members of groups and further see these groups as fundamental are going to behave in ways that are different than individuals who see individuals as fundamental. But, this does not require the group to be the unit of analysis. We can instead provide an analysis of this difference purely at the individual level, as it is the individuals who see themselves in this way. What is needed, then, is the means to describe the differences in individual perspective, and to characterize some individual features as those that are indicative of membership in a certain group. In the following section, I develop an account of what properties of individuals we ought to consider the basic building blocks of diversity. It is with these same building blocks that we can construct groups, by treating them as portfolios of individual properties. These portfolios, as will be explained in more detail later, can be thought of as the constraints on an individual's attributes such that the individual could be considered a member of a particular group. A trivial example is that Red Sox Nation members would have the constraint of having to hate the Yankees and love the Red Sox. Without this constraint, one could not have much claim to group membership.

It is by means of the portfolio account of groups that we can still discuss groups, and strong individual feelings about groups, while still relying on individuals as the appropriate unit of analysis.

Individual-level Diversity

Though we have already established that we are not looking at group-level diversity, and instead are focusing on individual-level diversity, we have not yet examined the relevant kinds of individual diversity. Of course, there are many ways in which individuals are different – the challenge in this kind of analysis is to identify the differences that can account for much of the dynamics of both the private and public spheres. We have diversity of our preferences or goals in both a political and personal context. Further, even if we agree on the same goals – the desire to lower crime, for instance – we can often disagree over means to our ends. Finally, we have diversity in our abilities, perspectives, and our knowledge. We bring different skills to the same problems, and a problem can be easier or harder depending on how we approach it. Let us consider each of these relevant kinds of diversity in turn.

Preference diversity is fundamental to both economic and political models of individuals. In these models, individuals have what is called a preference ordering – a complete list of the goods or outcomes that an individual is interested in, in order of how strongly that person is interested. So, for instance, one can imagine a (partial) preference ordering for fruits, in which I prefer oranges to apples, apples to bananas, and bananas to peaches. Preferences are assumed to be transitive, which is to say that, given this ordering, I must prefer oranges to peaches. It is clear how

diversity can enter into the framework: while I might prefer oranges to apples, it is very easy for someone else to prefer apples to oranges. While a diversity of preferences over fruits is fairly innocuous, when preferences are over social outcomes, like which public goods to provision, or which social policies to promote, preference diversity can be extremely worrisome.

The literature on social choice demonstrates how much of a problem preference diversity can be, stemming from Arrow's impossibility theorem. Arrow's impossibility theorem assumes at least three individuals want to vote over at least three possible choices. We further assume four basic constraints on the type of vote that we are allowed to use. The first claims that there is no dictator who decides the outcome irrespective of others' interests. The second constraint is that the voting mechanism should be universal – it should take into account all of each voter's preferences, and provide the same voting result given the same set of preferences. This just is to say that there is a deterministic voting mechanism that includes all of the voter's preferences. These first two conditions just make it clear that we are interested in democratic voting. The other conditions are meant to impose some basic consistency requirements. The first is that if everyone prefers A to B, then the collective vote should also prefer A to B, as the vote should recognize the dominance of A to B. If everyone prefers apples to oranges, then if we vote on which fruit to buy, the vote shouldn't tell us to buy oranges. The second is that if the vote prefers A to B, if we then expand the option set to include C, the vote does not then rank B

above A. That is, if we add bananas to the mix, that shouldn't cause us to suddenly prefer oranges over apples.⁹

Even prior to Arrow, political theorists saw that preference diversity was a problem to be solved – Hobbes solved it with the introduction of the Leviathan, but most political philosophers since have attempted to deal with it by means of the public-private distinction. Individuals can differ all they want in what fruits they prefer (and in fact it is good for them to differ, insofar as it enables more people to satisfy their preferences due to less resource competition), but in the public sphere, the claim has been made by social contract theorists that our rationality and public-spiritedness will lead us to consensus in what public goods we ought to pursue. Because of its centrality to the problems that political philosophy has tried to address, preference diversity will also hold a prevalent position in the present analysis.

Preference diversity is sometimes confused with a distinct notion: diversity over the choice of means to ends. While preference diversity as I've defined it ranges over *outcomes*, there is also diversity over the *means* to achieve those outcomes. This is often overlooked in the analysis of political conflict, but is

⁹ For a more detailed discussion, see Arrow, "A Difficulty in the Concept of Social Welfare", *Journal of Political Economy* 58(4) (August, 1950), pp. 328-346. Gibbard and Satterthwaite have shown similarly damning problems wherein voting systems are either dictatorial, have a choice that cannot succeed, or are open to strategic manipulation. See Allan Gibbard, "Manipulation of voting schemes: a general result", *Econometrica*, Vol. 41, No. 4 (1973), pp. 587–601 and Mark A. Satterthwaite, "Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions", *Journal of Economic Theory* 10 (April 1975), 187-217.

essential to understanding when apparently conflicting preferences can in fact be reconciled. For instance, if one party is interested in investing resources in bigger prisons and more stringent laws, and the other party is interested in better education and more social workers, it may be that their outcome preferences are in line, but their preferences over means diverge. In this case, both sides might ultimately be interested in reducing crime. It is rare that someone has an interest in prison construction for the sake of it. If analysis can help demonstrate that the goals are in fact the same, the conflict can transition from an ideological one, which can often be intractable, into an empirical issue of which method has greater success in efficiently reducing crime rates. When goals agree, and preferences over means differ, the role for social experimentation becomes clear: it allows a society to resolve disagreements that are over means.¹⁰

Social experimentation finds support from an area other than preference diversity: it also finds support from skill diversity. Individuals have notable differences in skills, and these skills will lead them to engage in different professions and projects. Economic theory has had much to say about the benefits of skill

¹⁰ A challenge for treating instrumental methods as distinct from ends is that the familiar Aristotelian argument could be used to argue that all but one of our ends are actually best understood as instrumental means to a final end of happiness. In practice, I do not consider this to be a grave concern. Aristotle himself recognized the distinction between a good that was merely a means towards some other end, and a good that is also good in itself. The difficulty, then, is to identify these distinct cases. While many goods could be either an end or a means, the good in the context of its advocacy provides sufficient information to distinguish between the two cases. For instance, a prison would not be argued for on the basis of its aesthetic appeal, whereas a museum would be.

diversity, as we will see in later sections. A key question to investigate, however, is what the relationship is between an often-beneficial diversity of skills, and an often-contentious diversity of preferences.

Just as skill diversity is more clearly beneficial to society than preference diversity, the diversity of perspectives and knowledge is also an area where we can more readily expect benefits to accrue. Scott Page and Lu Hong have argued that perspective diversity – the difference in how individuals categorize the world and analyze problems – is beneficial to solving many kinds of difficult problems. In fact, they argue that diversity of this kind can trump individual ability in terms of its epistemic value.¹¹ A common example of perspective diversity is how one would go about categorizing a variety of foods. Different people often choose different methods of categorization. For instance, some may organize them by whether or not they are organic, some by country or region of origin, some by whether it is frozen, canned, or fresh, some by color, some by nutritional content, some by whether it is a fruit or a vegetable, some by price. Many other categories are possible. Most likely, no one would opt for all of these categories at once, but for any given category, some subset of the population would find it salient. For instance, a poorer individual might be inclined to focus on price, while a wealthier person might focus on whether food is fresh or frozen. These categorizations fit into a

¹¹ For a further elaboration of this argument, see Lu Hong and Scott Page, “Groups of diverse problem solvers can outperform groups of high-ability problem solvers” *Proceedings of the National Academy of Sciences*, November 16, 2004, and Page’s book *The Difference* (2007).

larger background of the kinds of problems that these individuals typically have to solve. If one is blind, for example, color distinctions are not particularly relevant.

Somewhat tied up with the notion of perspectives is individual knowledge diversity. Our background knowledge shapes and is shaped by our perspectives, but knowledge is still distinct. Some specialized knowledge is also an ability, such as knowledge of mathematics or physics, but much specialized knowledge is just knowledge of particular facts. For instance, knowing about the batting average of a particular baseball player, or knowing about a local shortage of machine parts. Knowledge can be of a fairly trivial nature, like baseball statistics, or have grave consequences, like the predictions of climate models. It can also simply be culturally enriching, like the knowledge of Mozart or Dostoevsky's works. This kind of knowledge aids in shaping perspectives, even those at a fairly high level. One's exposure to literature and music, for instance, is likely to help shape whether one considers oneself a humanist. Greater awareness of history is likely to result in a greater tendency to frame current events in the context of past events. Knowledge of the principles of economics might encourage a particular view of how individuals make decisions. Thus we find that knowledge diversity can have a significant role in understanding individual diversity.

Now that we have discussed the different axes on which one can consider individual diversity, we can briefly return to the common notion of identity or cultural diversity. Though we have seen that cultural distinctions are not "basic" in the sense that they are not appropriate units of analysis, it does not mean that we

have no possible means of discussing culture and cultural diversity. We can instead describe a cultural identity as a portfolio of preferences, skills, perspectives and knowledge. Under this model, cultures can be characterized by a partial ordering of preferences, a set of commonly associated skills and perspectives, and what cultural knowledge is essential to membership. These can be seen as constraints on individuals insofar as their individual identity must suitably coincide with these requirements in order to be able to lay claim to the relevant cultural identity. Even though constraints are imposed, a great deal of flexibility remains. This portfolio approach allows one individual to simultaneously belong to multiple groups, as groups are not complete determinants of individual characteristics. It also allows us to investigate cultural change, as we can easily examine the dimensions across which individuals might change over time, including what incentives might induce them to do so. Most important, however, is our ability to promote diversity without making the mistake that doing that entails ossifying current social categories.

The Social Contract Approach To Diversity in The Public Sphere

Social contract theory has a long and multifaceted tradition, but here I am focusing only on the role of diversity in the public sphere. Furthermore, I am considering only Rawls and Rousseau, and I will treat them as representative of the

larger tradition.¹² This is heavily restricting the subject matter, but as the goal here is not historical analysis, but instead conceptual analysis, focusing on representative claims ought to be sufficient.

Diversity as such is not a primary point of discussion in the social contract tradition. If one were to classify Hobbes as the first social contract theorist, at least in the Modern period, the claim could be made that Hobbes is an exception. He defines people as equals only in the sense that the weakest can always kill the strongest, and the least intelligent can still conspire against the most intelligent. In other respects, people are rather different from each other, and importantly, they are likely going to have differing interests. This difference of interests is so fundamental and so problematic that he shapes much of his theory around trying to mitigate the challenges it creates. Namely, the institution of a Leviathan must be imposed so as to offer people the means of offering enforceable assurances that they will not violently try to impose their will on others. In his analysis of the forms of government, a monarchy is determined to be the best kind of government because it is the only one in there is no possibility for civil strife caused by disagreements among political leaders.¹³ So a certain kind of diversity drives his view, and forces him to create political institutions whose focus is to regulate and mitigate the

¹² I do not focus on Locke here because his method of argumentation, both in the *Treatise* and in *A Letter Concerning Toleration*, relies on one's relationship to God and religion. This is not found in other social contract theories, and as such, there is less to be gained by using it as an example.

¹³ See *Leviathan* Book 2, Chapter 19 for a detailed discussion of how only a monarch can lead, and not an assembly. The focus of the argument is on the potential for disagreement and different interests.

negative power of diversity. After Hobbes, however, this kind of argument fell out of favor with social contract theorists. The underlying assumption of a kind of moral egoism was replaced with a notion of the Reasonable, which was found in Locke and all subsequent theorists. Once this move was made, the “threat” that diversity posed was largely eliminated. Because of this, I claim that there was no longer any need to focus on it as a theoretical worry – instead focus could be turned to the ways in which individuals are similar. This is not to say that most thinkers ended up claiming that everyone was identical. Rather, my claim is that post-Hobbes, diversity was viewed as largely irrelevant from the point of view of the public sphere. With Hobbes, diverse interests could clash with one another unless a social contract provided some method of adjudicating between them. After Hobbes, a social contract became a deliberation on the common good. This is a significant shift in the perception of the public sphere.

One might suppose that the story can break down when we consider something like the liberal value of tolerance, which is explicitly about what position public individuals ought to take in the face of diversity. However, arguments for tolerance, like Locke’s, do not argue for the value of a diversity of views, nor do they argue that diversity is negative. Instead they argue that attempting to force someone to, say, convert to one’s choice of religion undermines one’s religion. States are to stay out of religion for similar reasons – the threat that one’s preferred religion would be outlawed, or diluted by imposing it on non-believers. This kind of argument has nothing to do with diversity per se, as it is not an object of analysis.

Instead, it focuses on individual strategies for maintaining one's own beliefs. As such, accounts of toleration that take this form are not themselves a counter-example to my claim that social contract theorists consider diversity to be irrelevant to any considerations in the public sphere.

To provide further support to this claim about the nature of the public sphere in the social contract tradition after Hobbes, let us first examine Rousseau's distinction between the General Will and the Will of All, along with Rousseau's account of the role of the General Will in the state. Using this as a baseline, we can then move on to an analysis of Rawls' Original Position. I claim that both of these are examples of an attempt to establish a public sphere that is not grounded in the private sphere.

Rousseau and the General Will

The will of all, according to Rousseau, is the union of all of the private interests of the populace. This is, then, just an agglomeration of individual wants and desires, and not something that Rousseau considers relevant to the public sphere. The reason for this is that the perspective undertaken by individuals so as to create the will of all is not public-minded. The will of all is the union of private goods, not the public good. This is a crucial feature of the will of all for Rousseau, as the individual wills that comprise the will of all need not consider any others in

society, and may even wish goods that could be disadvantageous to others. The general will, on the other hand, requires a different perspective.

The general will, unlike the will of all, is made from the perspective of individuals as public-minded citizens, not private individuals. It is from this public perspective that individuals attempt to determine what is in the public good, separating themselves from their private interests. Instead, as Rawls argues in his lectures on political philosophy, they consider what their fundamental interests are, derived from their basic humanity. As everyone has these interests, the general will is something that one ought to be able to endorse due to the fact that the public good is then also individual good. Rather than the general will being a union of individual interests like the will of all, the general will can be understood as the intersection of public interests.¹⁴ Note however, that the difference is not merely the substitution of the intersection operation for the union operation in how we collect individual views. The individual wills being collected are of a different nature. The general will is operating with public-minded views of the public good based on our fundamental interests.

¹⁴ Rousseau here is looking to take advantage of something like the Condorcet Jury Theorem, or (as Rawls describes it) Bernoulli's law of large numbers: so long as any given individual is at least 50% likely to be correct as to the true nature of what the public good is, then the general will will prescribe it, provided that there is a large enough population. But for this to be applicable, there cannot be any large interest groups, but more importantly, there can be no communication between individuals. This strange requirement is due to the fact that Bernoulli and Condorcet rely on the independence of statistical samples. If individuals communicate, their assessments of the public good will become correlated, and then the theorems are not applicable. For Rousseau's usage of these theorems, see *The Social Contract*, book 2, chapter 3.

It is here that we can see the shift made by Rousseau that Hobbes does not make but Rawls will also make: individuals are to determine the general will not by their knowledge of their private interests, but by their knowledge of our fundamental shared interests as humans. But to some degree, that is what we are trying to determine in the first place. The public stance taken by the individuals cuts individuals off from their personal conceptions of the good, which is the only readily available source of information in understanding what our shared needs might be. A further question is at what level of analysis do we find the general will. In his lectures on Rousseau's political philosophy, Rawls claims that the general will is a method of public reason.¹⁵ Because of this, each individual has the general will, as well as her own private will. But this does not seem to be consistent with Rousseau's discussion of the general will. If the general will is to be determined by the intersection of individual public wills, then it is something different from any individual will, or at least the vast majority of individual wills. The general will is also described by Rousseau as being corporeal – it is something that becomes independent of the individuals that generated it: "The body politic, therefore, is also a corporate being possessed of...this general will...which is no longer that of any individual".¹⁶ This does not seem consistent with an analysis of the general will as a method of discourse.

¹⁵ See Rawls, *Lectures on the History of Political Philosophy*, Samuel Freeman, ed. (2007)

¹⁶ *Discourse on Political Economy*, pp.132, 140

However, the general will is meant to be used as the basis of the public sphere.¹⁷ It defines the conditions of justice, freedom and equality, and governs the relationships between individuals. But as we have noticed, the general will does not utilize any account of individual differences – it serves only to strip them away. The private sphere is explicitly rejected as a source of the public sphere, as the general will requires a public stance that rejects individual interests. The general will itself is then made to exist independently of the particulars of the individuals in society who created (or found) it.¹⁸ Here we find parallels to our discussion of the relevant units of analysis in our reasoning. There we rejected a notion of group-level analysis, and instead embraced individuals as the appropriate unit of analysis. But here we find the general will as itself operating at a group level of analysis: it is an entity independent of its members, and it is formed by deducing their fundamental shared interests.¹⁹ Its independence by itself would be sufficient to suggest that there is a problem, but let us also consider the notion of fundamental shared interests, and show why it also would be a cause for concern.

The challenge of fundamental shared interests is the assumption that there is a proper set of fundamental shared interests that cover all individuals as they are, does not leave any important interests of individuals out, and is sufficiently well-defined such that the set of interests is enough to account for the public goods that

¹⁷ *ibid*, p. 132

¹⁸ *ibid*, p. 140, *The Social Contract*, book 2, chapter 3

¹⁹ “As long as several men in assembly regard themselves as a single body, they have only a single will...” (*The Social Contract*, book 4, chapter 1) Also relevant are book 4, chapter 2, as well as book 2, chapter 4.

the state ought to provide. But to claim as Rousseau does that such a set exists is just to beg the question. The public sphere, after all, is meant to structure society in such a way as to accommodate our needs. But if there are fundamental differences in what our needs are, or even in what our understanding is of these differences, the public sphere is where this discussion needs to take place. The general will, which underlies the public sphere, relies on the idea that there is a single unchanging set of fundamental human needs. As the fundamental interests under discussion are those interests that we take on as members of a society, and not as the pre-social brutes that Rousseau imagines in the state of nature²⁰, it is even more plausible that different individuals, or even different sub-groups, have different fundamental needs in society to satisfy their conception of the public good.²¹ But if this is the case, then we at least have plausible grounds for questioning whether there is a set of fundamental shared interests. Given the existence of this question, then the

²⁰ “The passage from the state of nature to the civil state produces a very remarkable change in man, by substituting justice for instinct in his conduct, and giving his actions the morality they had formerly lacked. Then only, when the voice of duty takes the place of physical impulses and right of appetite, does man, who so far had considered only himself, find that he is forced to act on different principles, and to consult his reason before listening to his inclinations. Although, in this state, he deprives himself of some advantages which he got from nature, he gains in return others so great, his faculties are so stimulated and developed, his ideas so extended, his feelings so ennobled, and his whole soul so uplifted, that, did not the abuses of this new condition often degrade him below that which he left, he would be bound to bless continually the happy moment which took him from it for ever, and, instead of a stupid and unimaginative animal, made him an intelligent being and a man.” *The Social Contract*, Book I, Section 8, Paragraph 1

²¹ Since these needs are not food and shelter, but rather social needs, one would expect that these would at least track social, economic and technological changes. Rousseau’s doubt about this is understandable, as at the time he was writing, there had yet to be significant technological changes that fundamentally affected the quality of life. These did not arrive until after the industrial revolution and subsequent advances in farming. For evidence of this lack of social change, see Gregory Clark, *A Farewell to Alms*. (2007)

construction of the public sphere ought to then pay attention to the diversity of basic needs and interests that are represented in society.

So we see with Rousseau that the formulation of the general will is problematic on a few grounds. First, we find that the general will is detached from the individuals who will it, as we saw in Rousseau's corporeal description. This leads to difficulties very similar to those found in the notion of treating groups as the unit of analysis. Insofar as the general will aims to describe the public-minded concepts of justice for the group at large, we have the basic problem of establishing what the group is, and whether it is well defined over time. This can be made even more problematic if different individuals have different conceptions of the group, such that each member has a different basis for determining group membership. For example, some Republicans might claim that one has to be in favor of small government to count as a Republican, while others might think that Republicanism is instead better defined in virtue of one's conservative social attitudes. These two criteria will create overlapping, but non-identical groups. Neither group can rely on being the "real" group, as they exist only in virtue of individuals joining up to create them. This can be resolved by deciding one's group is everyone, since there cannot be any misunderstanding over group composition. But while the extension of the group is then clear, the intension of the group is not – many groups are defined contrastively, and if there is nothing to contrast against, it is hard to have an in-group without an out-group.

Even if we suppose that we can resolve the issue of treating groups as the appropriate unit of analysis, we find that the process of willing is also problematic. This can best be understood in terms of contemporary work in social psychology on social identity. Social identity theorists, such as Marilynn Brewer, have argued that our identity is more than just our personal identity, but several social identities as well. One might identify oneself as a philosopher, and see one's identity tied up in the successes and failures of other philosophers. Similarly, one could generalize further out, and see oneself as an academic. This generalization could keep going, or one could start from a new basis for a social identity, such as being an amateur singer.²² In this theory, we have many different versions of ourselves, and with any of the social selves, our identity is tied up in the group. As such, the group's interests are our own interests, because we see ourselves as in terms of our group membership. It is here that we find similarity with Rousseau: the public perspective taken in constructing the general will is akin to taking on a social identity. However, we have many social identities. How do we all choose the same one? Brewer claims that certain identities can be made more salient than others: going to the voting booth, for instance, will tend to encourage one's identity as a citizen, whereas going to the stock exchange will encourage one's personal identity. But Brewer also claims that one's choice of social identity will be such that one satisfies competing interests for individuality and a need to fit in. If this were to be true, then viewing

²² For a more detailed discussion, see Brewer, Marilynn, "The Social Self: On Being the Same and Different at the Same Time." *Personality and Social Psychology Bulletin*, 17, 475-482. (1991)

oneself simply as a citizen is not going to be successful, as there is no room for any expression of individuality.

In social identity theory, there are usually two ways in which a social identity can be established or used: one is that there is a common set of interests, or that there is an out-group that defines the in-group. In the case of common interests, one finds that the group's success correlates with one's own success. An increase in prestige for philosophy in general increases the prestige of individual philosophers. But this has some limits: in commons dilemma situations, social identification can help increase cooperation, but not as much as it would if people truly saw the group's interests as being their own interests. In the case of the existence of an out-group, individuals will tend to do things that increases the advantage of their group members – even if the group division is essentially meaningless, like whether one is an over-estimator or under-estimator of the number of dots on a piece of paper.²³ Given the success of group division even in these meaningless contexts, it does not seem as if group identity as such is what is driving the results. After all, no one is going to claim that over-estimators of dots share common interests that are different from the rest of the population.²⁴

²³ See Tajfel, H., Billig, M., Bundy, R. P., & Flament, C. "Social categorization and intergroup behavior." *European Journal of Social Psychology*, 1, 149–177. (1971)

²⁴ An alternative account for explaining cooperation and social coordination can be found in theories of social norms. Instead of claims about social identity, a norms-based approach relies on the idea that individuals determine that some behaviors are expected of them in the population, and they care about meeting those expectations, potentially because of a fear of social sanction. This approach is more

Since the general will is meant to encompass a nation, if not all of humanity, it is difficult to see how this process of willing can be done. The group is too large for social identity to be strictly meaningful, unless there is an out-group to contrast against. And it is not clear that each individual's conception of the group is the same. But more basically, with such a large group, the social identity may not even exist in such a way that encourages strict identification of one's interests with the group's interests.

Finally, the general will relies on a notion of fundamental shared interests that goes without elaboration or discussion. As the general will tries to discover what the fundamental shared interests are in the first place, it is illegitimate to assume away the possibility that there is no single set of fundamental human interests. Particularly as the size of the population grows, the lack of social identification will make common interests increasingly unlikely.

It is possible that this line of critique leverages the fact that the general will is a fairly complicated theoretical entity that relies on an equally complicated social process, rather than getting at something more fundamental about the methodology employed by social contract theorists in the construction of the public sphere. To

able to account for the empirical evidence, as well as avoid the conceptual issues raised above. For a critique of social identity theory from a norms-based approach, see Bicchieri, "Cooperation and Communication: Group identity or Social Norms?" in N. Gold [ed], *Teamwork: Multi-Disciplinary Perspectives* (2005). For a detailed elaboration of a theory of social norms, see Bicchieri, *The Grammar of Society* (2006)

resolve this, let us turn our attention to Rawls' social contract theory, and the mechanism of the Original Position in particular.

Rawls and the Original Position

Rawls, like Rousseau, faces the challenge of developing an account of the public sphere. Rather than trying to use something like the general will, which requires individuals to take a public perspective rather than an individualist perspective, Rawls imagines instead that the parameters of the public sphere can be developed by means of an idealized, hypothetical bargain: the Original Position.

The Original Position can be thought of as the culmination of the basic ideas in the history of the social contract tradition. That is, we find a notion of bargaining that is prevalent in Hobbes and Locke, but from Rousseau we have an account of impartiality and independence from the status quo. The Original Position side-steps the fiction of a historical state of nature that supplies the justification for certain assumptions about the nature of humans and the influence of society on this basic nature, and in its place it explicitly provides an account of the ideal situation from which to reason about justice.

In the Original Position, we are to imagine individuals who have comprehensive interests and individual notions of the good that have gone behind the thick veil of ignorance. This veil radically alters their epistemic situation. Though these are meant to be actual people who live in a particular time and place,

have certain skills, are of a particular demographic, and have particular interests and affiliations, behind the veil of ignorance, they are not aware of any of these facts about themselves. This ignorance extends even to knowledge of what generation one is in: individuals literally know nothing about their particular situation. Their ignorance extends beyond this, however. Not only do they not know anything about themselves in particular, they do not know anything about the general shape of their society. They do not know the proportion of rich to poor, or what the distribution of jobs is, or anything of even a statistical nature about the demographics of society. Their ignorance, however, does not extend to all domains of knowledge: they know something about economic principles, natural laws, and other theoretical knowledge that is not dependent upon particulars.

Rawls also assumes that the individuals in the Original Position are rational in the sense that they obey the rules of rational choice as developed in economic theory. This is an important assumption: Rawls aims to show that instrumental rationality alone is enough to get his principles of justice.²⁵ The axioms of rational choice here are significantly constrained, as the Original Position is a position of massive uncertainty. What is striking about this model is that the agents are engaged in a bargain over how to determine the basic structure of society, which

²⁵ It should be noted here that Harsanyi developed an account of the original position independently of Rawls at approximately the same time as Rawls, though saw it justifying a form of utilitarianism. There is some difference, however, in the “thickness” of the veil of ignorance in Harsanyi’s account, as he allows some demographic information. See Harsanyi, J. (1953) "Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking", *Journal of Political Economy* 61(5):434-5

shapes the distribution of primary social goods,²⁶ without having any knowledge of what their bargaining position is. In this way, it is a very peculiar bargain – the parties involved not only have no idea of what their current position in society is, but they also have no knowledge of what the consequences of the bargain are for them as individual stakeholders.²⁷ This kind of uncertainty is such that Rawls argues a rational agent is forced to adopt an attitude of heavy risk aversion. Since one cannot tell who one is, one ought to maximally hedge one's risks by ensuring that the worst off in society is at least above a certain threshold of acceptable living conditions.²⁸ The bargain can then be seen as one over alternative conceptions of justice, which shape the distribution of primary social goods. Since there is an infinite possibility of alternative conceptions, he assumes that the agents are given a short list of the most plausible alternatives to his theory, with utilitarianism being of primary interest.²⁹

²⁶ Primary social goods, which Rawls articulates in section 15 of *Theory of Justice*, can be thought of as liberties and opportunities, income and wealth, and how they are distributed.

²⁷ For a further elaboration on this point, see Rawls' "Some Reasons for the Maximin Criterion", *American Economic Review* (1974)

²⁸ See the article referenced in the previous footnote on the details of this reasoning. Since risk aversion is a psychological trait, and one is supposed to be ignorant of such traits behind the veil of ignorance, I view this as a point of contention in Rawls' theory. Arrow has several powerful critiques of Rawls' analysis of the veil of ignorance in his review "Some Ordinalist-Utilitarian Notes on Rawls' Theory of Justice", *Journal of Philosophy*, 1973. However, I grant the assumption for the purposes of this discussion.

²⁹ I focus on the Rawls of *A Theory of Justice* here, and not the later Rawls. Rawls abandoned the Original Position as a basis for justification in the 1980's, which I interpret as a recognition of the serious challenges he faced from economists. He later moved to a theory of "overlapping consensus" that I claim still suffers from the same basic trouble, though instead of having a basis in decision theory, it shifted to Kantian constructivism. One can see this trajectory of his thought in "Kantian Constructivism in

What I would like to discuss is not whether his analysis of rational choice under uncertainty is correct: instead I would like to argue that the setup for this thought experiment is one that, like Rousseau's account of the general will, precludes the possibility of diversity in the public sphere. The first thing to notice is that multiple parties are superfluous to the bargain in the Original Position. Since each individual has no knowledge of who she is (or even whether she is a she), or where she will end up in any given distribution of primary social goods, she (by design) has nothing to differentiate her from any of the other individuals. On Rawls' conception, the only features of herself that she has access to are that she is rational, and that she is risk averse in the face of massive uncertainty. So, there is nothing to be gained by bargaining with another individual rather than just bargaining with herself. Neither party represents any particular interests, nor do they have any private information that can inform the bargain. As such, a single individual can be an advocate for all sides of the bargain at the same time. So trivially, there can be no diversity in the original position. No one has an opposing view, special knowledge, different cognitive abilities, or anything that differentiates them from the other parties. There can be plenty of diversity in the private sphere – Rawls' principles of justice are defined so as to allow for a great deal of tolerance of diversity in the private sphere. But at no point is there room for any diversity at the level of the public sphere.

Moral Theory", *Journal of Philosophy* (September 1980), 77 (9): 515-572, then "The Idea of an Overlapping Consensus" *Oxford Journal for Legal Studies* (Spring 1987), 7 (1): 1-25, and culminating in *Political Liberalism* (1993).

We have yet to see, however, why diversity might be something that we ought to pay close attention to. To better understand what diversity can do, we now turn to an economic understanding of the role of diversity.

Economic Perspectives on Diversity

Relying on economic examples to motivate the potential desirability of agent diversity has two main benefits: first, as has already been suggested, there are a number of such examples at the very core of economic theory, and second, economics might help point us in the direction of an understanding of how to allow diversity to thrive without the downside of increased conflict due to differences in individual preferences. Market mechanisms, when they are working, have an impressive ability to coordinate the actions of large numbers of people who agree on very little. In fact, the earliest arguments for a capitalist system were not the now-standard arguments from efficiency, but rather they had a moral perspective. As Hirshmann discusses in *The Passions and the Interests*, commerce and its mechanisms were seen as agents of positive change. Smith claimed that markets made agents more empathetic, as understanding what one's trading partners want and need can better facilitate exchange. Given that there is a monetary incentive for encouraging as many exchanges as possible, there is then a corresponding incentive for coming to better understand other people. Other kinds of interactions then are made easier because of the agents' more developed moral sentiments. Hirshmann

also identifies one's "interests" in monetary gain as a constraint on one's "passions" for more ephemeral goals, such as personal honor or glory. These latter goals tended to involve bloodshed and also a large amount of unpredictability. Commerce, and the interests it generated, encouraged both predictability and cooperation. Agents going after their interests may not agree any more than agents following their passions, but their diverse interests can be coordinated by a market without any real requirements for agreement. Although it will be addressed more directly in later chapters, as we consider the cases in economic theory in which diverse agents are either recommended or required, we can briefly consider how the agents are able to coordinate.

As we saw earlier, there are several ways in which agents can be diverse. In this section, we will focus on three: preference diversity, skill diversity, and knowledge diversity. These play a large role in economic theory. This is not without cause: what we will now turn to is why these cases of agent heterogeneity are of particular importance. First we will look at the general structure of markets, but we will then delve in with a more careful look at Ricardo's theory of trade and Smith's division of labor. From there, we will turn to Hayek's understanding of the role of local knowledge in organizing a society. Finally, we will briefly turn to the role diverse agents can play in the provisioning of public goods.

A market can be considered a collection of rational buyers and sellers, and the transactions that take place between them. Buyers have a particular demand for a good or bundle of goods, and sellers have a certain supply. What is to be noticed in

this description is that it immediately pre-supposes an asymmetry between buyers and sellers. That is, sellers have something that the buyers do not, and vice versa. A basic assumption of any market model is that transactions between rational agents take place only when there is at least the expectation of mutual advantage. That is, both parties must expect to gain from the transaction, otherwise the agent who anticipates getting hurt in the deal would simply not agree to it in the first place. Because of this assumption of rationality, if a market exists, it must be the case that individuals are different in terms of their endowments, preferences, or skills.

So far this is a trivial claim – buyers have to want what sellers sell. But from this we can see more substantial claims are possible. If we are to find further conclusions, we need to add in more specifics about basic economic theory. To start with, we will consider the basic contribution of Adam Smith, in the first chapter of his *Wealth of Nation*: the division of labor.

The Division of Labor

Smith motivates the notion of the division of labor with his famous example of the pin factory. The problem he was trying to address is basic: in smaller societies, the needs to be provisioned are fairly small. Because of this, any given individual can take on a number of tasks and do them adequately enough to ensure that needs will be met. However, as societies get larger, eventually there are so

many different needs, and in such great amounts, that the previous method is no longer tenable. In a small society, the needs for pins are not so great, and so a single individual can be an artisan, and take on the entire manufacturing process. But in a large society, there is such a large demand that there simply are not enough individuals to provision pins in this manner. So instead, pin-making is broken down into separate tasks – drawing the wire, straightening it, cutting the wire, and grinding the ends. Any given individual is responsible only for laboring on a single task. So rather than make an entire pin, an individual will only cut wire, while others will perform the other elements of the manufacturing process.³⁰

This division of the tasks allows individuals to become specialized: that is, they will gain additional proficiency at performing their given task. Smith described this as gaining additional “dexterity” with their specialization. Another advantage is that time is saved, as workers do not have to constantly switch tasks – workers can spend more time being productive. A final advantage, which is not relevant to this discussion, is that specialization also encourages specialized machinery that can aid in a laborer's task.

It is the notion of specialization that is of particular interest here: the factory is able to become more productive because different groups of agents take on different tasks, and because one group dedicates itself to its task, it can more easily develop better techniques. An important consequence of this is that, as more opportunities for more fine-grained divisions of labor become possible, there are

³⁰ See *The Wealth of Nations*, Book 1, chapter 1 for the full discussion.

economic rewards to utilizing them.³¹ The division of labor, then, can be seen as an engine for the promotion of skill diversity. This has two components: first, it demonstrates how to take advantage of an existing set of skill-diverse agents, and second, it shows how economic pressures exist for increasing a population's diversity over time.

Smith's Division of Labor is seen as a major engine for growth. The material benefits of continuing the division are very large at a social level – it enables more goods to be produced at much lower costs, and so the well-being of every individual is indirectly improved through the availability of more consumer goods at lower prices. It is useful, however, at this point, to note that the advantages of the division of labor can be viewed either from the perspective of the society or the individual. Socially, it is theoretically clear that the division of labor creates huge gains, and there are always social interests in promoting further divisions and distinctions among agents and processes. However, from an individual perspective, there is a trade-off that is not represented at the social level. While individuals tend to be made more economically well-off because of their increased productivity, Smith notes well before Marx does that this can occasionally result in a loss of well-roundedness. An artisan or craftsman makes an entire product, and is able to have a more well-rounded set of skills. A specialist, on the other hand, is not able to make the whole product, and as such, is highly dependent on others. Marx pushes this idea further to argue for the alienation of the worker from the product of his labor.

³¹ Smith points this out in Book I, chapter 2.

It is for this reason that he seems to endorse an ideal of an individual as a generalist, rather than a specialist. These problems are important, and will be considered in later chapters. For present purposes, it is sufficient to show that the division of labor both supports existing diversity and promotes greater diversity in the future.

Comparative Advantage

David Ricardo extended the division of labor theme in his theory of international trade. In Ricardo's trade theory, he argued that the mere existence of differential production in a situation in which at least two goods are being traded led to the opportunity for wealth creation. To show how this works, let us first consider two terms: absolute advantage, and comparative advantage. Imagine two individuals, Alice and Bob, both of whom produce shirts and pants. Let us suppose that in a day Alice can make 12 shirts or 9 pairs of pants. If Bob made only 4 shirts or 8 pairs of pants a day, then Alice would have absolute advantage over Bob in making both shirts and pants. If Bob could make 5 shirts in a day or 15 pairs of pants in a day, then Alice would have absolute advantage in shirt-making, while Bob would have absolute advantage in pant-making.

The case of absolute advantage was accounted for by Smith. Absolute advantage merely takes into account who is more able to produce a given good given the costs (in this case, time). Trade theories based on absolute advantage

would claim that Bob and Alice would trade only in the second situation.

Comparative advantage, on the other hand, argues for considering the opportunity costs of production, rather than the actual costs. So, even if Alice has absolute advantage over Bob in the production of both shirts and pants, Bob has comparative advantage in the production of pants, as it costs him less in terms of the production of shirts to make a pair of pants than it does Alice. In this kind of situation, Alice and Bob can benefit from a trade that generates more wealth, even if Alice is better able to produce both goods. This comparative advantage should cause Bob to produce only pants, while Alice produces only shirts. An exchange rate would be determined to ensure that both parties are made better off. The exchange rate would favor Alice, but this is because she already had an absolute advantage, and is able to produce more, absent any trade. But even though Alice would experience a little more benefit than Bob, both are still made better off than had they not made any exchange.

The notion of comparative advantage is important for several reasons. First, just as with Smith, it demonstrates that material benefits can be found from the existence of agent diversity. It also helps further the account of how incentives exist to create more diversity and specialization. What is unique about comparative advantage, however, is its implicit argument that diversity is in some sense more important than absolute advantage. That is, it is a remarkable testament to the power of the dynamics of a diverse population that one can be made better off by trading with someone who is less skilled at everything. Agent diversity both creates

the opportunities for and encourages specialization, and specialization creates opportunities for material growth.³²

So we see that there are always incentives for agents who are different to take advantage of the fact that they are different, and more importantly, there are incentives for agents to become more different. At present time, this can be understood to be taking an existing comparative advantage and increasing one's specialization. This has an interesting consequence as well: while there are clearly individual benefits to improvements in one's specialization, there are incentives for individuals to help others improve their specialization, though to a lesser degree. This is because the overall skill increase translates into a larger social product from which all participants gain. While the incentives are smaller than what they are for the development of one's own talents, if they are viewed in the aggregate, they are almost always larger than the value of improving one's own skills. Imagine a talented baker, Anne, who is also a talented brewer. Next door to Anne is Carol, who enjoys brewing, but isn't very good at it yet. Since Anne wishes to enjoy both the bread that she bakes and the beer that she brews, she has a conundrum: the opportunity cost of brewing is very high, as it is difficult to acquire bread of the quality she is capable of elsewhere. Rather than simply grow despondent over her predicament, Anne decides to invest in Carol's brewing endeavors. It is in Anne's interest to invest her own resources in Carol's abilities, rather than her own, because her opportunity costs are much higher than Carol's. If Anne invested in her

³² I say material growth here given the economic context. However, I intend to argue for the cultural benefits of diversity and specialization later using similar arguments.

own brewing skill, she would have to forego all the baking that she could have done during that time. So instead of having to choose between bread and beer, Anne chooses to invest in Carol so that Anne can end up enjoying more of both, by means of trade. Comparative advantage, then, can provide a self-interested motivation for investing in the education and skills of others.

Outside of material benefits, however, does economics have any other means of supporting the claim that there are benefits to agent diversity? The epistemic situation that agents find themselves in provides a notable source of potential benefits. That is, individuals are exposed to a great deal of information, much of which is superfluous, and very little of which is complete. In some senses, we have far too much information to fully utilize it all, but in others, we lack a great deal of information that would be of great use to us if we were to have access to it. This epistemic problem stems from both the fact that we are cognitively limited, and the fact that our environment is very complex and varied.

Epistemological Challenges of Complex Environments

Hayek, in his “The Use of Knowledge in Society” argues that a central planner would be unable to properly manage an economy, because the epistemic failures of any individual would be too great. While he frames the discussion in terms of a central planner, we can think of the argument in terms of agent homogeneity. That is, rather than one agent making decisions, we have a lot of agents who have

homogeneous preferences, use the same decision-making procedure, and who have access to the same information. In this case, they will always all make the same decision.³³

There are several dimensions along which this homogeneity causes a problem. I will argue that across each dimension, agent diversity is incentivized. Hayek focuses on one dimension: knowledge and information. Given that the environment we are in is extremely varied, knowledge of local conditions could be extremely important. For example, Europeans attempting to use farming techniques developed in Western Europe in sub-Saharan Africa failed due to the fact that the local growing conditions were significantly different from those in Europe.³⁴ Less severe cases of this are true in general with something as basic and important as farming: rain levels, soil conditions, access to technology and infrastructure, the local population, local flora and local fauna can all impact what an appropriate farming decision is. Individual farmers have very good reasons to become experts in the conditions of their local area, but have little reason to develop such specialized knowledge about other areas. Even if they did have such an interest, it would be an impossible task, as there is far too much information to learn. Similarly, a mayor has good reason to know all about the political players in her city and to some extent her state, but has little reason to make such efforts to learn

³³ Of course, this would be false if there were some method of randomization in the decision-making process.

³⁴ Jared Diamond's *Guns, Germs and Steel* (1997) provides a great deal of evidence for the importance of the transferability of knowledge and skills across environments in determining larger socio-economic outcomes.

about the political players in other municipalities. A trivial example is knowledge of the roads in one's town. There are good reasons to invest time and energy in learning shortcuts, traffic stops, and other peculiarities of the roads that one travels on every day, but there is no reason to do that for somewhere else. Simply in virtue of their physical location, agents need to develop specialized local knowledge.

Hayek correctly argues that there are ways in which information can be aggregated for the benefit of society as a whole. In a market economy, the primary means by which this happens is the price system. If I were to find out that orange trees had severe frosting, I would buy as much orange juice as I could, in anticipation that the reduced supply of oranges would drive up prices. This act in itself would help drive prices up. The price, then, functions as an aggregate statement about the relative prosperity of the orange crop in comparison with other goods. Prediction markets can do this for a variety of areas in which we lack solid information. The classic case of this is sports betting, but now the same methods are applied to a number of areas, including the outcomes of political elections. People can profit from their specialized knowledge by betting before others do. By this means, the market incentives motivate people to make their private knowledge public, and we all benefit epistemically.

A second major avenue by which a complex environment incentivizes diverse agent populations is by means of encouraging a varied approach to decision-making procedures. In a market example, if we consider the huge amount of complexity involved in the changes in prices across different market sectors, it is in fact

impossible for a single investment approach to be successful. If everyone bought and sold the same stocks at the same time, prices would be subject to fluctuations that would result in huge extremes – prices of some stocks would become zero, and others would approach infinity.³⁵ Modern financial markets have begun to show some features of this, due to what is known as program trading. Program trading is the practice of programming a computer to make automated trades of several financial instruments at once based on certain threshold values. This strategy is heavily employed by hedge funds, which often have fairly similar thresholds for trades. This automated process, which moves huge amounts of money around at once, creates a great deal of volatility in financial markets.

Suppose we assume a much less-extreme version of this, where only the majority of investors apply the same decision procedure, and the remaining part of the population have their own investment strategies. In this case, the response from the remaining portion of the population would be to arbitrage away the gains that might be possible for them. As an example, imagine that the Chinese stock market were completely closed to external investment. Though some more affluent citizens might be investing in the market, the dominant investor is the state. If the state's investment strategy were discernable to other investors, they could just, for instance, purchase the stocks that the state would soon buy, and sell the stocks that the state would soon sell. Steps like these allow the other investors to treat the state

³⁵ Of course there are practical limitations to infinite prices, as there is not an infinite money supply. However, if everyone were to buy the same stock at the same time, with no one willing to sell, the price would be in effect infinity, as there is no market-clearing price possible.

as a money pump, arbitraging away any gains that the state might have made. In general, if one employs a known investment strategy, or one that is common, one opens oneself up to such arbitrage opportunity. Markets cannot “move” – that is, make trades – unless there are so many investment strategies being employed at once that any given strategy cannot be distinguished from “white noise.” This is a consequence of the No-Trade Theorem in economics.

While the No-Trade Theorem establishes that diversity is required for basic (stock) market function, there is a more interesting second claim that can be made, that is more in line with Hayek’s observation about knowledge. Because there are a great many unknowns about the structure of the information we receive from the world, it is difficult to know whether the decision procedures we are using are any good, let alone optimal. The simplest way of getting more data on the quality of one’s decision strategy is by means of performance comparison. So if a population employs a range of decision strategies, it is easier to see which of them have better outcomes. The population can then shift towards using the stronger decision strategies, and then attempt further refinements. This process of imitation and refinement can be fruitfully compared to an evolutionary mechanism. We can see that imitation can work only insofar as there exists variation in the population. What’s more, variation must be continually present. These varied strategies will have an associated fitness, based on how well they deal with the underlying epistemic situation. Better strategies will be selected for in the population, as

people can then simply copy what they see as the more successful strategies in the population.

It may further be the case that some decision strategies are simply better in certain contexts. That is, local variations in the information environment will favor a decision procedure that is not globally favored. For example, environments that can be characterized as having many micro-climates, such as the San Francisco Bay area, Seattle, or Portland might be ones in which it makes more sense to invest more heavily into regularly gathering weather information, whereas this would not make much sense in a location with a fairly stable climate, such as Los Angeles. As such, the environment itself can be a strong motivator of decision-making diversity.

Thus far we have found that within economics both from a material perspective and from an epistemic perspective, there are significant claims to both the fact of diversity and the value of diversity – even among ideally-rational agents. The final area we will consider from the economic perspective is that of collective action. In this case, we will take into consideration Hardin's work in this area.

Public Goods and Collective Action

In his *Collective Action*, Hardin provides a framework for judging how different groups succeed in providing public goods or preventing public bads. For our present purposes, we will only need to consider the basic tenets of his account. We can imagine a population of size P who potentially have an interest in

provisioning a good G . This good has a value g_i to each individual i in the population. There is also a single cost for provisioning the good, C . This cost is a fixed amount, independent of the size of the population. Due to the nature of public goods, one cannot prevent others from having access to the good – it is either provisioned to everyone, or it is provisioned to no one. The good is provisioned if there is a minimal sufficient subgroup K of the population P , for whom the value g_i is greater than C/K .³⁶ If this is the case, then each of the individuals in K would be willing to pay for provisioning the good G . If this is not the case, then the good cannot be provisioned without some kind of external authority that could coerce people into providing the public good.

Hardin considers a few extensions to this basic model relating to asymmetries in the agent population.³⁷ He conceives of asymmetries as differential benefits to a public good, or differential costs. These translate well to preference diversity and skill diversity, respectively.³⁸ An important feature of these asymmetries is that they can make the provisioning of public goods easier because

³⁶ This further depends on the size of K itself. I refer to K here as the minimal subgroup of P that meets the constraints described. Hardin has a detailed account of the importance of smaller K 's for successful provisioning of public goods. Here all we need to rely on is that smaller K 's are better able to successfully provision public goods than are larger K 's. This argument is developed in detail in Olsen's *The Logic of Collective Action*, and is also refined in Hardin's *Collective Action*, particularly in chapter 3.

³⁷ See pages 68, 111, and 119.

³⁸ Note that there is a difficulty in directly translating differences in costs to differences in skills, if we assume that individuals have a declining marginal utility for money. In this case, costs can be lowered either by means of the possession of some skill or by possessing a large endowment of money. Insofar as money can be turned into a technological advantage, or just a direct purchase of the good in question, it can lower costs the same way a skill can.

of how they can affect the minimal size of K . Take the case of preference diversity. Imagine two populations, P and P' . In population P , everyone has the same set of preferences, and as such everyone has the same valuation for the provisioning of a good G . In population P' , however, agents have different preferences, but we will assume that the mean³⁹ of these preferences in relation to good G is the same as the mean in P . Since preferences are diverse, we know that there must be individuals who have valuations greater than the mean. Just with this information, we can establish that, for any given K where $\text{Card}(K) > 2$, the size of K' of P' can always be smaller than K . If we have additional information about the distribution of preferences, we can make more accurate statements about how much smaller K' will be.⁴⁰

Hardin notes that a single asymmetric population is likely to be able to find support for a number of public goods simultaneously, since individuals can be highly motivated across many different dimensions of a preference space.⁴¹ It also creates opportunities for some amount of logrolling, in which different subgroups can trade support for their causes in order to achieve a sufficient number of contributors. It should be noted that we can apply similar arguments to the case of skill diversity – insofar as some skills can lower the effective cost of provisioning a given public good, even if we hold preferences constant in a population, the differences in skills create opportunities for smaller K 's, since there will exist

³⁹ This argument would also work if we employed the median rather than the mean.

⁴⁰ On page 68, Hardin speculates that the nature of, say, political views is such that the distribution will be sufficiently skewed to encourage the creation of smaller K 's.

⁴¹ See page 119 for his discussion.

subgroups for whom the effective cost of provisioning the good is lower than it would have been had everyone shared the same skill set. These lower costs increase the likelihood that the good can be provisioned.

A more difficult question to ask is what happens when a population has diversity of both skills and preferences at the same time. This will, to a large extent, depend on the covariance of skills and preferences. There are three classes of cases that are interesting to consider: positive correlation, independence, and negative correlation. In the case of positive correlation, we find that individuals who have a greater preference for a good G also have skills that would lower the cost of providing the good. Here we can imagine that people who like public gardens might also have a greater ability to garden themselves, or those who would like a community orchestra are also those that are themselves musicians. We can see how other public goods can have similarly motivated individuals, like in the case of a firehouse, a community watch group, and lifeguards at the beach. Because of this positive covariance, we should expect the minimal K to be even smaller, since perceived benefits are larger and costs are lower.

In the case of independence, we effectively reduce ourselves to the previous consideration of varying only one parameter at a time. Since skills and preferences are independent, we should expect minimal K 's to be roughly the same size as the cases in which only one parameter is varied, though perhaps more likely. There are three ways one can consider the composition of the members of a subgroup K : they could all be higher-skilled, they could all be more interested in the good, or they

could be a mixture of the two. In the first case, we can imagine that engineers are more capable of providing public goods like levees, even if they don't have any special interest in the construction of any particular levee. As with any case of those who are more skilled but perhaps less interested in provisioning the good, opportunities for logrolling are available. That is, those may trade their efforts to support one public good in order to get help on another good. In this example, perhaps engineers as a group have a greater interest in having recycling programs. They can trade their talents in levee construction for support of recycling centers. The second case to consider here is a group with skills no different from the rest of the population, but with a great interest in providing the good. A national army seems to be a candidate for this. Soldiers appear to place a great deal of value on the kind of service they perform for their nation. Even though there are plenty of able-bodied, and even athletic people that are not soldiers themselves, they do not place the same amount of value on the service. The final case to consider here is that of a K with a mixed population. This likely covers a broad range of cases, but one important example is that of the system of public defenders and prosecutors in the legal system. Some public defenders do it because it is the virtuous thing to do, and some might do it as a career stepping stone – it is a means to demonstrate their legal skills, and better position them for either lucrative jobs in private practice or consideration for a judgeship. Though the individuals might have different motivations, as a whole they enable the public good to exist and to persist through time.

In the final case that we can consider, we could have negative correlation of skills and preferences. In this case, we are forced into the logrolling scenario: individuals who have the most ability to achieve some end X, but do not want to achieve it, can trade their abilities to a group who wants X but cannot provision it themselves, in exchange for their ability to provision Y. How well this scenario works is dependent on the particular utility calculations of the different cohorts of individuals. Individuals will be willing to trade skills only so long as the trade remains utility-increasing. That is, the positive utility gained from getting their public good must overwhelm the potentially negative utility of working to achieve some other public good that their skill set happens to be suited to. This is further complicated by whether there exist opportunities for shirking if the trade is sequential, or one good requires some amount of maintenance, while the other has only one-time startup costs. Analysis of these situations is left to later chapters.

This analysis of asymmetry in collective action makes a significant claim about the value of diversity in either economic or political contexts. Agents in more diverse populations are more likely to have more public goods provisioned, even if their mean preferences and their mean skills are identical to a homogeneous population. More public goods provisioning results in higher average utility, and so provides material benefits to the agents. The last case considered, in which a diverse population had a negative correlation between skills and preferences, is by far the most interesting to consider, but it tells a potentially mixed story. In this case, it is no longer clear that each individual's utility has increased, if there are

many public goods and many skills, as some public goods may be viewed as public bads by portions of the population. An example of this could be the provisioning of a town bell to announce the five prayer times for a Muslim community, which may be viewed as a nuisance by a non-Muslim portion of the population. Since logrolling and trades could be sequences of bilateral agreements in a population with more than two groups of individuals, it could be the case that some bilateral trades might have negative externalities for third parties. At the same time, however, it is equally possible that there are positive externalities. However, the complexity of this particular case should not distract from the general result that more public goods can be provisioned, and in most cases, this will generate material benefits for the whole population, though the amount of benefit will vary across the population.

Summary of Economic Perspectives

What we have seen so far is that in four different cases in economic theory, we have arguments that support the value of diversity. In our beginning investigation of economics, we noted that diversity is required for the existence of a market. From Smith, we find that the division of labor both creates additional material benefits, and provides incentives for agents to differentiate their skill sets. This provides direct support for the value of skill diversity, as well as some evidence for its inevitability. We found that Ricardo further develops this thread by means of

the notion of comparative advantage. Combining comparative advantage with the ability to trade increases the value of skill diversity even further, as we find that even if one agent's skills are completely dominated by another's skills, there are still opportunities for trade that make both better off. Our discussion of Hayek showed that there are epistemic benefits to skill diversity, knowledge diversity and perspective diversity. Finally, an analysis of collective action demonstrated that preference diversity and skill diversity can provide material benefits to both individuals and a population as a whole.

This investigation allows us to see that a population of rational agents will not only always be diverse, but will frequently have reasons to become more diverse. Insofar as these kinds of diversity are correlated with differences in perspectives on the nature of the public sphere, an analysis of the public sphere will have to recognize these perspectives, or heavily restrict the private sphere so that it too becomes homogeneous. If we are to retain at least some of the goals of the liberal political tradition, and want to maintain a relatively free private sphere, then we are forced to acknowledge the legitimacy of diversity in the public sphere as well. A just social order will rely on understanding the nature of the world that we find ourselves in, understanding of our nature as humans, understanding how we are able to interact with each other, and discovering mechanisms to achieve the ends that we set for ourselves. This is contingent on too many factors that are discoverable only by rational deliberation informed by social experimentation. As

this chapter demonstrates, the presence of and thoughtful leveraging of diversity helps make these challenges more manageable.

Chapter 2: The View From Everywhere

Introduction

A fundamental question to ask, and indeed the first question to ask, in moral and political reasoning is from what perspective we ought to be reasoning. There are two prominent moral stances that are adopted in answer to this question. Nagel has famously called the more prevalent stance “the view from nowhere.” This position supposes that moral reasoning ought to be done from a completely neutral position, where particular features of individuals, including their particular wants and desires, are eliminated from the realm of reasoning. This stance was adopted by Rawls, and I would claim that it is the dominant approach in contemporary moral and political philosophy. The less common, but still historically important stance is what I will call “The view from somewhere.” This view supposes that rather than eliminate all perspectives in order to achieve neutrality in our reasoning, we ought to choose a particular perspective, and reason from there. Prominent traditions of the view from somewhere are found in Aristotle, Hobbes, and neo-Hobbesians such as Gauthier. In the neo-Hobbesian tradition, the chosen perspective is first-personal: each individual ought to reason about what would make that person best off, recognizing that others are trying to do the same thing. In this perspective, individuals are indifferent to other people’s well being, and care about others only insofar as they help or obstruct them in trying to achieve their ends.

These positions have dominated the debate over what our moral perspective ought to be, but they do not exhaust the possibilities. Notably absent from this debate is the other possible perspective, “The view from everywhere.” It is this view that I aim to defend here. The view from everywhere is both a rejection of the two alternative views and an attempt to incorporate their key virtues: namely, the goal of neutrality found in the view from nowhere, and the recognition of the importance of particular interests and desires found in the view from somewhere. It is because of this combination of their virtues that I will argue that the view from everywhere is in an epistemically superior position compared to the other choices of perspective.

Before we can see what the view from everywhere *is*, however, we must first see what it *isn't*. To that end, we will first look at the view from nowhere as Rawls and Nagel use it, and then look at the view from somewhere as Gauthier uses it. The goal in each section is to understand both what each perspective requires and what each perspective provides to the reasoner. Once we have seen the outlines of each, we are then in a position to see what is different about the view from everywhere. Armed with this knowledge, we can then see how the view from everywhere provides a superior perspective for moral reasoning.

The View From Nowhere

The view from nowhere is an attempt to obtain a purely objective perspective. This objective perspective is obtained only by shedding particularities of our subjective selves – our particular capacities, perspectives, and relationships. As Nagel says in his opening paragraphs to *The View From Nowhere*, “we may think of reality as a set of concentric spheres, progressively revealed as we detach gradually from the contingencies of the self.”⁴² This gets at two crucial features of his system. Namely, he is a moral realist, and he believes that if we are to discover objective reality, we must look to our objective selves rather than our subjective selves.

Nagel does indeed hold that we have both objective selves and subjective selves – our subjective selves have the features that we pre-theoretically associate with ourselves, such as particular wants and desires, views of the world, and interpersonal relationships. Our objective selves, on the other hand, are without individual perspectives. We each have our own objective self, which is meant to be us, but stripped of all of our subjective features, including a sense of what it is like to be ourselves. The objective self thus sees the world from a view that is outside of the world – nowhere. Our subjective selves are in this world, but our objective

⁴² T. Nagel, *The View From Nowhere*, Oxford University Press, 1986 , p. 5

selves in the view from nowhere have no particular attachment to our subjective selves.

To get at moral truths, Nagel believes that we ought to reason from the view from nowhere. After all, in this view, we have eliminated the subjectivity that can cloud our vision. From this objective point of view, we not only ensure that the outcome of our moral deliberation will be objective, but that it will be responsive to objective features of the world. This is not to say, however, that this objective perspective does not see and attempt to respond to differences and particularities in individuals. It is simply that these particularities do not also come with the force of a subjective experience.

This position is contrasted with a different “view from nowhere,” which Rawls's work embodies. Rawls moves away from any metaphysical claims, and instead opts for a procedure that similarly aims to strip away those features of ourselves that might bias us in our process of moral reasoning.⁴³ The procedure Rawls has in mind is deliberation in the Original Position, behind the “veil of ignorance.”

Behind the veil of ignorance, one (or one's delegate) is self-interested, but does not know any particular features of one's self. Reasoning behind the veil of ignorance is in a sense a more stringent condition than what Nagel proposes. This

⁴³ Indeed, while Nagel's view from nowhere is theoretical philosophy, Rawls' Original Position is meant as a piece of practical philosophy. While these views have substantive differences between them, I lump them together to show that they have important common features and suffer from the same epistemic problems.

reasoning is from a position of complete ignorance: ignorance of oneself, one's place in society, and indeed, the makeup and nature of that society. Individuals in this position are rational, but do not know what their ends might be. As the individuals in the Original Position attempt to determine which rules ought to govern the basic structure of society, their rationality in the face of ignorance leads them to choose those rules that ensure the worst off in society is made as best off as possible.

Now that we have seen brief accounts of both Nagel's view from nowhere and Rawls' original position, we can consider what the common features are that are desirable for a moral stance to have. Most notably, they both embody a perspective of neutrality. Reasoning from these perspectives does not allow one to privilege oneself as more important than others in moral deliberation. This is done procedurally – there simply isn't anything that agents in these moral stances could use to become biased. The procedures that are introduced both aim to remove any information that is morally irrelevant.

While this position of neutrality is a very desirable quality in a moral stance, their approach to achieving it presents two potential epistemic challenges. The first worry is that it is unclear that we know how to reason from such a “perspectiveless” perspective. Since we have always viewed the world from our own perspectives, we must not only be able to evaluate a novel decision scenario, but do so from an alien perspective. The common technique for attempting to reason from such a perspective is to use counterfactuals. That is, we engage in thought experiments

where we take ourselves to have adopted the desired perspectiveless stance. It is within the confines of these thought experiments that we attempt to determine the consequences of these perspectives that we cannot ourselves take on.

Just as Nagel thinks we cannot know what it is like to be a bat, and Rawls thinks that the Utilitarian need for an interpersonal comparison of utilities is doomed, I suspect that we cannot know what it is like to reason without our own perspective.⁴⁴ This is different in kind than some other thought experiments: in the famous Trolley Problem, people tend to have clear reactions to the scenarios. Though few people have ever had to decide whether to kill one person to save five others, the scenarios ask us to reason as we normally would. So though the situation presented is novel, the methods by which we are meant to evaluate it are the same as what we use in our everyday lives. It is this feature that makes the trolley case, and other scenarios like it, so useful in moral reasoning. The examples can be designed to put us in new situations that might reveal something about how we ought to think about situations that are more typical. But they all assume that

⁴⁴ It might be contended that mathematics or the hard sciences embody precisely this kind of objective, “perspectiveless” perspective. Features of individual scientists do not matter in their work, and so they just drop out, and all that is left is the science itself. I will deal with this more substantively later in this chapter, but I will note here that this seems to be a false characterization of the practice of science and mathematics. Feminist critiques have amply demonstrated that even the hard, data-driven sciences show significant evidence of bias. Even mathematics is in an important sense deeply perspective-laden: there are often multiple ways to prove the same theorem – the way that is chosen can be based on notions of elegance, utility, concision, explanatory power, or any number of other criteria. This is readily evident when one compares how a human proves theorems to how a computer proves theorems – the approaches are drastically different, even though humans programmed the computer. Computers are not reliant on the semantics of axiom systems, and so consider portions of the solution space that humans ignore, and as a result often find novel proofs.

we reason as ourselves. This is the key difference between situational counterfactuals and perspectival counterfactuals. Situational counterfactuals remain grounded in our experience precisely because we use the same perspective that we normally would in a novel situation. Perspectival counterfactuals, on the other hand, offer no such grounding. Even if the situation has many external features with which we are familiar, we cannot account for how to relate our experiences to a novel perspective. This is due to the fact that our experiences are mediated by our perspectives. Though I have spent a great deal of time in Philadelphia, I have no sense of how a squirrel would experience the same place. I may be able to imagine myself as a squirrel, and even get a great deal of the key features of such an experience correct. It could even be a very vivid thought experiment. The worry is, however, that these kinds of thought experiments generate vivid false positives: these thought experiments might feel really right, but this feeling would be the same whether or not the thought experiment was carried out well. I may have left out some key feature of being a squirrel, and would never know it, because my thought experiment appeared to be successful. This poses a basic challenge to the claim that we are able to conduct the thought experiments that Nagel and Rawls propose, which would inhibit our ability to adopt their moral stances at all.

Let us suppose that we can know what it is like to reason without a perspective (or at least, from a foreign perspective), which would put our first worry to rest. We still have to deal with our second epistemic worry. Namely, how

do we know when we have rid ourselves of our own perspectives and successfully adopted this perspectiveless perspective? The starting point is clear enough – we start with ourselves as we are. Then we shed those features of ourselves that are subjective, until all we are left with are the objective, universal aspects of ourselves. It is this second step that poses a challenge. Imagine that you are trying to eliminate those features that are subjective (or more generally, morally irrelevant). What Nagel proposes is to run through all of your personal commitments, relationships, beliefs about yourself and your surroundings, and decide, for each one, whether or not it is morally relevant. If it is, you keep this attribute. If it is not, you throw it away. But to carry out this procedure, you must already know what you want to keep.⁴⁵ More troubling, however, is that you do not have a way of determining if you have carried out the procedure appropriately. People make mistakes, there needs to some way of checking to see how well you did in carrying out the procedure. But there is no independent standard available to check against. One may suggest turning to Reflective Equilibrium to solve this problem. But Reflective Equilibrium is calibrated against one's own theoretical commitments and moral judgments – these are precisely the things in question. So it can't provide an independent standard. Thus, it seems that while you can try to carry out the procedure for eliminating morally irrelevant features of one's perspective, there is no way to know if the procedure has been successful. This

⁴⁵ A procedure of this kind would rely on the idea that the view from nowhere is just a restriction on what information one has access to. But the view from nowhere is not merely a restriction of information, but an alternative perspective. (See upcoming discussion of perspectives.) In this sense, then, the problem is much more difficult than is being expressed.

suggests that we cannot rely on reasoning that uses such a procedure as its standard for justification.

Habermas raises a similar challenge to Rawls in the famous Rawls-Habermas debate in the *Journal of Philosophy*:

Things are different when the veil of ignorance constrains *from the beginning* the field of vision of parties in the original position to the basic principles on which presumptively free and equal citizens would agree, notwithstanding their divergent understandings of self and world. It is important to see that with this initial abstraction Rawls accepts a *double* burden of proof. The veil of ignorance must extend to all particular viewpoints and interests that could impair an impartial judgment; at the same time, it may extend *only* to such normative matters as can be disqualified without further ado as candidates for the common good to be accepted by free and equal citizens. This second condition makes a demand on the theory that is difficult to meet, as is shown by a brief reflection. Following the justification of the principles of justice, the veil of ignorance is gradually raised at the successive steps of framing the constitution, of legislation, and of applying law. Since the new information that thereby flows in must harmonize with the basic principles already selected under conditions of informational constraint, unpleasant surprises must be avoided. If we are to ensure that no discrepancies arise, we must construct the original position already with knowledge, and even foresight, of all the normative contents that could potentially nourish the shared self-understanding of free and equal citizens in the future. In other words, the theoretician himself would have to shoulder the burden of anticipating at least parts of the information of which he previously relieved the parties in the original position! The impartiality of judgment would be guaranteed in the original position only if the basic normative concepts employed in its construction - those of the politically autonomous citizen, of fair cooperation, and of a well-ordered society, in the specific sense Rawls attaches to these terms - could withstand revision in light of morally significant future experiences and learning processes. (Habermas, p. 118)

What I want to focus on with this passage is his point that the theoretician must have already built in all of the normative contents into the moral stance in the

first place.⁴⁶ These challenges echo my second epistemic challenge: that we have to already know what all of the morally relevant facts are before we even begin to engage in our moral deliberation. The Original Position, then, is not a very neutral position to take – instead, it places heavy constraints on what the outcomes could be, independent of the questions of justice at stake. And while the outcome is a good one, this suggests that the procedure itself isn't doing the work, but rather what went into setting up the initial conditions for the procedure carry a significant burden. I raise this as a concern precisely because it relies on the theoretician already knowing the morally relevant facts to derive the result. It is not a process of discovery, but rather a kind of justification.⁴⁷

The basic worry with the view from nowhere approach that I take to be embodied by both Nagel and Rawls is that this presupposes that we (whether we are theoreticians or the agents in their thought experiments) would have to know so much about normative concepts and facts before we even begin that the view itself does not do any work. While the goal of neutrality (and in the case of Nagel,

⁴⁶ In his “Rawls on Justice,” Nagel makes a similar point: namely, that the Original Position is far from neutral or even objective in its construction, as it has a heavy individualistic bias. The allocation of primary goods themselves can also be non-neutral, depending on the actual life plans of citizens. Primary goods are not equally valuable across various conceptions of the good. (Nagel, 228) Rawls would contest the final point by noting that the set of primary goods are meant to enable any conception of the good, rather than necessarily be taken advantage of by all conceptions of the good.

⁴⁷ I do not think that Rawls would disagree with the claim that the original position is meant to be a justification rather than a discovery process, but I do think it is relevant to how useful it is as a moral stance. Habermas has another critique of this process in (Habermas, p. 128).

objectivity) is laudable, the method by which they attempt to achieve it is highly questionable.

Now that we have looked at the view from nowhere, we can turn to its most significant rival, which I will call the view from somewhere.

The View From Somewhere

The characteristic difference between the view from nowhere and the view from somewhere is that, while the former aims to achieve neutrality between individual preferences by adopting a stance that is “outside” the world of preferences, the latter makes no particular attempt at achieving neutrality, and instead adopts a perspective that responds to particular preferences and beliefs. For the purposes of discussion here, I will consider a particular strain of the view from somewhere, which is rational egoism. I choose rational egoism here because within social contract theory, it is the primary alternative to a Rawlsian project.⁴⁸ Within the wider realm of ethics, however, the view from somewhere applies to much more than just rational egoism. Aristotelian virtue ethics, for example, would fall under the classification of a view from somewhere. Nothing I say here relies on features found only in rational egoism, and should be read as addressing all theories that fall under the view from somewhere.

⁴⁸ Utilitarianism does not neatly fit into the divisions I suggest. Of the extant moral theories, it most closely shares some characteristics with my own view in the sense that it is both impartial between agents and interested in the subjective interests of the agents themselves. But it differs from my theory in how it aggregates those subjective interests.

The primary feature of a view from somewhere is that it abandons the goal of neutrality. Rather than try to shed all of one's particular beliefs and attitudes in order to attain an impartial perspective, the view from somewhere instead argues that one should view morality from one's own perspective. After all, one can reasonably suppose that their own experiences, beliefs and desires have helped shape their moral attitudes. Those theories that put character as a central feature of what it is to be moral must eliminate impartiality, precisely because one's own motivations and dispositions are central to such a conception of morality. In either the case of a rational egoist or of a virtue theorist, one's own interests and dispositions are taken to be informative of the correct thing to do. So not only is impartiality not achieved, but it is actively dismissed, due to the fact that it would require a significant loss of morally relevant information. The first-person perspective is thus shaped by one's own interests, relationships, and beliefs. This perspective is well known to the agent, and as such, it is clear how one reasons from it. Thus, there is a trade off between neutrality and epistemic feasibility.

This tradeoff can be seen in Gauthier's *Morals by Agreement*. The moral stance developed there is completely from the perspective of the individual agent. Gauthier argues that morality follows from individual rationality. Neutrality is completely abandoned. In fact, Gauthier's theory requires that agents be indifferent towards each other: their utility is completely determined by the satisfaction of their own desires.⁴⁹ This constraint does not stem from a moral point of view, but

⁴⁹ The conception of non-Tuism is developed in (Gauthier, p.311).

more of an epistemic one: neither the sadist nor the altruist has reliable epistemic access to the experiences of pleasure and pain of others, so they are not able to count those inaccessible experiences as part of their utility. This limitation does not mean that other individuals do not weigh on one's moral considerations, however. While we may not be able to ascertain the particular utilities that others derive from particular actions of ours, we are able to see how they respond to those actions. So the sadist can take pleasure in people's screams, or the altruist can take pleasure in expressions of happiness. This external manifestation of apparent internal attitudes is readily available to anyone, and so can be used in utility calculations. Agents are thus able to reason as if they were Robinson Crusoe, whether or not they are alone. This is consistent with the view from somewhere's interest in epistemic feasibility.

We find a similar situation with Aristotelian virtue ethics. Insofar as our friendships with others are important for the development of our own virtue, we must look to see which of those friends act in ways that demonstrate their own virtue. Associating with people who have noticeable character deficiencies makes it more difficult for us to maintain our own character. Similarly, one must be generous to the right kinds of people, and show respect to those that are deserving of it.

On this view, neutrality is not only lost, but wholeheartedly abandoned: agents act only from a position of their own interests (whether they are defined by utility or attainment of virtues), and care about others only insofar as they help or hinder their ability to attain those interests. In neutrality's place, we have a much more plausible set of epistemic constraints. The view is in some sense less desirable

because of the lack of traditional moral concerns that might be embodied by neutrality, but it is far more feasible to achieve in the sense that it does not place any unrealistic epistemic burdens on either the theoretician or the agents themselves.

In our survey of the two most prominent choices for moral stances, we find that they have represent the two ways of responding to a tradeoff between neutrality and epistemic feasibility. While the view from nowhere can make a claim to neutrality, it can only do so by means of heavy epistemic demands. The view from somewhere is a stance with fairly straightforward epistemic requirements, but it lacks neutrality. While this tradeoff seems inescapable, given that neutrality seems to require adopting a viewpoint that is not naturally available to us, I propose an alternative moral stance that sidesteps this tradeoff, and allows us to attain both neutrality and epistemic feasibility.

The View From Everywhere

The view from everywhere is an attempt to reconcile the positive features of both the view from nowhere and the view from somewhere. In particular, it aims to achieve the neutrality of the view from everywhere while maintaining the epistemic feasibility of the view from somewhere. To do this requires a significant departure from both views: most notably in that the moral-political stance is not something to be taken on by a lone individual. Instead, it is fundamentally an aggregative social viewpoint. In addition, the view aims to be a project of discovery rather than

justification. Rather than build in and justify substantive moral claims, the view is designed to allow these claims to arise from collective deliberation.

The basis of this view is the conviction that individual beliefs and values matter for moral reasoning, but only insofar as they provide support for what moral principles we ought to adopt. Paying attention to individual beliefs rather than “universals” makes moral disagreement much more likely, which on the face of it would suggest that this is a problem. However, it is important to note that individuals not only can have different beliefs, but also different perspectives. These different perspectives are an essential aspect of the view from everywhere: where diverse beliefs might cause disagreement, diverse perspectives provide the opportunity to understand where these beliefs come from, and how they might be combined to determine an answer to that disagreement.

Thus far we have seen that the view from everywhere can be seen as a search process that has three basic elements. First, we have a collection of everyone’s beliefs and desires. These serve as the support for our procedure. There is a worry, however, that beliefs are on much weaker ground than the kind of evidence gathered in the hard sciences. I claim that beliefs, though potentially unreliable, are the best indirect measure for evaluating particular moral principles. Consider a case of fair division. Alice and Bob are trying to determine how to divide 12 avocados and 12 grapefruit. We may have several evaluative beliefs about this case, such as what would constitute an equitable division, but these beliefs do not stand on their own. They are instead informed by factual beliefs. For instance, perhaps Bob has a

nutritional deficiency that would be better served by avocados than grapefruit, or Alice doesn't like avocados. These facts will shape our evaluative beliefs even though they are not themselves value-laden.⁵⁰

However, beliefs alone would not suffice, so we need to introduce a second element in the theory. This second element is the collection of everyone's perspectives.⁵¹ These perspectives allow us to evaluate the status of given beliefs by means of how many independent lines of support they have, which helps to solve some of the perceived problems with other aggregation measures like the Condorcet Jury Theorem. Thus, instead of considering raw beliefs and the number of people that hold them, we want to evaluate beliefs in terms of their support from different perspectives.⁵² Finally, we have a method for taking these beliefs, desires, and perspectives as input, and output impartial statements about the evidentiary status of particular moral principles. This method aggregates beliefs in terms of their support from perspectives, and performs a version of the Condorcet Jury Theorem over belief-perspective pairs, as opposed to belief-person pairs.

⁵⁰ See M. E Yaari and M. Bar-Hillel, "On Dividing Justly." *Social Choice and Welfare* (1984) 1:1-24 for more examples of this nature, and a detailed discussion of how our factual beliefs shape our evaluative beliefs.

⁵¹ I later define perspectives as having two components: an imposed ontology, and an evaluative heuristic.

⁵² For example, belief B_1 might be consistent with perspective P_1 , whereas belief B_2 might be consistent with perspectives P_1 and P_2 . In this case, we would say that B_2 is better supported than B_1 .

The Role of Beliefs

As was the case in the view from somewhere, we are interested in first-person information, including one's own beliefs and desires. Individual beliefs matter in virtue of the fact that they matter to the individuals who hold them. The difference that we have in the present case is that we are interested in the beliefs of not just one individual, but all individuals.⁵³ The content and nature of those beliefs is not restricted: we are not interested in only so-called "impartial reasons", but rather any beliefs and desires that individuals might have. The aim is for all of these beliefs and desires to be potentially used in moral deliberation, not just some pre-determined set.⁵⁴

The beliefs of individuals matter on this view because beliefs are a primary source of support about the status of the moral principles that we are trying to learn something about.⁵⁵ This learning is something that must take place over the course of a lifetime: one forms working theories about the world, and new experiences may

⁵³ The scope of "all individuals" is a relevant concern, as it could be fairly limited or rather extensive. For the purposes of this chapter I take it to be referring to the relevant community that is seeking to engage in a social contract.

⁵⁴ I say "potentially used" rather than "used" here because the point is that there is no a priori restriction on what can be used in moral deliberation. I do not wish to hold the view that everything *will* be used, as that is almost surely false: my beliefs that my shoes are tied and I am sitting down are unlikely to be productive contributions to moral discourse.

⁵⁵ Other relevant sources of information are scientific, such as biological facts about humans.

cause us to update these theories. On this view then, an a priori deduction of the principles of morality is as open to modification as any empirical hypothesis.

A potential challenge to the claim that we should use people's beliefs and interests as sources of support is that this is unlike the hard sciences. Biologists don't attempt to confirm the theory of evolution by means of marshaling evidence about people's beliefs about the facts of evolution (at least not before said facts are discovered): it is foolish on the face of it to attempt confirmation through such means. Similarly, the beliefs of individuals do not carry much weight in a justificatory process for physics or chemistry. Beliefs of scientists matter only insofar as they are based on empirical and theoretical findings. So, why should we pay attention to beliefs and interests in the case of moral theory?⁵⁶

The answer to this challenge lies in the fact that we have not yet found an agreed-upon method of directly measuring goodness or badness, or rightness or wrongness, or even resolved whether such a thing could be possible. There is not an identifiable physical property for which we can create measurement instrumentation. What's worse, not only do we not have some physical property to point to, but we have no agreed-upon non-physical property that we have some extra-physical access to. So, as we do not have a target to directly measure, we are

⁵⁶ This is true both of laypeople and moral philosophers. Though much theoretical work is done in the hard sciences, empirical testing remains an essential standard against which theories are judged. String theory, for instance, is under increasing levels of attack from other areas of physics precisely because it remains without any empirical validation. In philosophy, no such checks exist against the intuitions of philosophers – except competing intuitions – precisely because the intuitions cover those things for which we are unable to get scientific data.

forced to find alternative means of measurement. One such approach is by attempting indirect measurements: find proxies for the properties that we are after, and inferring from those proxies to claims about our target properties – goodness and rightness.

The most easily available such proxy is our evaluative beliefs and interests. One at least hopes that our beliefs, attitudes and interests are correlated with what we think is right and good. One feels it when they have been wronged, has some sense of when they are flourishing, and can find qualities in others that they admire.⁵⁷ It is not clear if there are physical or material correlates that are reliable: good people aren't more or less attractive than the general population, and it is doubtful that they are notably richer or poorer, taller or shorter, or manifest any other physical differentiator. Because there is a dearth of physical proxies for our attempts at measurement, and it seems that we have at least some traction with beliefs and desires, we are thus pushed towards the use of evaluative beliefs and desires as our source of evidence in moral theory.⁵⁸

⁵⁷ Of course, this need not be (and is not!) a perfect correlation, and some may be more reliably sensitive to these issues than others. There may well be psychological tendencies or limitations that help determine such sensitivities.

⁵⁸ As previously discussed, beliefs can be either factual or evaluative. Here I take it that factual beliefs can be derived by empirical or scientific means, which support, but do not necessarily logically entail, evaluative beliefs. More formally, I am treating factual beliefs as generated signals, and evaluative beliefs as interpreted signals. The world generates signals that we can gather, and after some processing, we produce interpreted signals. Thus, we all receive the same noisy generated signal, but there may be a diversity of interpreted signal output.

Now that we have established that the relevant sources of evidence in moral theory are our evaluative beliefs and desires, we should take stock of how this might align the view with other existing theories. In particular, because the view seeks to aggregate the beliefs and interests of individuals, one might find some similarity to omniscient observer views, such as those found in Hume, Smith and Bentham. On these accounts, there is an ideal observer who can see what everyone wants and what would maximize the utility of the whole population, and on this basis determine what the right thing to do would be. Though there are some similarities between my view and the omniscient observer tradition, such as the fact that both aim to aggregate individual interests, the view being presented here is importantly different. The view from everywhere seeks to avoid the fiction of an individual who can know and see all: in this view, the beliefs of individuals are aggregated, but there is no additional information available. We are not to imagine some observer who can see all of the morally relevant facts, and then determine what we ought to do based on what this observer would calculate as best for everyone. Just as we have no ability to directly measure goodness or rightness, we cannot assume that some fictitious entity can make those measurements for us. We are limited by what we collectively know, and any inferences that we might be able to make on that basis. Even if there are pertinent facts that could be known, if no one has any beliefs about them, we do not have any access to them.

Now that we have a better understanding of what the evidentiary role of beliefs is in the view from everywhere, we can now turn to the next component: the role of perspectives.

The Role of Perspectives

That the view from everywhere aims to take beliefs and interests into account is not startlingly novel: Rousseau outlined a belief aggregation procedure in *The Social Contract*, and since then the concept has at least been known, if not always embraced. And since then, a regular worry about trying to take these beliefs and interests into account has been that there are no judgment aggregation procedures that can guarantee to arrive at a single result – the will of the people – that satisfies very basic constraints. Diverse beliefs and interests create problems precisely because it is so difficult to find a means to aggregate them into something coherent. We are interested in beliefs and interests because they are supposed to provide us with some evidence of what we ought to believe.⁵⁹ But just the beliefs

⁵⁹ This can be seen most clearly in a market setting. Consider the spot price of a commodity like wheat. This price is determined based on hundreds of thousands of small-scale exchanges that determine the future value of wheat. This includes information about how large people think the harvests will be, what demand will look like, whether there are any weather or pest concerns, and a host of other relevant considerations. Though any single individual has a more or less reliable belief about some of these factors, the final price represents what one ought to think about the future availability of wheat. Similarly, a jury's aim is to determine what we ought to think about the innocence or guilt of the defendant. The aggregation of the views of the individual jurors has greater weight than their individual views.

and interests themselves leave us with a morass. To begin to work through this mess, we must take individual perspectives into account.

Perspectives are simply the filters that we use to view the world. This supposes that we do not take in the world “as it is”, but instead (consciously or unconsciously) choose to group certain features together, choose to ignore certain information while focusing on other information, and choose systems of representation and interpretation. Perspectives are thus mental schemata that embody those choices.⁶⁰ A perspective can be as trivial as standing on one side of the room rather than another, or it can be as complicated as taking on substantive theoretical commitments. Cultures will often share a great many perspectives. For instance, an individualist perspective is central to much social reasoning in the

⁶⁰ The notion of mental schemata is found in the psychology literature. Schemata are understood as representations of general knowledge. A schema is a basic theory about the world. This theory is meant to apply to many situations, and includes an associated web of beliefs, attitudes, emotions and expectations for the situations to which it applies. On a beach in the summer one would be unnerved if there were a group of people in business suits having a meeting – it does not fit into one’s schema for beach attire or behavior. The thought of a beach conjures up an image of a generic beachscape, with people swimming, sunbathing, making sandcastles, and enjoying leisure time. One might also think of sunscreen or a Frisbee, as these concepts are associated with beaches. When one encounters a beach, one immediately tries to categorize the activity one sees in terms of these beliefs and expectations. This might even result in people discounting, ignoring, or misinterpreting objects or activities that are inconsistent with this schema. People who share these schemata will then have similar beliefs and attitudes about beaches, and largely interpret a beach scene or beach experience in the same way. If, however, people have different schemata relevant to their interpretation of, say, a philosophy classroom, they can interpret the same scene in very different ways. While some might expect to find a rewarding intellectual exercise, and believe that the room will be filled with interesting and thoughtful people, others might expect to be bored, tired, or otherwise disengaged. Though they are exposed to the same stimuli, they expect, and wind up experiencing, very different things.

United States, where this is much less true in socialist or communist countries.

Group membership might even entail the adoption of certain perspectives that bind the group together – for instance, Sierra Club members share particular perspectives about the importance and salience of environmental issues, Republicans tend to use small government and low taxes as frames for how they interpret political questions, and Americans tend to view international affairs from the perspective of how events affect the United States and its standing in the world. Shared perspectives are a source of social cohesion – they provide a framework for mutual understanding and interpretation of shared events.

In order to better hone the account of perspectives that is required for moral reasoning, we will first consider examples of perspectives in simpler contexts. Once we have seen how perspectives operate in those cases, we can then complete the analysis for what the relevant conception of perspectives is.

Let us first consider how we come to learn about the shapes and spatial location of three-dimensional objects. Imagine a cup on the table. One can come to know the shape of the cup and where it is located on the table by looking at it with two eyes, and possibly moving one's head, if not one's whole body, to examine the cup from different angles. We do this so naturally we do not notice. Each different angle we examine the cup from imparts new spatial information. If for example, we had an eyepatch on, and could not move, we would be limited to a single viewing angle. In this incapacitated condition, we would not necessarily be able to tell if it were a small cup that was up close, or a large cup that was farther away, or even just

a two dimensional image of a cup. Each viewing angle of the cup is a different perspective that we could take on, and each one provides us with information that any other single perspective does not. Since cameras have only one “eye,” photographers and cinematographers can create visual illusions using a technique known as “forced perspective.” It can make objects seem larger or smaller, or closer or farther away than they actually are. If we could look at the same set-up with a slightly different perspective, the illusion would immediately go away, and we would be able to see things for what they really are.

This example is in many ways an apt analogy for what the view from everywhere aims to do. We often have nuanced and disagreeing views on what the appropriate set of moral principles for society would be. In the face of this, we need to gather as many perspectives as possible so as to be able to get a better picture of what these nuances are. Unfortunately, however, vision is in some sense an inapt analogy: after all, we have already seen that our beliefs and interests aren’t direct reflections of our moral principles. We don’t have some special ability to see moral principles as bare concepts. Instead we can only see their effects on other things that we can see. Plato seems to be right – we’re only seeing shadows on a wall. If this is the case, how can we think of perspectives in a way that might be able to give us the same kind of traction that we found in the case of vision?

For such an example, we shall turn to theoretical commitments. Even in the hard sciences, which are meant to be an archetype of the view from nowhere, we can find examples of perspectives (in this case theoretical commitments) shaping

our beliefs. Consider the example of understanding the evolutionary history of the female orgasm. Standard accounts of the female orgasm assumed that it must be adaptive – that is, it was selected for on the basis of some function it had that increased evolutionary fitness. However, Lloyd has demonstrated that all of the proposed adaptationist accounts lack sufficient empirical support, and in fact the evidence points towards the “byproduct” or “Fantastic Bonus” theory, which suggests that the female orgasm “came along for the ride” from the strong selective pressure for the male orgasm.⁶¹ This demonstration was not made on the basis of new, crucial evidence that had not yet been uncovered, but instead was made on the basis of adopting a feminist (and anti-adaptationist) perspective. Most theorists assumed that the female orgasm was similar to the male orgasm, and thus encouraged the production of offspring. It is not surprising that these theorists were also men. Adaptationists argue more generally that any trait must have been *selected for*: each trait exists because it independently enhances the fitness of the organism. The task is merely to determine what function it serves so we can understand why it was selected. Traits like female orgasms or male nipples, however, do not seem to serve any evolutionary function, but are merely harmless byproducts of other traits that were in fact selected for.

It is on the basis of a new perspective that a breakthrough was made. It was not from a maximally neutral position, nor from knowledge of universal truths, nor was it on the basis of new evidence. It was based on the examination of the same

⁶¹ For a detailed discussion of this, see Lloyd, Lisa. *The Case of the Female Orgasm: Bias in the Science of Evolution*. (Harvard University Press, 2005)

evidence that everyone else had access to, but that evidence was seen in a new light. This is not even necessarily to say that a feminist or anti-adaptationist perspective is always going to be most useful: an adaptationist would be able to provide a very good explanation of the evolutionary history of female nipples or male orgasms. Further, it is entirely likely that neither of these two perspectives – at least by themselves – would be adequate for understanding some other phenomenon. As with the example of vision, it is not the claim that there is a single privileged perspective, but rather that any single perspective is going to be inadequate to fully understand some phenomenon. Additional perspectives can often reveal new information, even if they rely on the same evidence.

This idea that there is no single best perspective for approaching all problems is well known in the computer science literature. The No Free Lunch Theorem in search and optimization demonstrates the impossibility of a single best procedure for problem solving across all problem domains. In fact, there is no procedure that can do better than any other procedure on average, when considering all problem domains.^{62,63} This is a basic constraint that has not been fully appreciated in less technical contexts. Importantly, however, it helps us more

⁶² “We show that all algorithms that search for an extremum of a cost function perform exactly the same, when averaged over all possible cost functions. In particular, if algorithm A outperforms algorithm B on some cost functions, then loosely speaking there must exist exactly as many other functions where B outperforms A.” Wolpert and Macready, “No Free Lunch Theorems for Search” (Technical Report SFI-TR-95-02-010). *Santa Fe Institute*, 1995

⁶³ Hong and Page’s “Problem Solving by Heterogeneous Teams” *Journal of Economic Theory* 2001 demonstrates that there are equivalences between encodings of problems and heuristics. They show that across all problems, there is no best encoding, given a fixed heuristic.

fully see the role that perspectives play. As we have seen, perspectives are mental schemata – they are bundles of expectations, judgments of salience, interpretive norms, and emotions for classes of situations. One can see this as a combination of two basic parts: first is the imposition of an ontology to a class of situations, and second is the heuristics or procedures one uses to evaluate those situations.⁶⁴ The imposition of an ontology is just the determination of how one individuates and groups objects and phenomena, and the basic classifications that are set. As an example, foods could be classified based on size, shape, color, texture, nutrition, freshness, price, or any number of other dimensions.⁶⁵ Ontologies can be fine- or coarse-grained, and may be a mixture of the two. For instance, an ontology of people might include very fine-grained distinctions for members of the in-group, but very coarse-grained distinctions for everyone else.⁶⁶ There may be many cases of two ontologies A and B, where B includes all the distinctions of A but also includes additional distinctions. In these cases, the perspectives that include these

⁶⁴ The first systematic discussion of categorization and schemata in the context of social and moral philosophy is in chapter 2 of Bicchieri's *The Grammar of Society*. There she develops a detailed account of how these two mental processes pre-consciously shape our behaviors, particularly in relation to norm adherence.

⁶⁵ Note that there is a possibility for one ontology to be a superset of another – that is, for one ontology to contain all of the distinctions and categories that the first makes, but also additional finer-grained distinctions. These two should be viewed as distinct ontologies, and thus the basis for distinct perspectives. There is the potential for considerable overlap between these two perspectives, which has consequences for their independence. I deal with this in the next section.

⁶⁶ This provides a framework for looking at stereotypes. Stereotypes involve very coarse-grained generalizations about classes of “other” people. A simple example is how the media discusses “the Hispanic vote” as if it were a single bloc of voters with identical interests, when often the subgroups have conflicts with each other and recognize a great number of distinctions. These individuals don’t typically define themselves as “Hispanic,” but instead as, say, “Dominican-American,” “Puerto Rican,” or “Mexican-American.”

ontologies may well disagree when those finer distinctions are relevant to the question at hand – even if the same evaluative function is in use.⁶⁷

Once an ontology is established, heuristics can be employed for determining, say, which fruits to buy, which to paint, or which to produce. We might notice that a classification based on nutrition would be fairly useless in deciding which food to paint, but might be very useful when determining what to eat. What's more, even when a single ontology is appropriate for two different cases – say, price is an important factor in deciding what to eat or what to produce – different heuristics become relevant. When consuming, cheaper food might be more appealing, whereas more expensive food might be more appealing when choosing what to produce. Thus in a sense our perspectives function very much like search algorithms: they are specialized for use in particular contexts and from within the ontologies that they are bundled with. As a result, they cannot be seen as general-purpose: other perspectives will outperform them in other problem domains.

⁶⁷ For example, even within a single religion, there are different levels of commitment that can create finer or coarser ontological distinctions. A category of “food” might be further refined to include “acceptable food” and “unclean food”, marriage customs might have finer distinctions, etc. In this scenario, there is no reason to make the claim that one or the other ontology is the “correct” ontology, such that one perspective is subsumed in the other. The perspectives are surely related, but there are evidently important differences between them. There are no a priori lower or upper limits to the number of possible distinctions or categories in a given perspective's ontology, outside what our cognitive limitations are at the upper end, and some sense of utility at the lower end. As such, it is unlikely that perspectives would be useful with ontologies that have only a single category, and it is unlikely that humans can cognitively handle perspectives that make millions of distinctions.

A perspective, then, can be seen as a mental schema that has two identifiable components: the classification stage where the general ontology is imposed, and the heuristic or judgment stage where choices and evaluations are made. We have seen that there is no single best perspective, nor can there be, due to the limitations imposed on us by No Free Lunch theorems. So now we can see that we can benefit from multiple perspectives because of the need to cover multiple problem domains. But does this mean that we ought to only have a single “best” perspective in each domain? This was not true in the case of visual perception, and it is not true generally. While we may be able to find the single “best” perspective in a given problem domain, it is not to say that we cannot do better by augmenting that perspective with alternative perspectives that rely on different ontologies and different heuristics. These other, “lesser,” perspectives will not fully overlap with the “best” perspective, and thus potentially be able to improve our overall judgments by covering areas missed by the “best” perspective. We can consider the example of diagnostic testing to see why this is the case. Imagine that there are two diagnostic tests for a dangerous disease. Test A has 80% accuracy, and Test B has only 40% accuracy. But if Test B works on different principles than Test A, taking both tests can boost the accuracy of diagnosis because they each capture a different aspect of the diagnosis. So perspectives working together can provide greater epistemic benefits than just relying on a single perspective.⁶⁸

⁶⁸ In fact, it is also true that combining many poor performers together can outperform the single best performer. As Scott Page and Lu Hong have put it, “Diversity trumps ability.”

Now that we have seen how perspectives can impact our daily lives and the scientific enterprise, we can turn to how they affect our moral reasoning. As in the scientific case, we are trying to marshal evidence to understand what moral principles we ought to endorse. And just as we cannot directly observe a trait's evolutionary history, we cannot directly observe moral principles. Instead, in both instances we must find proxies for what we are interested in. In the case of the female orgasm, if the adaptationist story were true, we ought to find women who orgasm more easily would have more children, but this is not the case. And as we have seen in the last section, our proxies in the case of moral reasoning are our beliefs and interests.

So, we must rely on our perspectives to help make sense of our moral beliefs and interests. This can prove to be a difficult thing to do, as our perspectives help to shape our moral beliefs in the first place. It is very difficult for a single individual to be sufficiently self-reflective to be able to tell that they have caught themselves in a mutually-reinforcing relationship between their moral beliefs and their perspectives – especially when those perspectives are theoretical commitments. If our theory helps us interpret our evidence, and we take our evidence to confirm our theory, then how are we to break out of this cycle?⁶⁹ The only path out is by means of introducing new perspectives. This can be done either by an individual or by a

⁶⁹ This is characteristic of the worry about the reliability of reflective equilibrium as a method of determining or justifying one's moral beliefs. There is nothing in the process of reflective equilibrium that can provide a check against potential biases imposed on one's thought process by the perspective one has. The challenge is being able to identify those cognitive tendencies and work around them if they are liabilities. But this cannot be done without some external check that determines whether such biases are there.

group. While individuals learn to take on new perspectives over the course of their lifetimes, I would claim that most often, these perspectives come from exogenous factors. Going to college, taking a class in an unfamiliar subject, traveling, and meeting new people are what I suspect are the typical ways one learns to foster new perspectives. While it is surely possible for one to develop perspectives on one's own, this is usually a significant breakthrough, and not an everyday phenomenon.⁷⁰ So while I do think that any given individual has a toolkit of many different perspectives to draw from, I do not think that it is possible for any individual's toolkit to be so large so as to incorporate every possible (or even currently-existing) perspective. As a result, we must rely on groups to be able to generate as many perspectives as we can so we can improve our epistemic abilities.

Even if we have successfully managed to break out of our own narrow perspective, and are no longer stuck in a rut of having our perspectives create our beliefs which then confirm our perspectives, we still have not yet dealt with how the view from everywhere takes these beliefs and perspectives and outputs impartial moral claims. For this, we must now turn to a discussion of the aggregation procedure.

⁷⁰ Being able to conceive of, say, how the universe looks to a photon, as Einstein did, is not a common exercise. At least for some, meeting someone from a different background and learning about how they see the world is.

The Aggregation Procedure

It is in the aggregation procedure that we can see how the view from everywhere can be taken to be a moral stance, in that it provides a framework from which to reason about moral questions. Thus far we have seen why beliefs are the evidence that we have access to in moral reasoning, and perspectives have a large role in both generating and interpreting those beliefs. Unlike in the view from nowhere, which only takes in objective, impartial information as part of the reasoning process, so that the “output” is also impartial, here we take in biased, subjective information. The task ahead of us is to outline a procedure that can take this subjective information and use it to create impartial moral claims.

For simplicity’s sake, what I will describe assumes that there is only one issue in discussion at a given time, and that potential beliefs about this issue can be mapped onto a linear scale. For example, an action that is under moral consideration could be “downloading music without paying”. We can then imagine a scale from “the worst possible moral offense” to “supererogatory action” onto which our beliefs about the activity can be mapped.

If we were to follow Rousseau and employ something like the Condorcet Jury Theorem, what we would do next is collect everyone’s beliefs, and determine the value to which they converge. The idea behind this is that if we assume that everyone is at least slightly better than chance at knowing the truth, we will find that as we add more people to our jury, we will become increasingly more likely to

arrive at the true answer.⁷¹ While probability theory confirms the veracity of the theorem, the challenge is determining when we have the appropriate conditions for applying it. Namely, the jury theorem requires two things of the agents participating: first, that they will be better than chance at picking the true answer, and second, that each individual's choice is independent.

For many questions at least (though certainly not for all), we can safely claim that individuals are better than chance at picking the true answer. But on the second condition – that individual choices must be independent from each other – we are much less fortunate. There are two sources of difficulty here. First, as we have seen, people who share perspectives will have correlated beliefs. But even more basically, people talk to each other. Even in juries, jurors are required to talk it out and try to convince each other that they are right. This will tend to violate the independence condition. Nor does it seem particularly desirable to take steps towards trying to enforce such a condition. If people are not allowed to talk to each other or substantively interact, there doesn't seem to be too much of a point to having a society in the first place. So it seems that, even though the Jury Theorem is epistemically a very desirable procedure, we cannot employ it and continue to enjoy living in society.

⁷¹ The Condorcet Jury Theorem was meant to be used for binary choices about matters of fact. Here, however, we are dealing with evaluative beliefs, which make no truth claims themselves. As a result, we are forced to shift to an alternative standard, "well-supportedness." The aggregation procedures discussed in this section thus make claims about which evaluative belief has the most support, which is independent of a truth claim.

Though we cannot adopt the Condorcet Jury Theorem, we can learn from where it falls down to develop a procedure that can capture most of its epistemic virtues while also having less burdensome assumptions. Most importantly, this means giving up the independence assumption. In doing this, however, we open a fairly large can of worms. Namely, even if everyone in society agrees on the same moral principle, we may not be able to endorse that principle. To see why this is, recall the concept of forced perspective. If everyone looks at something from the same perspective, they will arrive at the same false conclusion about what is going on. Everyone watching Lord of the Rings thought Gandalf was much taller than Frodo, but this was just the effect of the actor playing Gandalf standing closer to the camera.

For these reasons, we cannot take just beliefs into consideration: instead we must look at belief-perspective pairs, and further, the distributions of these pairs. As we have seen, perspectives provide both the ontology and the method of evaluation for a given situation. Thus, we can think of perspectives as lines of support for beliefs. They provide a framework for understanding and justifying the beliefs they produce. Mill provides a precedent here for why relying on beliefs alone is not enough: his defense of the freedom of speech in chapter two of *On Liberty* argued that the role of speech was to help us learn the arguments for our beliefs. We ought not just hold dogmatic beliefs – even true beliefs – but must also know why we should hold these beliefs. Perspectives play a similar role: they tell us why we hold certain positions. What's more, they play a role in our judgments of

epistemic normativity: given that one holds such a perspective, one should be led to certain kinds of beliefs and judgments. So when we seek to aggregate our beliefs, we must tie that aggregation to the perspectives that helped to generate those beliefs.

The conceptual shift I aim to make here when departing from the Condorcet Jury Theorem is that, where the Jury Theorem provides one vote to each individual, I aim to provide one vote to each perspective.⁷² Because a perspective is both an ontology and an evaluative procedure, each perspective is compatible with only one set of beliefs. Each perspective receives equal weight, regardless of how many people share it. If an entire population shared only one perspective, for example, there would be only one vote to count. This looks vehemently anti-democratic on the face of it, but recall this procedure is meant to be epistemic, and not political. The goal is to determine the set of beliefs that have the most independent⁷³ lines of

⁷² This move is meant to account for the problem of independence with individual voting. By counting perspectives, which essentially represent lines of argument, we can guarantee that we are counting independent judgments. This continues to capture the spirit of the Condorcet Jury Theorem, but instead identifies perspectives as the carriers of beliefs instead of individuals.

⁷³ This claim of independence relies on our ability to identify different perspectives as being different, and further, that different perspectives are non-overlapping. The first challenge to this – that we are able to identify different perspectives, is an empirical assumption. Even if one cannot identify one's own perspective, it is likely that someone with a different perspective will point it out. The more difficult question is whether or not any two given perspectives are always independent of each other. One way of assessing this question is by recalling that perspectives have both an ontological component and an evaluative component. If two perspectives have different ontologies, then they are independent. If they share ontologies but have different methods of evaluation, then statistical comparison across numerous cases will determine their level of correlation. In these cases of dependence, then the value of the vote for each dependent perspective should be reduced to match the increase in information it provides

argumentation supporting them, not those that are most widely held. These are the moral principles that we can then take for granted at the beginning of the political process. Part of the task of the political process is to determine whether to act on those principles for which there is not much epistemic support. The task of the next chapter is to develop this process.

Now we can turn to a discussion of the different possible cases of the distributions of perspectives and beliefs. First, let us consider a case of agreement. As discussed earlier, this may be due to the fact that everyone has the same perspective. So though there is agreement, when we have agreement based on the same single perspective, we ought not to place much credence on the fact of agreement. After all, we essentially only have one vote. Jury Theorems require large numbers. As a result, we should have no faith that a single vote is reliable. As such, rather than endorsing the agreed-upon moral claim, agents in the view from everywhere would have to seek out alternative perspectives to provide confirmation that their agreement would continue in light of novel perspectives. With only one perspective, we have no way of testing whether our beliefs are robust against alternative ways of thinking about and evaluating the situation in question. Robustness against alternatives is the sole check the view has against moral relativism. A moral belief needs to be able to withstand the judgments of multiple perspectives to have any weight. Just as scientific practice finds value in robust

over those perspectives on which it depends. As this is a statistical phenomenon, it is likely that these values will have to be adjusted over time, as better information allows for more refinement.

results which are supported by multiple models, so too does the view from everywhere buttress moral beliefs by means of multiple perspectives. A moral belief supported by a single perspective is only as strong as the weakest assumption in the perspective. A moral belief supported by multiple perspectives is stronger than the status of any individual perspective.⁷⁴ A robust moral belief, then, remains stable against the introduction of novel perspectives, and the elimination of faulty perspectives.⁷⁵

Alternatively, we could find agreement amongst individuals with several different perspectives. The different perspectives would then provide a consilience of evidence in support of the same belief. In this case, the agents in the view from everywhere can claim that their agreement is evidence of a robust moral belief. Now that we have several active perspectives, we have a jury that is in agreement. Our beliefs are robust against perturbation in how we classify the situation and how we try to evaluate it. This suggests that we have found principles that are and will remain well supported. These cases are straightforward, but this is largely because

⁷⁴ This idea was developed in the context of scientific models in Wimsatt's "Robustness, Reliability and Overdetermination", found in his *Re-Engineering Philosophy for Limited Beings* (2007).

⁷⁵ It is important to emphasize that the moral beliefs at issue here are token beliefs, not type beliefs. A belief like "everyone should be treated equally" needs to have accompanying definitions of "everybody" and "equally" for it to count as a moral belief. Sen importantly raised this issue in chapter 1 of his *Inequality Reexamined* (1992). All mainline political frameworks have some notion of equality, but they pick different substantive theories on which they define equality. So while they are structurally similar and may agree on a type belief, they will disagree on token beliefs, since those token beliefs will have fleshed out the concepts involved. Thus type beliefs cannot be robust because they are not substantive – only token beliefs can be subject to robustness checks.

agreement is easy to deal with. The more interesting cases are those in which agents disagree with each other.

For simplicity's sake, I will consider disagreement to be represented as a bimodal distribution in beliefs. This is of course but one way in which substantive disagreement could manifest itself, but it is a fairly common case: it represents two opposing positions that each have a large portion of the population as supporters. An empirical example of this can be found with attitudes of fairness in the Dictator Game. In this game, one of two players – the dictator – is given \$10, and told that they can give as much as they want to the second player, and keep the rest. The dictator's decision is final and determines the distribution of money. Cristina Bicchieri and Jason Dana asked students enrolled in introduction to philosophy at Carnegie Mellon University what they thought was a fair amount of money for the dictator to give in a Dictator Game. The results were a bimodal: 68.3% of students thought that equal split was the fair allocation, while 21.4% thought that not giving anything was perfectly fair. Slightly less than 20% thought that not giving anything was unfair.⁷⁶ Given this bimodal distribution of beliefs about fairness, we can imagine two possible scenarios: first, that these beliefs are correlated with particular perspectives, or that there is no correlation between beliefs and perspectives.

In the first case, we have substantive moral disagreement, but each side marshals a different set of arguments. This is precisely the case in which moral

⁷⁶ Details of this study can be found on pp. 137-139 of Bicchieri's *The Grammar of Society* (2006).

deliberation ought to be most productive. The agents come to different conclusions, but this may very well be because they see the world differently. Just as our case with perspectives in science, there is an opportunity here for an exchange of perspectives that might help to settle the moral dispute. It could be that one side is correct, or it could be that both sides are correctly capturing some portion of the situation, without realizing that there are other components. In the first instance, we find ourselves in a situation not unlike the case of the female orgasm that we considered above. One side's perspectives are in fact ill-suited for the particular situation at hand, while the other side's are much better calibrated. In the second instance, we find ourselves in a case similar to that of medical diagnostics: each set of perspectives is useful for some portion of the problem, but only by combining them together do we get a more complete picture. It is thus possible that the exposure to the other perspectives can enable the side with less robust moral beliefs to see how fragile their beliefs are in the face of new perspectives, or for both sides to use the information gleaned from the other's perspectives to find that there is a better, third position that they both can agree on. As such, deliberation amongst people with these diverse views, not over which perspective is best to take in every case,⁷⁷ but in a particular instance, has great potential for getting to the correct view. This does not mean that there will be agreement in every instance, but from a purely formal epistemic perspective, this is the case for which there is the most useful evidence to leverage. As the numbers of

⁷⁷ After all, as we have seen already, the No Free Lunch Theorem shows us that there is no such perspective.

active perspectives increases, we in fact fall back into our modified version of the Condorcet Jury Theorem. As such, we ought to have increasing confidence in the fact that a straightforward aggregation mechanism will yield the most reliable result.

In our final case, we find that we again have a bimodal distribution, but this time, beliefs do not correlate with perspectives. Both positions, then, have what appears to be a consilience of evidence in their favor, except for the fact that the other side has similar levels of support. That is, both sides marshal the lines of argumentation to defend their beliefs. They both appear to be endorsed by the same set of perspectives. Here is where the view from everywhere differs most drastically from the view from nowhere. Since each perspective appears to endorse multiple beliefs, no Jury Theorem can operate.⁷⁸ It would be like allowing individuals in the Condorcet Jury Theorem to cast two opposite ballots – there could be no convergence in the aggregation. As there is no formal procedure to use that can leverage perspectives to help us to inform an understanding of where our beliefs came from, and thus enlighten the discussion, we are stuck at a situation of substantive moral disagreement. The agents in the view from everywhere have no logical means to convince each other, since they share perspectives already, and so further moral deliberation will prove to be fruitless.

⁷⁸ Note that this is not a case of a finer-grained perspective disagreeing with its coarser-grained counterpart. Instead, these are identical perspectives that ostensibly say two different things.

In this situation, we find a deviation in the supposed epistemic relationship between beliefs and perspectives. In the three other cases we considered, we supposed that a perspective, combined with epistemic rationality, constrained the individuals who took on those perspectives in what they could believe. That is, a perspective, as taken to be a mental schema that provides an ontology and a method of evaluation for a given class of situations, will tend to cause individuals to have a certain set of beliefs. Those beliefs can be thought of as being entailed by the perspective.⁷⁹ In this case, however, the constraints on beliefs that perspectives typically provide are somehow violated. Thus, perspectives lose their role as guarantors of evaluative independence, as this entailment relation is broken.

As such, all the view can do is represent what the views that are under active consideration are, and describe the nature of the disagreement. These positions are in some sense reason independent – given that the same strains of evidence support both positions, it suggests that the disagreement is not actually based on the evidence itself, but just some kind of fundamental arational disagreement.

⁷⁹ The relationship must actually be something other than entailment, as a perspective is not a set of propositions. Further, there is likely to be at least some natural variation amongst individuals who share the same perspective. I see perspectives as constraining what that variation can be, and providing the justificatory framework for a given set of beliefs. Epistemically rational agents will then tend to have those beliefs for which the perspective provides justification.

Conclusion

As we have seen, the view from everywhere aims to take everyone's biased views about the world, and create from that an impartial moral claim. By taking advantage of both our beliefs and our perspectives, the view aims to uncover the whole picture out of the pieces of it that we all have. In doing so, however, the aim does not suppose that we have answers for everything – in fact, different instantiations of this view will have varying abilities, depending on the makeup of the individuals participating. What's more, the view can suggest to us that we either do not know enough to make certain moral claims, or that our support for some positions might not actually rely on any evidence or rational support. This is an unpleasant result, but the value in this kind of result is that it can point to the limitations in our moral epistemology. Rational deliberation cannot always ensure that we will end up agreeing with each other, nor can it assure that we will always hone in on the best principles. The view from everywhere merely identifies for us where we can come to know what moral claims are well supported, and where we are currently unable to proceed.

Chapter 3:

Justice Without Agreement

Introduction

In the previous chapter, we examined the epistemic benefits of leveraging multiple perspectives in our efforts to find well-supported moral principles. This forms the basis for our task in this chapter, which is to find a procedure for determining the set of rights afforded to individuals in society. While it is possible that individuals participating in the view from everywhere will find themselves in broad agreement over the basic structure of their society, this chapter does not assume that such agreement has been found. In fact, the aim of this chapter is to develop a procedure for determining the set of rights afforded to individuals when those same individuals agree on very little. In fact, I will consider the case where people not only disagree about what their preferred outcomes are, but they also disagree about what the set of outcomes even looks like.

Sen's Insight

In the first chapter of *Inequality Reexamined*, Sen noticed something quite important. When one considers the major conceptions of justice, every one of them has some notion of equality. It is just that each account of equality is different – each

theory is measuring something different. Libertarians are concerned with equal negative liberties. Egalitarians are concerned with equal outcomes. Utilitarians are concerned with equal consideration in utility maximization. If an Egalitarian were to look at either a utilitarian or libertarian regime, she would likely find significant inequalities in individual outcomes, but when those societies are viewed by their own members, they might see them as models of equality. It is not that the egalitarian is mistaken in her view, but simply that she is using a different standard of equality than the utilitarian, who is in turn using a different standard than the libertarian.

With this insight, Sen shifts the debate. Rather than arguing about whether equality matters, we find that everyone is in fact arguing about which political values matter. Libertarians assert that freedom is the most important value. Egalitarians think that socio-economic outcomes matter the most. Utilitarians claim that happiness is the most important value. Each theory, then, develops a way of looking at the world which accords with the values that matter the most to that theory, including standards of comparison. Unfortunately, once Sen makes this remarkable insight, he goes on to simply propose a new measurement of equality - capabilities. While his capabilities approach is interesting in its own right, much more can be mined from the original insight.

A way to look at what Sen noticed is that each political theory is a representation of a particular *perspective*. This perspective categorizes the world in terms of the values that the theory holds dear. It also provides a judgment

apparatus built around this categorization. As such, these political theories carry with them perspectives that shape our reasoning. Not only do these perspectives shape our preferences over potential political outcomes, but they also determine what we see as the outcomes. For example, a Libertarian might see a political debate in terms of two outcomes: “loss of liberty because of increased government intervention” and “increased liberty because of less government intervention” and thus prefer the latter option. Meanwhile, an Egalitarian might see “free government-provided health care for all” and “elimination of social safety net,” and then prefer the former option. Both preferences are informed by political values, but more importantly, those preferences appear to be over *different outcomes*. The libertarian and the egalitarian not only disagree about what’s a better choice, but how they see those choices in the first place.

The previous case is not even a particularly problematic example: both parties here agree on where the contours of the two outcomes are, but see them under different descriptions. This will not always be true. For instance, consider a vegan, a vegetarian, an omnivore, and an omnivore that keeps kosher. These individuals wouldn’t share the same preferences about what to eat for dinner, but more importantly, they would have different categorizations of the possibilities. A vegan would largely see the world in terms of “things that are vegan” and “things that are not vegan.” Vegetarians would draw the line differently. Omnivores have no particular need for such distinctions, while someone keeping kosher would look to distinguish between meats, based on which are prepared and served

appropriately. A vegan doesn't even have to think about whether something is kosher or not – everything that isn't kosher is already not vegan. These categorizations will naturally affect preferences over choices, but they also cause the different parties to see different choices altogether. Whereas the omnivore can have developed views about whether pork belly is better than bacon and pork chops, the other eaters would simply lump all of those options under a “can't eat” category. Likewise, if someone wanted to require that all butcher shops in the country must be halal, no one but the omnivore would have reason to care, as it would only affect foods that they already wouldn't eat.

An implication of this view of perspectives is that once we have chosen a particular political theory as the basis for our political discussion, we have cut off the ability of people who have a different perspective to even properly represent their own views. It is not merely a matter of what their preferences might be, but how they see the problem. Choosing to privilege a single perspective above the others means that the values embedded in those other perspectives cannot be fully articulated.⁸⁰ This means that once a perspective is chosen, much of the substance

⁸⁰ Donald Davidson argues against this position in “On The Very Idea of a Conceptual Scheme” *Proceedings and Addresses of the American Philosophical Association*, Vol. 47, (1973 - 1974), pp. 5-20. He claims that if we think of conceptual frameworks as languages, then the claim is either trivial or incoherent. If we merely mean that we can only assign truth claims within a linguistic framework, then this is a trivial claim. However, going more in the Kuhnian direction is incoherent, since he claims that we can only make sense of different points of view by reference to a common “coordinate system” which allows us to map the differences. But then, given such a coordinate system, we can establish a translation from one conceptual scheme to another. I would respond, however, by suggesting that Davidson is himself stuck within his own perspective, that of linguistic-turn philosophy. Instead of considering conceptual

of political disagreement has already been shut down. Once this has been done, the way we consider what the outcomes are and what our options look like have already been determined. Following Sen, this means that the debate over which values matter the most has already largely been completed. It is no doubt a fact of psychology that individuals are going to be more prone to view the world in terms of one perspective or another. If anything, this should bolster our interest in preventing a single perspective from being the only appropriate frame for political discussion in the public sphere. While many may share the same perspective, it would prevent others who have a different perspective from being able to fully participate in political discourse.

What this suggests is that instead of choosing another measure of equality, or another set of values to prioritize for social organization, we need to turn to a model that allows for political agreement amongst individuals who hold any arbitrary set of perspectives. It is important to note that this is not to say that developing new perspectives is not valuable: the view I present commits me to the idea that we can benefit from the development and presentation of novel perspectives. However, anointing a single perspective as somehow neutral or primary is problematic for the reasons I have described. As such, we must then move towards a model of political

schemes as synonymous with languages, we could imagine them as mathematical objects known as projections. A projection takes a higher-dimensional object down to a lower dimension, and in the process, some information is necessarily lost. Even if we start with the same high-dimensional object, and then take two different projections of it, we cannot translate between them because even though they started from the same source, they each differ in what information they retained. So no direct mapping between them is possible. And thus, had Davidson had an alternative conceptual scheme to work from, he could have seen that they are very much possible.

agreement that attempts to combine these disparate perspectives in a way that allows all participants to be enfranchised.

What is a Social Contract?

Before we proceed, it is important to articulate what exactly it is that we are aiming for in a model of political agreement. In particular, we are interested in developing a social contract. So we must first ask ourselves what the goals of a social contract are, and then we can see how to go about satisfying them. In doing so, we can identify the necessary components of a social contract theory, and then best determine how to satisfy them, given our commitment to enfranchising all perspectives.

If we consider what is most fundamental about a social contract, we can see it as at minimum an attempt to design a set of rules and assurances to further the goal of allowing individuals to successfully live together in society. In particular, we can imagine ourselves as seeing the potential value of living together, but we need to establish the basis for our social cooperation. It is this need to establish the basis for social cooperation that led the classical social contract theorists such as Hobbes, Locke and Rousseau to postulate a “state of nature” where humans are pre-contract and in turn, pre-social in the sense that there is no cooperation.⁸¹ These

⁸¹ Note that in Locke’s conception, the state of nature is not entirely pre-social. In fact, individuals have basic property rights and other things that we would normally associate with a social contract.

philosophers were different in how they pictured the state of nature, but they all argued that individuals would all be motivated to enter into a social contract with their fellows. From this, we can draw our first conclusion about the nature of the social contract: that we can view the social contract as a mutually-advantageous agreement designed to facilitate social cooperation.

Mutual advantage might initially seem like a loaded term aimed at skewing the discussion, but I want to argue it is merely a minimal, necessary condition of a social contract. If this basic condition were not being met, it would imply that some subset of the population is being harmed by remaining in society. This suggests that either those people would not have entered into the contract in the first place, or would leave if given a chance. A rational agent would not join a society if it harmed her unless, like a slave, she were forced to join society. For Hobbes, the absence of mutual advantage makes social contracts impossible, as it would undermine his own argument. The key legitimating feature of the Leviathan is that she enables cooperation amongst the ruled, and an escape from the “nasty, brutish and short” life in the state of nature. Even Rawls needs mutual advantage, as individuals in the Original Position are still rational agents, and would not choose under the veil of ignorance a social contract that leaves the worst off in a worse position than some alternative. Even the rich would want to be in some social contract, as it assures them protection of their property and their persons against attacks of those less well-off, as well as insurance against dramatic reversals of fortune.

The social contract must enable rational agents to cooperate with each other and in doing so make themselves better off than they would have been otherwise. This is a purely individualistic requirement. It follows from agents being rational that they will not accept a deal that makes them worse off than before.⁸² Mutual advantage is thus really a collection of individual advantages, since no agent needs to be concerned with whether the other agents view themselves as better off than before – each agent watches out for her own interests.⁸³

Mutual advantage, thus minimally defined, also implies that individuals will abide by the contract only so long as it is to their advantage to do so. If subsets of the society would be made to be better off if the society split up into multiple societies, then society will split. The classical social contract theorists do not discuss this, but it is a consequence of the idea of mutual advantage, and one that I will return to at greater length later.

Mutual advantage is a nice principle, but in a sense it is without content for the purposes of developing social contracts. After all, we need to specify *how* we can enable cooperation. Social contracts enable cooperation by means of providing a set of rights to citizens. Even in the minimal Hobbesian model, the Leviathan ensures

⁸² In the following chapter we will see this principle, which I call the no-subsidy principle, elaborated in greater detail.

⁸³ It should be noted that this individual focus is independent of whether agents view advantage in a *relative* or *absolute* sense. That is, whether individuals compare themselves to their neighbors to assess how well they are doing. Even if an agent has a relative sense of advantage, she does not have to care about whether her neighbor considers himself to be better off – just that she sees herself as better off. So even if she is concerned with relative losses or gains, she does not have to care about her neighbor's perceptions of loss and gain.

property rights and rights to bodily integrity. Rawls has a much more expansive account of rights, but what is common is that in each social contract account, we see that rights are the guarantors and enablers of social cooperation. Social contracts, then, always embed sets of rights. They specify the limits and scopes of rights, including both what the government cannot interfere with and what the government owes its citizens. It is by means of these rights that individuals can assess whether they are worse off or better off with them compared to some other bundle of rights. For example, individuals can assess whether a scheme of individual property rights would be better or worse for them than a scheme of communal property. Individual property rights have a number of incentives-oriented advantages over communal property, though that could be balanced against reductions in inequality and a more communitarian spirit engendered by communal property. Individuals can assess which system of property better suits them.

Now that we can conceive of social contracts as embedding bundles of rights, we have two questions to answer. First, we must think of a procedure for allocation. This determines what rights we allocate as well as who gets them. It is here where we can find the most significant source of differences in theories of the social contract. I will consider two alternative approaches to this question: deliberation models and bargaining models. But we are not done there: even though we may be able to say something about the process for allocating rights, we still need to clearly

specify what we mean by rights, and how to interpret them.⁸⁴ I will proceed by first comparing deliberation models of the social contract to bargaining approaches, and show that bargaining approaches are superior. From there, I will develop a conception of rights to be bargained over. Finally, I will turn to the particular bargaining model that I wish to advocate.

Bargaining versus Deliberation approaches to the Social Contract

The social contract theories of Rousseau and Rawls emphasize a conception of public reason. In this model, moral agents have substantial moral agreement. In particular, the General Will articulates the will of citizens qua citizens. It is meant to be a neutral statement of what the good of society is. Rawls develops a similar conception in the original position, wherein people identify those goods that everyone needs. While Rawls tries to avoid the requirement that agents share in the same comprehensive moral doctrine, he does believe that there are only certain forms of public justification that are appropriate. Public reason is then those justificatory arguments that can be made such that they can be made from a common political point of view. Public reason, like the General Will, is a common political conception of the good, and a system of political justification.

⁸⁴ Consider the example of free speech. We can say that we have a right to free speech, but this right must be fleshed out by specifying the particular instances of speech that are deemed protected. I will elaborate this in detail in a later section.

Rawls saw public reason as the political solution to pluralism.

Comprehensive moral doctrines could not be used for public justification, because in a pluralist society there is no assurance that others will share one's own comprehensive moral doctrine. As such, they would be very unlikely to accept the form of justification, and then there would be no political agreement. By turning to a common political rather than moral conception of justification, Rawls was aiming to suppress the private sphere differences in a society and focus solely on common political commitments. This is not unlike Rousseau's goal of forming the General Will by having citizens reason *as citizens* rather than as self-interested individuals. This is a proposal that has many merits, but ultimately, it has to confront the problems that we have considered above: namely, that this reasoning "as citizen" is not neutral, but instead is privileging a certain conception of values that may not be universally held, particularly as societies become more diverse. This problem will arise with any conception of public reason, precisely because there is no assurance that there will be a sufficient set of overlapping consensus about what may be described as a common political point of view. Our comprehensive moral doctrines and our other private-sphere interests play a significant role in shaping our perspectives, and in turn our perspectives shape our political interests. Though taking part in the view from everywhere might help to establish such a public point of view, we have no guarantee that the overlapping consensus is sufficient for the kind of robust political agreement necessary for public reason.

This problem of diverse perspectives poses a significant challenge to the goal of political agreement. Even though we may be able to uncover some areas of consensus, there will be much that will have no consensus. There are simply too many different areas of political life that require independent consideration to expect that a justificatory framework that can be established in one area can be then extended to another area. Even if we all agree that murder is wrong, that does nothing to help us agree on questions of distributive justice. Likewise, even if we agreed on the principles of distributive justice, that does nothing for establishing freedoms of speech or assembly. This is why a common political point of view would be so valuable if we could have it: it resolves all of these issues, and grants us a systematic framework for political justification.

Thus we are stuck with a problem. We need some method for reconciling competing political goals without having recourse to a neutral political perspective that everyone endorses as a system of public justification. I propose that this problem can be resolved, but to solve it we must move away from a conception of public reason, and instead move towards a model of bargaining. Bargaining has the advantage of requiring very little agreement between the parties involved. When I negotiate the price of a shirt in a Chinese market, I do not have to share many political values with my counterpart: we just have to find a price such that she finds it to be at least equal to her valuation of the shirt, and I find it no more than equal to

my valuation of the shirt.⁸⁵ In fact, my counterpart and I barely need to have a common language beyond numeracy and some display of agreement or disagreement. We certainly don't need to agree on the finer points of the Cultural Revolution, the Great Leap Forward, or for that matter, the Great Society. Nor need we even discuss our views on fairness: the exchange either happens or it doesn't. This is a significant deviation from the public reason model in several respects. Most importantly, we do not have a common system of justification, even though we are typically able to find agreement. Neither of us has to make a claim about some neutral position from which we can evaluate the bargain. I can seek to better understand the needs of my counterpart so I might be able to better determine her reserve price, and she likewise can do the same with me, but in these cases we aren't finding an outside standard—instead we are attempting to better understand each other's own standards. Each party must be convinced on his or her own terms.

This bargaining view has several consequences. First among them is that bargaining outcomes are going to be sensitive to the particulars of the perspectives of the agents involved. This is a natural consequence of the elimination of a neutral standard that both parties share, as each bargain is going to be affected by the

⁸⁵ It may be argued that there are some other implicit agreements that we both have: that money is a reasonable unit of exchange, that you don't just kill the other person and take what they have, etc. These assumptions can be eliminated without much difficulty – we could barter and trade goods instead of using money, and we can both be armed with guns or other weapons that make unilateral attack not terribly likely to succeed. Even in US prisons, where individuals are more likely to be violent and prone to disagreement than the general public, market systems emerge with agreed-upon units of exchange. Originally the monetary unit was a cigarette until they were banned, and now it is canned mackerel. See Justin Scheck's "Mackerel Economics in Prison Leads to Appreciation for Oily Fillets", <http://online.wsj.com/article/SB122290720439096481.html>

pairings of particular perspectives rather than a single “neutral” shared perspective. Since different perspectives view outcomes differently, there is no reason to suspect that they will all identify the same outcome as the appropriate solution. This is easy to see in normal flea market cases: the same good is often sold at different prices even by the same vendor. In the case that I look to employ a bargaining model, this means that the particulars of the social contract depend heavily on the makeup of the particular society in which it is being enacted. More specifically, the social contract depends on the perspectives of the agents creating the contract. To return to the food example, a society of people who all keep kosher could very sensibly ban pork and shellfish so as to prevent accidental pollution of their food. But this may change if people with no eating restrictions are also in the population, as their categorization of those foods would now be included in the bargaining process. Another example might be the case of a feminist perspective in articulating what workplace equality might entail: rather than just making jobs available to both men and women, as might be suggested by a male perspective, government policies requiring maternity and paternity leave, and some subsidy on child care might be required to ensure equal access to employment. In later sections, this perspective-sensitivity will be developed at greater length.

Supporters of public reason and deliberational models of political agreement might want to challenge a bargaining position on the ground that bargaining does not promote a system of reciprocity and mutual understanding and respect, but is instead emblematic of a *modus vivendi*. On this way of thinking, a bargaining model,

presuming an agreement could be reached, would merely allow individuals to coexist, but would not do anything to promote developing stronger social ties. In this way, a bargaining-based social contract would be the shallowest form of social union: it would merely allow individuals to promote their own self-interest in a society of other rational maximizers. This would be a far cry from the ideal of public reason, in which public debate relying on commonly held political justifications would serve to draw people closer to each other and foster a greater sense of shared identity that would promote a system of social reciprocation. I want to challenge this objection on two fronts. First, that bargaining just *promotes* a *modus vivendi*. The bargain itself is independent of the depth of social ties. While it is true that bargains can be reached between people who will interact only once and have no particular interest in each other, bargains do not have any features that require individuals to have no social ties. It is true that a bargaining model would permit a *modus vivendi* arrangement, but this seems to be an advantage: a bargaining model might find success where none would be possible in a public reason model. A public reason model requires greater commonalities among individuals. And as I question the premise as to whether these commonalities often obtain, a bargaining model seems more appropriate.

There is the further question of whether a bargaining model does anything to promote stronger social ties. Public reason models have a great deal of traction here: they do encourage fostering a shared political identity, and this would seem to encourage reciprocity. I would contend that a bargaining model does much to

promote social ties, but in a manner that is quite different from the public reason model. In particular, a bargaining model encourages individuals to better understand the interests of their counterparts. Better negotiation outcomes can be had if all parties can better understand what the other side wants, so they can be in a better position to give it to them, and in turn be able to demand other concessions. In other words, bargaining encourages empathy.⁸⁶ What's more, this empathy is based not on some shared political viewpoint, but instead an attempt at understanding the perspective and values of one's counterpart. For example, cities that have historically been more trade-oriented, such as port cities, tend to be more tolerant than cities with little active trade.⁸⁷ Rather than suppress a "private sphere" version of oneself for a common "public sphere" identity, citizens can bring their full selves to bear on social concerns, and be engaged with on those grounds. This engagement legitimates each other's perspectives, as the bargaining process acknowledges them as a relevant set of considerations to be reckoned with. This validation has no counterpart in a public reason model, as only the shared political perspective is considered. In a bargaining model, alternative viewpoints are

⁸⁶ It should be noted that bargaining encourages empathy *in the long run*. If two people are randomly paired to bargain only once (as we might imagine in a flea market case), there is only so much that one party can do to come to understand the other. However, in cases of repeated bargains, one can get a better sense of what the other party responds to, and has an incentive to investigate this further. Taking on a cost to learn more about the other party may not make sense in a single-shot case, but it will often pay dividends in a repeated interaction. This argument has a long tradition, most notably in Hume and Smith. A detailed discussion of this view can be found in Hirschman's *The Passions and the Interests*. (Princeton University Press, 1997).

⁸⁷ A recent case study of this is found in Jha, Saumitra, "Trade, Institutions and Religious Tolerance: Evidence from India" (January 10, 2008). Stanford University Graduate School of Business Research Paper No. 2004 Available at SSRN: <http://ssrn.com/abstract=948734>.

embraced as relevant, even if they are not assured of carrying the day. This promotion of empathy and the procedural respect shown to alternative perspectives would do much to promote social cohesion, precisely because it serves to enfranchise those individuals who might otherwise have had their perspectives excluded from the political process. This is not unlike the *doux commerce* theory advocated by many of the early economic thinkers. Smith argued that market players do best when they are honest and empathetic. Self-interest drives agents to develop reputations for being fair and considerate, for if they are not, no one would want to do business with them.⁸⁸

Though the source of justification of a bargaining outcome is ultimately individual rather than communal, this form of justification assures us that individuals from every perspective are willing to endorse the outcome. This is not to say that this will assure everyone the best possible outcome, but as we will see in later sections, it will ensure that everyone will get the best outcome that can be gotten given the constraints that everyone places on each other.

A Conception of Rights

Citizens of both the United States and France enjoy religious freedom, which includes the separation of Church and State. These rights, however, are interpreted in a quite different manner in the two societies. In the United States, the state does

⁸⁸ Hirschman traces the history of the *doux commerce* thesis in his *The Passions and the Interests* (Princeton University Press, 1977). Adam Smith develops his account in *The Theory of Moral Sentiments*

much to protect the right of individuals' religious expression. When we interpret the rights of church-state separation, we see it in terms of "anti-establishment." The state cannot do anything to promote one religion over any other. But we allow for a market of religions – citizens are free to choose their religion and level of religiosity, including their outward expressions of their faith.

France, however, has a different approach to interpreting these rights. In part, this is due to the way that the French interpret the concept of the citizen. Namely, individuals are "French first" and every other cultural identity comes second. Individuals are asked to subsume every other aspect of themselves to their Frenchness, and in return, the state grants citizens the right of choosing whatever religion they want. This has had some success – in particular, Post-Napoleonic France was historically more welcoming of Jewish citizens than many other European countries so long as they saw themselves as primarily French. However, this interpretation of religious rights differs in an important way from the United States. Namely, it is not acceptable to have outward displays of religiosity, particularly in public institutions such as schools. This has in recent years become a significant issue with France's growing Muslim population. Though students can wear necklaces with religious symbols on them under their shirts, they are not allowed more outward expressions of faith, such as hijabs. Muslim girls have been expelled from school for wearing a hijab to class. Teachers have refused to teach so long as students wear them. The Muslim community claims that these restrictions are a violation of their rights, but the state responds that it is applying the same standard

to everyone. If anything, France is trying to cling to its “French first” interpretation of rights more aggressively now, as it faces increasing ethnic diversity. The state, and many in it, see this interpretation of rights as a way of promoting national unity. Many Muslims see this interpretation of rights as discriminatory.

What makes the French case particularly useful is that we can see how one right can affect many others. A right to religious expression for Muslim girls can be seen by other French citizens as a violation of gender equality. Having this violation of such an important right occur in school can have an undermining effect on other non-Muslim girls who might take in a conception of women as submissive or as objects of sexual desire that must be covered. This right of expression can be seen as *promotion* – that the state is now endorsing the practices of a particular religion. This can make others who are not Muslim uncomfortable or even feel pressure to conform with Muslim practices, which would violate their own religious autonomy. What is important to notice here is that our choices about interpretations of certain rights can impose external costs – those costs are limitations on how we might be able to interpret other rights. If we want to take a strong stance on rights of gender equality, we have to scale back our interpretation of a right to religious expression.

The goal of this section is to set the groundwork for understanding how we can apply a bargaining model to the social contract. Most basically, the bargain is about the rights and privileges that form the basic structure of society. These rights, as we can see, can come into conflict with each other, and we need to find a way to

come to a common agreement about which rights we ought to endorse over others in particular cases. But we need to understand what exactly a right is. As we can see from the example, not only can we interpret what appears to be the same right in two very different ways, but also each interpretive choice has many externalities that can affect how we interpret other rights. Thus, we are here concerned with the core freedoms and duties that citizens enjoy in society.⁸⁹

Let us consider what a right is. Rights are apportionments of freedoms and limitations on actions. They determine what individuals are allowed to do, how they are allowed to do it, and what constraints of action they can demand of each other. Importantly, rights also determine what actions cannot be constrained by others. It will here be useful to introduce a standard distinction between kinds of rights: there are negative rights and positive rights. Negative rights assure your freedom from *restraint*: the right to free speech says that I can say what I want without anyone being able to silence me, and the right to habeas corpus says that I can't be imprisoned without cause. Positive rights incur obligations on others to assure your well-being: The last two of FDR's Four Freedoms are an example of this. The freedom from want and the freedom from fear are obligations that we all have to each other. Positive freedoms require some measure of government or social intervention – a freedom from want implies a social minimum, for example. So

⁸⁹ Though I am only discussing rights allocation here, this discussion can easily be extended to an analysis of economic allocation. I do not develop this here as I am primarily concerned with questions of political power in this chapter. In the final chapter, material considerations will play a much larger role in my discussion. Very loosely, this chapter is addressing what Rawls saw to be the domain of the first principle of justice, and the next chapter addresses the domain of the second principle of justice.

while negative rights tend to be prohibitions of government intervention, positive rights outline government obligations towards its citizens.

The distinction between positive and negative rights is an extremely useful one, but it is not the only distinction that we might make. Importantly, we can also categorize rights based on whether they are rivalrous⁹⁰ or non-rivalrous. We can also categorize based on the orthogonal criteria of whether rights are excludable⁹¹ or not. Combining these two distinctions, we can determine whether rights can be understood as public goods⁹², or if they come short in some respect. Rights that are rivalrous imply that when one person “consumes” the right, it takes away resources from another. Many positive rights have this character. For example, guarantees of a social minimum, or a right to health care would be rivalrous goods, simply because they involve the consumption of other material (rivalrous) goods.⁹³ Other rights do

⁹⁰ Rivalrous goods are goods where consumption reduces availability of the good to others. Normal physical goods, like chairs or food are examples of rivalrous goods. If I eat a sandwich, you can't eat it, as it's now gone. Non-rivalrous goods do not have this character. When I breathe, I am not depriving you of any air in any meaningful sense. Likewise, if I swim in the ocean, I am still leaving plenty of ocean for you to swim in.

⁹¹ Excludable goods are goods that can be access-controlled. I can let some people into my home, and keep others out. However, non-excludable goods don't have this character: it would be impossible to limit the light from street lamps to only those people who drive American cars, for instance.

⁹² Public goods are those goods that are neither rivalrous nor excludable. This means that there is zero marginal cost to providing another person with the good (this condition is non-rivalry), and that if some people have access to the good, there is no way to restrict others from having the same access (this condition is non-excludability). A lighthouse is a classic example of a public good – it does not matter how many ships use its light, nor can the lighthouse prevent ships from seeing its light if some ships can see it. A normal manufactured good is the opposite of a public good – it is both rivalrous and excludable.

⁹³ One could try to make the argument that the cost incurred in providing social minima or health care is offset by the savings in other areas of government spending, such as emergency services or prison expenditure. This may very well be the case, but it would

not have this character, however. Free speech is nonrivalrous: my ability to speak my mind does not in any way inhibit your speech. In fact, if anything, my free speech can promote speech from you, whether you want to agree or disagree.

An interesting question is whether any rights are non-excludable. Having a right to clean air is non-excludable, in the sense that if you've supplied it to one group, others in the same area can't be denied the same right. A right to protection from foreign invasion is similar: one can't protect only some people in a country and not others from some foreign invasion. Either you have an army or you don't. However, most other rights do not seem to have this same character. Police protection could be selectively granted, for instance. Even a right to speech could be selectively granted: one class of individuals could have that right, while others could be denied it. This is true of rights to practice a given religion, right to assembly, or a right to health care or a social minimum. Non-excludable rights, then, are only those rights that are also public goods, in the economic sense of the term. So we have seen that while most rights are excludable, many of those excludable rights are non-rivalrous. This means that though these rights can be extended to many or few people with no direct tradeoffs, there are ways of preventing these rights from being extended widely.

Why might someone want to not extend rights to a group, especially when there are no direct costs associated with such an extension? There are three

have to be a perfect offset in every case for such an argument to work. Being that we are in the end dealing with material goods, which are always rivalrous, this argument is not going to be successful.

plausible reasons that I will consider. First, one might get satisfaction from the denial of rights to others. Second, one might have the paternalist inclination that it is best for the group in question to have a restricted set of rights. Third, one might find some aspect of these rights in conflict with some other set of rights that one takes to be more valuable. The first reason requires a particular psychological disposition: not only does one care about others, but one cares about them insofar as one can hurt them. This cannot be a preference of an egoist or an altruist. So while this preference applies only to a limited portion of the population, it is a possible preference. In particular, members of hate groups could be classified as having this kind of preference.⁹⁴ They oppress people simply because they get satisfaction from the oppression itself. Otherwise, there are no relevant benefits.

The case of paternalism is simply one in which the paternalist must assume that they know better than others about how to best help them. A restriction in rights stems from the belief that a certain class of people, left to their own devices, will make poor choices. This can have ugly consequences, as can be seen by some paternalist arguments in favor of racial subjugation.⁹⁵ However, in other instances, there can be cases of paternalism that have more justification, such as mandatory

⁹⁴ Though this classification is possible, it may be incorrect. Another plausible reason for being a member of a hate group is that one believes that once you grant rights to the group whom you are oppressing, these newly-granted rights will inhibit or diminish your own rights. Which would then move us to our other reason for not extending rights.

⁹⁵ A striking example of this is found in Thomas Carlyle's infamous essay "Occasional Discourse on the Negro Question" *Fraser's Magazine for Town and Country*, London, Vol. XL., (February 1849). Mill famously responded by challenging the paternalist approach to a denial of rights in his essay "The Negro Question" in 1850 in the same magazine.

seat belt laws and smoking bans, wherein minor rights restrictions come with a fairly sizable public benefit.

The more interesting case to consider is where rights are not extended fully because they come in conflict with other rights. This is the case that requires a more nuanced understanding of what rights themselves are. There are two cases to consider: restricting the number of people who share in access to a particular right, and restricting the scope of the right itself. To illustrate the first case, consider voting rights. In the United States, we have a significant history of restricting voting rights to only certain classes of people. Originally, only landowning white males could vote. While women and minorities are now granted voting rights, we still significantly restrict voting rights. This mainly happens in two populations: minors and felons. Individuals under 18 are not allowed to vote, no matter how politically active and informed they may be, simply because they are not yet old enough. This is true in all states. In some states, felons who are out of prison are not allowed to vote either. We restrict the votes of minors and felons in part because (rightly or wrongly) we believe that they cannot be trusted to put the public good first in their vote.⁹⁶ In this instance, rights are restricted even though they are non-rivalrous because the rest of society deems that the disenfranchised population has not earned those rights.

⁹⁶ In the case of felons, another likely source of justification is that their loss of suffrage is part of their punishment for committing serious crimes.

In Belgium, however, rights are restricted to particular groups not because the rest do not deserve them, but rather different cultural groups have interests in different sets of rights. In particular, cultural-linguistic rights are managed by Community rather than by the federal government. So the differential distribution of rights is meant to reflect underlying differences in the populations, not a statement about power differences or equal treatment. However, the Belgian case can be explained in terms of rivalry. If a right is such that there are costs incurred for providing it to the marginal individual, then it is reasonable to supply the right only to those who want it at least as much as the marginal cost. Providing the right to everyone only serves to reduce the number of other rivalrous rights that one could otherwise provide that people would value more. So, when one is considering rights that are rivalrous in nature, it is entirely consistent to have both differential distribution of rights and equal treatment of all citizens. But when rights are nonrivalrous, differential distribution of rights needs further explanation.

Let us now consider the second case of limiting rights, where we limit the rights themselves, rather than the scope of people who have them. To illustrate, let's consider the right to free speech, and the right to personal bodily integrity. We have free speech rights, which should entail that we can make whatever claims we may want. But if we shout "fire!" in a crowded room, we will cause people to trample over each other. Similarly, if we shout "Terrorist! He has a bomb!" on a plane, the person we point to will likely be tackled and beaten. In both of these circumstances, allowing for the fullest version of a right to free speech would

adversely affect the rights of others – in particular, their right to maintain bodily integrity. So we have a conflict between free speech and bodily integrity. This conflict raises an important issue about the nature of rights. We need not now decide that there is no right to free speech, or that of bodily integrity. Instead, we can choose to *reinterpret* one of those rights so as to eliminate the conflict.

Reinterpreting rights suggests that they can be broken up into component parts. So rather than having a single entity, “freedom of speech,” we can instead think of this right as being a bundle of many allowances. One such allowance might be the right to run a press, another might be to give a speech at a street corner, and another might be inducing panics in crowds. Though these activities are all related to each other, they are not mutual requirements for one another: it is perfectly consistent that one might be allowed to speak at a street corner, and not allowed to induce panics in crowds. Similarly, property rights can be absolute, or we can limit those rights in such a way that make taxation acceptable while still prohibiting stealing. So the concept of a right must shift away from being an absolute, and more towards a bundle of many affiliated allowances for action or guarantees that may be considered at least somewhat independently of each other.

This conception of rights as bundles of allowances and guarantees makes them firmly a social conception, as opposed to a metaphysical one. We agree to what of each bundle we take to be important, and what we can discard. This is evident in the differential allocation of rights across countries. Even Western liberal countries vary amongst each other in their rights allocations. If we viewed rights as

something other than an agreement amongst citizens, we would expect to have active debates between friendly countries so as to determine the “correct” set of rights, just as we do with scientific disputes. There is no French consensus on the atom that differs from the British or American one – there is just the scientific consensus.⁹⁷

What we will turn to next is how we can develop an account for making these kinds of decisions. It is easily seen, however, that this requires that we view rights as subject to revision, and more importantly, subject to tradeoffs. Though we might care deeply for the freedom of speech if we think about it in general, not many among us would support the idea that the right to induce panic in crowds is more important than bodily integrity. So, we are willing to modify our conception of the freedom of speech to eliminate such a right so as to get rid of the conflict. Thus, the conception of rights I develop here requires that rights be bundles of associated allowances and guarantees, and these are all subject to tradeoffs amongst other bundles, since we cannot have two rights that conflict with each other.⁹⁸

At this point, we have seen that a bargaining model has some advantages over a public reason model of a social contract. We have also outlined a conception

⁹⁷ This is not to say that evidence of debate shows that there is no underlying truth. Rather, I suggest that a lack of debate in the face of differences suggests that these differences are not caused by different attempts to uncover hidden truths, but instead a recognition that the differences represent different social agreements.

⁹⁸ Having conflicting rights is incoherent: if rights are to be respected or even just always followed, then if two rights require conflicting sets of actions, at least one will not be followed. But then rights would not be notably different than just sets of guidelines that we can choose to follow if we so wish. For rights to have the inviolable character that we think they have, they must be mutually consistent.

of rights that allows for categorization along the economic conceptions of rivalry and excludability. Employing these distinctions, we can see how rights can be traded off amongst each other – and in fact that they must be traded off so as to assure a consistent set of rights. With these concepts in mind, we can now turn to a full elaboration of a bargaining model of justice.

The Bargaining Model

As we have seen, the starting assumption of a public reason model is that there is some suitable perspective that everyone can adopt to reason as citizens. This is problematic precisely because individuals do not all share this perspective, and since different perspectives can have different associated categorizations, individuals might not even be able to accurately represent their own political interests in a non-native perspective. Our goal here is to correct this main deficiency of the public reason model by allowing individuals to engage each other with their own perspectives. In particular, in this section I will lay out a bargaining procedure that allows individuals to agree on a particular distribution of rights without having to agree on a common perspective from which to debate. To do this, I will rely on economic accounts of bargaining, though I will extend them to satisfy multiple perspectives.

The first step is to examine the consequences of perspectivism for an economic description of bargaining. In the previous section, we established that

rights can be treated as bundles of associated guarantees. In normal economic theory, we would first define states of the world by partitioning state space: we divide up the state of the world into many pieces, and each partition identifies a particular scenario. Bundles of rights are the set of actions that we have that can apply to particular partitions of the world. The scope of a given right can be thought of as how many partitions of the state space it applies to – when we are allowed to perform the given action. An assumption here is that everyone sees the same partitions of state space, and can see how a bundle of rights apply to the world.⁹⁹ When two bundles of rights affect overlapping portions of state space, we can have conflicts between the two bundles. That is, rights can come into conflict with each other when they apply to the same state of the world, and describe different, incommensurate sets of allowable actions.¹⁰⁰ In the case of hijabs in France, an account of religious freedom interferes with gender equality. So a right has two components: first, a list of the allowable actions, and second, the list of partitions of the state space that these actions are allowed in. In those cases where there are conflicts between sets of allowable actions, the agent can then have preferences over these rights in the particular partition. For example, one can rate free speech in

⁹⁹ This is slightly incorrect. In fact, the standard axioms require that if Alice and Bob are playing a game, Alice's partition set must be a subset of Bob's. This means that one agent may make finer distinctions than the other player, but he recognizes all the distinctions that she does. An example of this would be that Alice recognizes sunny days and cloudy days, but Bob recognizes sunny days, cloudy days and partially cloudy days. Bob can make finer distinctions since recognizes the difference between sunny days and partially cloudy days, and thus make more careful choices about his actions than Alice can, but he does not lack any categories that Alice has.

¹⁰⁰ In standard economic language, this would be one's strategy set. A strategy set defines the set of actions (also known as strategies) available to an agent.

crowded rooms as less important than bodily integrity. For partitions of the state space where there is some potential conflict, agents have to decide if they prefer right A to B, B to A, or are indifferent between A and B. Finally, one can have preferences over the partitions themselves: that is, one can determine which cases they deem to be most important, and what cases they care about the least. One may strongly care about free speech when it comes to the ability to have a printing press, but not care much when it comes to the ability to say whatever you want to your elementary school teacher in class.

Under normal economic theory, we have sufficient pieces to describe a bargain. Individuals have complete preference orderings over all the cases of rights conflict. Agents can then attempt to concede on some cases that they care less about to win concessions on those cases that they care the most about. Each agent's aim is to find the bargaining outcome that maximally increases her own utility with the constraint that her bargaining partner will agree to the arrangement. At the bare minimum, this means that neither party can come out of the agreement worse off than how they started. But this leaves open a huge range of potential bargaining outcomes. If we take seriously the requirement that both parties seek to maximize their returns subject to being able to get agreement from the other party, we are in fact left with a single solution – the Nash Bargaining Solution. The Nash solution is the bargaining outcome where the product of the utility gains is maximized. Maximizing the product rather than, say, the sum, helps to ensure fairness: each agent is gaining at as close to the same rate as the other agents as is possible. While

this might seem like a convenient mathematical trick, it represents a deep thought: it captures the willingness of agents to accept a particular bargain. The Nash solution is the unique efficient¹⁰¹ solution that satisfies some basic axioms that support this condition of maximization subject to agreement.¹⁰² So each agent is maximizing their utility as much as they can, subject to the constraint that others have to agree to the bargain as well. If any agent proposes a bargain that increases their own utility, at least one other agent would be losing relative to the Nash solution, and thus reject the bargain. The solution treats all agents equally, and so serves as a focal point for the bargain.¹⁰³

This bargaining approach has some nice features. In particular, the Nash Solution provides the best deal that any individual can get subject to being constrained by others. But it is done without having to do any interpersonal comparisons of utility. Each agent can have utility on whatever scale they want, and at no point do they have to try to interpret someone else's utility. They just have to worry about someone accepting the outcome. As we saw in the earlier discussion of mutual benefit models, this is taken care of in virtue of the interaction of rational

¹⁰¹ In this context, "efficient" means that there are no pareto moves left to make. That is, any gain must be associated with a loss by someone else.

¹⁰² The first axiom is invariance to change in how one represents their utility (invariance to affine transformations), which means that it does not matter how you describe your utility function, just what it is. The second axiom is that it is pareto optimal. The third axiom is that it is independent of irrelevant alternatives. This means that if you prefer Apples to Oranges when those are all that's available, you won't all of the sudden prefer Oranges to Apples when you can also choose Bananas. The fourth axiom is symmetry.

¹⁰³ For a much more extensive justification of the Nash Bargaining Solution, see chapter 7 of H. Peyton Young's, *Equity: In Theory and Practice*, Princeton University Press, Princeton, NJ 1995

agents itself. But we still run into problems. In particular, all the agents have to share the same partition sets. In effect, this is the same as saying that they must have the same categorization, which seriously constrains their abilities to have different perspectives. So we must find a way to modify this bargaining procedure to account for multiple perspectives.

There are two changes that become necessary. First, we must find a way to deal with the fact that different individuals can have different categorizations. Second, we must weaken our claims about the completeness of preference orderings.

If two agents have different categorizations, this means that they will not divide state space up in the same way. This is simply because they categorize the world differently, and so they have different conceptions of salience. As we saw in the previous chapter, they are seeing the same thing, but interpret it differently. As there is no neutral representation of the state space, we cannot take one partition as better than another, and so we cannot just pick one and require everyone to use it. This much is clear. But we also cannot just take the union of everyone's partition sets and expect everyone to share this larger partition set. This is because there is no assurance that people can *see* the distinctions that are in other perspectives if they do not themselves hold them. So just as we do not have the same accounts of value (and hence utility functions), we do not have the same partition sets. But even though we cannot expect that individuals can adopt this union of all partition sets,

we can describe the overall bargain by making reference to this union set, so long as we do not assume that the agents in the bargain themselves are aware of it.

Though we can use the union of partition sets in our model of the bargain, it may be useful to consider an alternative representation of the same idea. In particular, we can think of the categorizations of a perspective as dimensions of some high-dimensional object. Each dimension is a criterion by which one can make distinctions about states of the world. So, one dimension might be weather, another might be day of the week, a third might be how many people are present, and so on. We can aggregate all agents' perspectives by simply adding dimensions for every way of categorizing that these perspectives have. So, we have a single high-dimensional space that each individual can "see," but only by taking a lower-dimensional projection, as represented by their particular perspective. This high-dimensional space is not a metaphysically real entity – it is just the agglomeration of all the agents' perspectives. This space allows us to see how everyone is still in the end talking about something shared, even if they do not agree on how to describe it.

What is important to notice is that, as we each can see only a projection of this larger space, when Alice walks into the store and sees just food; Bob can walk in and see meats, grains, fruits, vegetables and dairy; and Carol can see canned food, frozen food and fresh food. They are all still looking at the same objects, and they are each making legitimate categorizations – they just don't recognize each other's categories. This captures something important about political discourse. Disagreement does not stem only from differences in preferences. It can also come

from differences in how people interpret the world. So it is essential to capture this source of disagreement if we are to create an accurate model of political negotiation.

Thus, the bargain itself takes place over disagreements in preferences about rights allocation across this space, even though each agent can see only their particular categorization of the disagreement. Agents do not see themselves as engaging in the same bargain – they would describe the bargain rather differently from each other, based on their own perspective. Further, because these categorizations are different, preferences are going to be over different things as well. Though each agent can have a complete preference ordering of rights across states of the world on their own perspective, they must have only a partial order¹⁰⁴ across the whole high-dimensional state of the world, since it includes many distinctions that the agent does not recognize. It is impossible to state preferences about objects that one does not recognize as existing. One consequence of this is that if Alice does not recognize some distinction that Bob makes, Alice is going to be more than happy to let Bob have what he prefers in that case, since Alice literally can't see what Bob wants.¹⁰⁵

With these two modifications, we now have a bargaining model that allows for multiple perspectives. Each agent has a partial order over rights in the union of

¹⁰⁴ A partial order is one where an agent cannot say whether they prefer A to B, or B to A, or are indifferent between the two of them, for some A and B in their choice set.

¹⁰⁵ This is true in those cases where Bob's preference does not affect any outcomes that Alice can recognize.

all partition sets, which defines their utility function. They also each have a partition set that is a strict subset of the union set. Even if agents do not agree about what they are bargaining about, they can carry out a bargain.

What enables these agents to carry out a bargain is that there is always a price to allocating a given right. Even if different agents interpret rights differently, a price can be established, since everyone has a defined utility function. In particular, this analysis of a social contract in terms of a bargain allows the society engaging in the bargain to discover exactly how much certain rights are worth to people who hold a given perspective. To discuss this important idea, I would like to examine two issues. First, how we find a price, and second what this price mechanism means for rights allocation.

That rights have a price is a perhaps uncomfortable claim to make. Rawls and other contract theorists deny citizens the right to sell their rights and freedoms for other benefits. One is not allowed to make any tradeoffs at all. I too would like to restrict the sale of rights. However, that does not mean that there cannot be tradeoffs *between* rights. As mentioned earlier in the description of the model, I assume that agents place value on particular interpretations of their rights. I further assume that the utility one gains from rights is fungible, so that tradeoffs between rights are possible. This is just to say that I might be willing to accept a limitation on my right to bear arms if I can have increased rights to free speech.

With just these two assumptions, we can establish a price mechanism for rights. The price of a right is just its social cost. The social cost of a right is defined as the utility loss incurred by the conflicts with other rights bundles. The social cost of a strong right to gender equality in France is then the utility loss suffered from the perspective of those who want girls to wear hijabs to school. Some rights may have no social cost, and are thus free. A very basic version of free speech is likely to impose no social costs, as it does not impinge on any rights. It is even possible that it could have a negative cost if it helped enable other rights. But the right to hate speech can have high social costs, as it impinges on various rights to equality.

This account of a price for rights requires nothing more than a normal market discovery mechanism. The agents bargaining will accept or object to various rights bundles based on their personal utility assessment, and as such, they will reveal their valuations of these rights over the course of the bargain. The market price of a given right is just that price at which the social costs of the right are covered. It is only at this point that other agents in the bargain will accept the imposition of that right, since they are otherwise losing more utility than they are gaining. The price of a right is the point of indifference between having and not having the right. Bundles of rights can then be used to have rights that offset each other. A weaker right to property might be bundled with a stronger right to health care, for example. So rather than spend money or some physical resource to “pay”

for a given right, we can think about paying for a right with the allocation or removal of other rights in a particular bundle of rights.¹⁰⁶

This price mechanism, unlike the public reason model of social contracts, allows agents with different perspectives to engage with each other in a way that doesn't privilege any given perspective. Since price is just the social cost of a right, this can be evaluated from the perspective of any given agent – if the agent is losing more than she gains, then she pulls out of the agreement. But she never has to agree on what it is that she is bargaining about. This is not possible under a public reason model. In addition, as opposed to just talking about “the balance of reasons” supporting a particular right, we can determine how much a particular interpretation of a right is valued by looking at what a society might be willing to give up to get it.

Additionally, the price mechanism embedded in a bargaining model allows us to treat different kinds of rights differently. In particular, we can identify those rights that are public goods and those that are excludable and rivalrous. Rights that have the character of public goods, or are even just non-rivalrous are those to be allocated to everyone in society at once. Public goods by definition can only be provided to everyone, and non-rivalrous goods cost the same to provide to

¹⁰⁶ This particular form of payment is known as “logrolling” in the political science literature. Logrolling occurs when legislators agree on voting patterns across a series of bills. An example of this is when Alice agrees to support Bob's budget bill if Bob will then support Alice's climate change bill. Of course, logrolling can involve many more than two people and two bills. The idea is that legislators need to get enough votes on their bills, so they agree to exchange votes that they would not normally make so as to get what they want at a price they are willing to pay.

everyone, so there is no reason to not allocate them fully. But rivalrous goods are also excludable goods, and those rights that have this character can be more selectively allocated. So, only those individuals who value rivalrous rights will choose to pay for them. This will result in different people having different rights, but as discussed before, this is not an indication of discrimination or inequality. Instead it is a recognition of diversity. What assures this claim is the price mechanism itself. Those that want “additional” rights must pay for the social cost of these rights. While rights that are public goods must be supported by everyone, rivalrous rights can be exclusive, and so we can allow for those agents who want them to have them, and those who do not to not have to pay for them.

Conclusion

Building on the account of perspectives laid out in the previous chapter, we can see that a public reason account of the social contract is unable to accommodate diversity in a way that is neutral across perspectives. As a social contract is most importantly an embedding of a set of rights, we examine what we mean by rights. Once we adopt a conception of rights that considers rights to be bundles of associated interpretations of the right, and a classification of rights based on rivalry and excludability, we can define a bargaining procedure that can fully accommodate perspectival diversity. The resulting social contract is a Nash Bargaining Solution, which ensures that it is the outcome that ensures that everyone received as close as

they could get to the set of rights that they most wanted, given that everyone else was out to do the same thing. This kind of justification is different in kind than the justification found in public reason models, as it is individual rather than social.

In the past two chapters, we have looked at a two-stage process for determining moral principles and establishing a set of rights. These two stages are the foundation for an iterated search process for social improvement. But we need one final stage to provide the necessary information to feed back into the beginning of the process. Namely, we need to examine how we take our given moral principles and rights allocations and then apply them in actual social cooperation. The next chapter is thus concerned with issues of social cooperation and distributive justice in diverse societies. I take this next stage to be where our thought experiments become actual experiments in living. It is from their experiences in actually trying to follow what they have determined in the previous chapters, along with the distributive issues that they face that agents can determine whether they want to re-evaluate the social contract, or continue to abide by it.

Chapter 4: Experiments in Distributive Justice

Introduction

In the previous chapter, we examined a mechanism for achieving agreement about the allocation of rights in society. This bargaining mechanism is designed to work even if different parties disagree on the terms of the bargain itself. The allocation of rights that results from this bargaining mechanism should be thought of as an agreed-upon and large-scale experiment in living – it is something that will be tried, but is expected to be subject to refinement or even larger revision. The particular allocation of rights may turn out to be close to ideal for everyone, or some aspect of it may turn out to have unintended consequences that are unacceptable to at least some part of the population. If the system of rights is found to be wanting, then the different groups in society can collectively determine whether it is best to return to the stages of determining a social contract found in chapters 2 and 3, or split the society into multiple independent units. As we are epistemically limited beings, the only way that the suitability of the allocation of rights to the particular individuals in society can be determined is by means of actually implementing it and examining the consequences.

What I aim to do in this chapter is demonstrate how this procedure of evaluation might look by examining how distributive questions can be negotiated once a scheme of rights has been established. As property rights are social

constructions, the range of possible property regimes is quite broad. However, given that the structure of the bargain is such that there is full information available to all agents about their current status, and that the bargain supposes that all rights are priced based on the harms imposed on others, we can assume that the actual choices that could be made about a scheme of property rights is more limited than the set of all logical possibilities. In particular, more “extreme” versions of property rights, by which I mean complete collective ownership, or absolute private ownership, would not be selected. This is for two reasons. First, it is very unlikely for a population to choose either scheme of rights. Second, it would be extremely difficult for a society to shift from some other more moderate scheme to an extreme system of property rights. To see why this is the case, let us consider the bargaining implications of such rights regimes.

A Lockean or Nozickean scheme of “absolute” property rights¹⁰⁷ would likely be favored by the richest in society, as it would do the most to assure the protection of their own wealth. However, such a scheme of rights would impose a great deal of harm upon the poorest in society¹⁰⁸, and at least some harm on those in the middle, because of the loss of the state’s ability to fund social programs. So the only way that this property scheme could be awarded is if the rich sacrificed a great deal in other ways – either by giving in to allocating rights to others that they would

¹⁰⁷ By “absolute” I mean that the state does not have rights to tax or expropriate property. Under such a property rights regime, market outcomes are final.

¹⁰⁸ The poor are harmed in this scheme because the state has no wealth to redistribute toward them. They are left with market outcomes, and the state is unable to guide additional resources to provide a safety net.

otherwise reject, or giving up on rights that they themselves would otherwise advocate for.¹⁰⁹ By the same token, a strong version of collective ownership of property would also be very difficult to achieve. It is conceivable that the poor could view this scheme as preferable to many of the other possible rights allocations, and thus have some interest in pushing for it. However, the rich would lose so much that the only way they would agree to the bargaining outcome is if they earned concessions on a wide variety of other rights. As such, the structure of the bargain would generally support only more moderate schemes of property rights, which would involve (limited) private property.¹¹⁰ These schemes of property rights would involve fewer harms imposed on a given class of individuals, and so require fewer concessions on some other aspect of the overall bargain.

More important than the particular scheme of property, however, is that for the purposes of this discussion, we can take for granted that there exists a current distribution of property rights in the population. The procedure outlined in the previous chapter results in a full system of rights, and the discussion in this chapter

¹⁰⁹ For example, the rich could lose their ability to gain political office, or lose some speech rights, like being able to donate money to campaigns or purchase political advertising. So while their property rights would be increased, the value of those rights would decrease in that their economic power could not be converted into political power.

¹¹⁰ It is possible, of course, that in particular societies, we could find instances of wide agreement on more extreme versions of property rights. For example, a society made up of Ayn Rand adherents would likely want absolute property rights, regardless of current wealth and income. I submit that this would be a very uncommon phenomenon, if we assume that states have even moderately sized populations. However, on those occasions when the particular makeup of the population would support these extreme versions of property rights, the remaining analysis is unaffected.

assumes that system of rights, and examines the constraints on a system of distributive justice in a diverse society.

Two Competing Images of Social Bonds

I would like to open with two images. First, we have the Original Position – easily one of the most compelling thought experiments in political philosophy. In this thought experiment, Rawls asks us to imagine ourselves stepping behind the thick “veil of ignorance,” where we lose knowledge of not only current facts about the society we live in, but who we are in this society – our upbringing, our education, our race, sex, gender, our interests, our social and economic status, and even our conception of the good. We are then asked to imagine how we would pick the rules governing the basic structure of society. Rawls argues that, given our massive ignorance, our rationality would compel us to choose that basic structure that would maximize the position of the least well-off members in society, since we would be right to fear that we would step out of the veil of ignorance and find ourselves stuck in an intolerable life. The Original Position is thus used to justify the choice of the difference principle as being an essential part of Rawls’ second principle of justice.

For this thought experiment to work, Rawls must make the further claim that when we reflect on our actual position in society, and on both those whose position is worse than ours and those whose position is better than ours, we recognize that

much of this difference is due to factors that are morally irrelevant. In particular, we do not do anything to deserve our parents, or how we are brought up, what resources and opportunities were made available to us, our natural talents. That we become better or worse off in our lives is largely due to the “natural lottery” – all those things over which we have no control, but nonetheless shape our lives. Alongside this claim, Rawls claims that we are free and equal moral agents. That is, despite the difference of where we have ended up on the scale of socio-economic status, each of us has equal rights and claims over the primary goods that enable us to lead the kinds of lives that we value. Critically, not only must we in fact *be* free and equal moral agents, but also we must *recognize* that we all are free and equal. Thus, a consequence of the thought experiment is that we must treat each other as equally deserving of the resources needed to live the lives that we find valuable.

The second image I would like to consider is Reagan’s condemnation of Cadillac-driving “Welfare Queens” as an argument for why the United States ought to drastically cut back on its social safety net. Left implicit in this image, and many others like it offered by conservative politicians, is that the so-called welfare queens are minorities.

Reagan was very successful with this rhetoric, and this style of argument can be seen throughout many policy discussions, particularly when there is an identifiable Other that can be attacked.¹¹¹ In current political discussions,

¹¹¹ This rhetorical strategy is often more effective if the social safety net is not universal, but rather targeted at the poor. If the comparatively well off only lose money on welfare,

immigrants bring diseases, are more criminal, are lazier, or are even potential terrorists. In European countries, which have more generous social safety nets, welfarist policies are being trimmed as immigration increases. The argument for cutting them is usually that immigrants take advantage of these programs in a way that majority culture never would. In this imagery, the Other has a series of negative qualities for which they are morally blameworthy. This can not only serve to remove our moral obligation to insure that they have adequate resources to conduct a life that they find valuable, but it can serve to instigate a desire to view resource allocation as a zero-sum game in which whatever they receive is a loss to the rest of us. We then find ourselves with the view that we must protect ourselves from this parasitic group, and work to deny them any resources that we might otherwise have been able to allocate to ourselves.

These two images conflict with each other in important ways. The psychological attitudes on display towards out-groups that we see in political discourse push adoption of policies that make the least well off in society even worse off. What's more, these policies often make the rest of society worse off as well, as I will demonstrate later. But our most compelling philosophical model of how to think about arranging the basic structure of society leaves these psychological attitudes out, precisely because they are deemed morally irrelevant. The real political actors are unmoved by the Rawlsian argument precisely because they do not see themselves as morally equal to members of the groups that they

and see no personal benefits, it is easier to be convinced that those who do receive welfare benefits are different in a morally relevant way.

disparage. In part, this is because “morally equal” tends to mean something like “is similar to me” in common practical usage.¹¹² As we can identify salient differences between ourselves and members of other groups, we can quickly decide that they are not our moral equals.¹¹³ This is only made worse when we consider regions of the world with even more inter-group conflict, such as the Middle East or Africa. But even if we limit ourselves to reasonably well-ordered societies, such as those found in the West, we find these basic conflicts either preventing or eliminating the realization of basic structures of society that are designed to promote a more egalitarian distribution of wealth and income.

What I want to argue, then, is that the question of distributive justice cannot be considered in isolation from the question of tolerance, particularly when we are faced with a pluralistic society. We cannot begin to try social experimentation in good faith if we cannot assure ourselves that we are aiming for improving outcomes for everyone. This becomes increasingly difficult if minorities do not feel as though their interests are being fairly represented. An important premise in the original position is that we can in fact see ourselves and others without our group identities,

¹¹² As has been argued in earlier chapters, there is a great deal of psychological literature supporting the claim that people treat members of their own in-group preferentially, even if the “in-group” is entirely arbitrary. An example of this literature can be found in J Sidanius, F Pratto, M Mitchell, “In-Group Identification, Social Dominance Orientation, and Differential Intergroup Social Allocation.” *Journal of Social Psychology*, 134(2), 151-167, 1994

¹¹³ A non-standard source of evidence for this claim can be seen in political cartoons. Representatives of other groups are often depicted as sub-human. This often takes the form of some simian features being introduced, or simply wildly exaggerated features that are stereotypically associated with that group. A collection of such political cartoons from Britain during the Irish Potato Famine can be seen here:
http://www.nde.state.ne.us/SS/irish/unit_2.html

and instead see ourselves as bare moral agents. What I want to suggest is that, at least initially, plural societies are not going to have members that are willing or able to do this. This is because we may not see an Other as a Person, in the sense that they are entitled to the same moral considerations as members of our own in-group. In these circumstances, we cannot expect an original position style argument to even get off the ground. Because of this, it becomes necessary to design the basic structure of society in a way that aims to foster the attitude that all should be treated as moral equals, rather than design institutions that already assume that this attitude exists. This bears importantly on the stability of a particular choice for the basic structure of society: if we need to create the bonds that tie diverse individuals together, we need to insure that we have done nothing that unduly increases the strains of commitment. In making this argument, I hope to make clear that toleration itself is an insufficient, and if anything, dangerously limiting, concept in liberal thought. What we need in place of toleration is a demonstration that diversity promotes an individual's own ability to conduct her life as she wants.

To proceed, I will first argue for a more limited concept of equality to ground the rest of the chapter. From there, I will demonstrate that economic matters within a state tend to have a positive-sum nature, and from there, show that a more diverse population is better at creating wealth than a homogeneous one. With these production-oriented questions out of the way, we can turn to the question of what the constraints are on making distributional decisions in a diverse society. In particular, I will argue that distribution is tightly aligned with contributions in

allocation, but that this alignment is different than traditionally thought. These distributional constraints limit the options for notions of distributive justice, but do not themselves present a model of distributive justice. Instead, they provide the necessary conditions for social experimentation that does not undermine social stability. Finally, I will show that while on the allocation side, more diversity is always better, the distributional requirements will place limits on how diverse a society can become before it gets unstable.

Minimal Equality

We have already seen that different perspectives are going to provide different accounts of equality, but there we were considering equality in terms of distributions of goods. Here we must first address what we mean by equality of persons. Rawls provides a thick account of equality when he uses the term “free and equal persons” – in particular, we are all equally valuable moral agents who treat each other as such. While we might have special relationships with some people and not others, we still view people equivalently qua persons. This is true whether or not they share the same group identifications or other individual traits.

This claim is normative, rather than descriptive. And insofar as it is a regulative ideal, I find it compelling. I would prefer to live in a society in which we treat each other as free and equal citizens who are engaged in a system of social cooperation. However, I disagree that we can work from this regulative ideal in our

design of the basic structure of society, since it appears to be abundantly clear that descriptively, the claim is false. As the welfare queen example was meant to demonstrate, we systematically treat members of other groups worse than we treat members of our own group. This is either inconsistent with a belief in moral equality, or a demonstration that the concept does not get us much.¹¹⁴ What is worse, our previous discussion of equality demonstrated that there are many different ways to define “equal,” many of which are actively employed to the exclusion of others by some subset of the population. This suggests that we need to instead find some more minimal account of the equality of persons that can be utilized across the spectrum of perspectives.

Our first step will be to abandon attempts to find an account of equality that is developed in the abstract. As we have seen, these kinds of accounts do not connect well to our actual concerns about assuring pro-social behavior. Instead, we will turn to a more empirical account, where we can directly engage in the question of whether individuals have an effective voice in political discourse.

¹¹⁴ The Spanish Inquisition, for example, could be interpreted as believing in moral equality in that they wanted to assure everyone’s entry into heaven. Which, they believed, required a great deal of torture and suffering on earth. Moral equality is also consistent with arguments for reducing the claim’s scope. Slave owners could very well have held that all persons are moral equals, but their slaves didn’t count as persons. Finally, we might find that our strong idea of moral equality leads us to conclude that if a minority group is systematically behind socio-economically, it must be due to personal traits, like laziness or lack of intelligence. So then, given that we are all equal, those individuals do not deserve assistance because they have chosen their position in life. Moral equality and feelings of superiority are not mutually inconsistent.

To find a concept of minimal equality, we will start with Hobbes. Hobbes argued that all men were equal because there was no man who was so strong that no one could kill him, and there was no man who was so intelligent that no one could outwit him.¹¹⁵ Notice that this definition of equality is descriptive, and not normative. It just is the case that no person is so powerful so as to be immune from attacks from others, whether they are violent or merely clever. Hobbes does not argue that this *should be* the case; he just argues that it *is* the case. I think Hobbes is correct in this approach.

What is so appealing about Hobbes's approach is that he recognizes that humans have a basic need for sociality. In order to protect themselves, they must join society. Otherwise, they will never be able to assure their own safety – there will always be some coalition that can destroy them. While a common reading of this is that it is just one of many examples of Hobbes's dim view of humanity, I find it to be a key element in what we want out of an account of equality of persons. I take Hobbes to have adequately argued for this account for self-defense reasons, but I would now like to extend the account into an area more relevant to our current inquiry. In particular, there isn't an individual who can be economically self-sufficient. Even if you are the most talented person in the world at each and every human endeavor, the limits on your own time impede what you can achieve. In particular, if you were a veritable Robinson Crusoe, and never engaged in any economic activity with anyone, and truly provided everything for yourself, you

¹¹⁵ In particular, no one is so powerful that they can resist any *coalition* of others.

would almost assuredly be very poor or very dead, depending on your circumstances. Having to start from scratch, with no tools, no access to outside learning, and no food reserves would be nearly, if not completely, impossible.

So this sets us up for the claim that we may have some interdependence, but we are not yet at a claim of equality. Here we can very naturally extend Hobbes to our more economic case. In particular, no individual is going to be more economically productive than any other competing economic coalition.¹¹⁶ Bill Gates might be very rich, but if you add up the two next richest people's wealth, he is overwhelmed. More importantly, however, is that his wealth was not of just his creation—to get it, he had to create a company with tens of thousands of employees, build up an entire industry around his company's products, and sell them to a very large piece of the world's population. Though Gates might have had a key ingredient in the creation of all of this wealth, he could not have done it on his own.¹¹⁷

What we find is that every individual is equal insofar as her economic bargaining power can be equaled by some competing coalition. It is important to recognize two features of this. One, more productive agents create more wealth,

¹¹⁶ This claim is restricted to be about individuals. It is possible that one could form an economically dominant coalition of individuals, but this is a separate issue from whether there is an individual who is singularly economically dominant.

¹¹⁷ The example of Bill Gates is complicated by the fact that his wealth is largely derived from the peculiarities of our legal system's treatment of intellectual property. While I do not want to make claims about what intellectual property rights might be chosen in as a result of the bargaining process outlined in chapter 3 (other than perhaps claim that it would be unlikely to match the current system in the US), I do want to suggest that what is relevant in this example is just that Gates' incredible wealth is not something solely of his own making, but is in fact a result of a great deal of social cooperation.

and as such have a superior bargaining position relative to less productive agents. But two, these more productive agents have a lot more to lose relative to the median agent if they are unable to find a bargaining partner, so relative to the set of possible partners acting as a coalition, they are in a weaker position, especially in the long run.¹¹⁸ A similar claim can be made about current wealth endowments. One can wield one's existing wealth as a bargaining advantage only up to the point when the other side decides that they would be better off if they just took it away.¹¹⁹ While wealth may be used in a one-shot game as a significant bargaining advantage, as it shifts the disagreement point in favor of the wealthier player, in long-run games initial wealth endowments are overwhelmed by the potential future earnings of both parties, and so the "shadow of the future" becomes more important than the disagreement point of the first round of bargaining.

Any discussion of comparative bargaining power must take place in the context of what kind of bargain we envision. As is often the case in strategic situations, one-shot games have different considerations than long-run, indefinitely

¹¹⁸ This basic bargaining insight explains why trade unions are good for employees. As an individual worker who is easily replaced by another worker, the individual has very little bargaining power. But the class of workers has a great deal of bargaining power, as without any of them, the companies in which they work cannot function. This is why union solidarity is so important: the union's power is reliant on maintaining a complete coalition of workers in order to match the power of the employer. Absent a complete coalition, scabs can undercut wages, and the union is undermined.

¹¹⁹ Seizing someone's assets ought to be no more difficult than killing a person, and as we've seen already, no person is immune from some coalition attacking her. So we can assume then that some coalition is capable of taking away someone's money. In less violent terms, this happens regularly when democracies decide that inequality has become too extreme, and create a more progressive tax structure. In essence, the wealthy can have short-term bargaining advantages, but these advantages can be corrected for.

repeated games. In particular, in one-shot games, it is much easier to use perceived bargaining advantages to one's benefit. However, in repeated interactions, this becomes more difficult, because the nature of repeated interactions allows for both parties to use defection¹²⁰ from cooperation as a means of "training" the other party over time to behave in a way that is more beneficial to the defector. As living in society is clearly a long-run, repeated interaction of many parties, we should understand ourselves to be thinking about bargains in the repeated context. Because of this, we can safely focus on the long-run bargaining power of the agents involved, rather than power in the one-shot case.

By defining equality in terms of bargaining power, we can now also see how this account of equality can accommodate multiple perspectives. We have left the concept of equality as an extremely minimal concept, and as such, have not included anything that would conflict with a thicker account that a perspective might provide. Instead, the concept of equality simply describes the relations of individuals to each other in the context of decisions about distributing resources. Not unlike in the previous chapter's discussion of a bargain of rights, here we find that bargains over the product of social cooperation can be treated in the same way. Individuals may

¹²⁰ When I use the term "defection," I refer to choosing to not bargain, or holding out for an outcome that the other would not accept. Importantly, I am also lumping together actual defection with the threat of defection. Often the threat of defection (or retaliation for the other party's possible defection) is sufficient to enable cooperative behavior. The classic example of this is the strategy of Mutually Assured Destruction. Neither party intends to use their weaponry, but both sides must have sufficient weapons to completely destroy the other. With this explicit threat of complete destruction, neither party will engage in even small-scale attacks on the other, for fear of massive retaliation.

differentially value goods and services, or even currency itself, but that does not affect how we account for their relative bargaining power.

Minimal equality, then, can be understood as a descriptive claim about the status of agents in a long run bargain. No individual can opt out of the bargain, and no individual can dominate the bargain. As such, every individual is faced with a similar set of constraints, and can be forced to respect those constraints by the other agents in the bargain.

One may be suspicious of this claim, and simply assume that the social bargain is going to be made up of coalitions primarily organized around economic class. One might further think that the capitalist class would be in a position to crush the worker class, and thus we would find ourselves with no working concept of equality. I think that such suspicions are misplaced. First, as we have seen in Chapter 1, there are no “natural” groups – it may well be the case that we have some coalitions that form around economic interests, but just as easily, coalitions can form around regional interests, religious interests, shared political goals, or any number of other focal points. But let’s assume that economic interests will dominate group formation, and we find ourselves with a coalition of capitalists and a coalition of workers. In a one-shot interaction, absolutely it would be the case that the capitalists would be able to force an extremely unfair bargain onto the workers. But as we have seen, in a long run bargaining situation, this does not appear to be likely. Unions are the natural response to excessive corporate power, and in the long run, the worker class can use its ability to defect to deny the capitalists the use

of their productive capacity. This defection can drive the capitalists to accept a more equal bargain. So, even in this worst-case scenario, where all of the economic capital is in one coalition, the long run dynamics of the bargain create opportunities for defection to be used to shift bargaining power towards the weaker party.

If one is coming from a Rawlsian, Kantian, or Rousseauvian context, this account of equality might come across as hollow. The account of equality of persons here makes no claim about their moral value, and instead substitutes in economic arguments about their bargaining power. But I want to reiterate that the reason for this move is that, even if I grant them that we want society to be made up of people who recognize that all are free and equal persons engaged in a system of social cooperation, we need to find a means to lead people there from here. So if our institutions assume that we've already achieved this, and further assume that people see themselves as equals in this thicker sense, then we will not be able to deliver on such a society, particularly in a plural environment. With this very strong assumption, our institutional design might be insufficiently robust to deal with individuals who do not believe in this thick notion of equality. This would decrease social stability, which could further undermine a fostering of the thick sense of equality. When we know already that people do not view themselves this way, we should not assume that they do, even if it seems like an ideal at which we ought to aim. However, I do want to contend that this more minimal account of equality is fully compatible with achieving those kinds of ends. In particular, I will argue in the next sections of this chapter that individuals who hold a view of minimal equality

will rationally learn to become more accepting of others and more willing to see agents who are different from them as valuable members of society. As we will see, however, there are limits to this approach in how accepting we can reliably become.

Creating Prosperity

The first challenge we face in understanding distributive justice in a plural society is determining what kind of situation we face. In real-world political discourse, people often understand the national (and world) economy as a zero-sum game, like poker. In these kinds of social interactions, every gain made by one person must be accompanied by losses by someone else. In a parlor game like poker, this is very clear. There is a fixed amount of money at stake, and so if I win a hand, at least one other person must have lost. Wealth only changes hands; it is never created or destroyed. As most of us have played poker or some similar game, this is often times our model for larger economic phenomena. Thankfully, it is the wrong model. I will first present some reasons for thinking it might be a compelling model, and then show why it is the wrong one.

The basic argument for the economy as one large poker game is that we do have, to a first approximation, a fixed pool of money. Central banks are not in the habit of printing money, and individuals cannot make their own money. So, then, if there is a fixed amount of money, then it must be the case that if I make more money, someone else is losing money. Further, we often hear about how

immigrants have “taken” the jobs that used to belong to natural-born citizens, or more prominently lately, that jobs have been moved to another country. On these descriptions, jobs are taken from one set of individuals and given to another set. So we can neatly point to the winners and losers, and see how the money changed hands, not unlike a poker game.

There is a problem with this description, however. Namely, that wealth regularly expands (and occasionally contracts). While central banks do not print money, the money supply does expand and contract based on how banks lend money. Even more important, though, is that money and wealth are not the same thing. We can understand money just to be the naïve notion of what one keeps in a bank account. But wealth is better understood as the range of goods and services one has access to. To see why this is an important distinction, let’s imagine a scenario in which one could have less money but more wealth.

Imagine immigrants come into a country, and exert downward wage pressure by increasing the labor supply, and having a willingness to work at substantially lower pay. Some of the previously employed workers will lose their jobs, while still more of them will keep their jobs but at a pay cut. These workers will have less money than they had before. But, if this downward wage pressure is occurring in more than just their industry sector, then they very well might find that their wealth has increased, because the prices of the goods and services that they are interested in have gone down. So while they have less money, each dollar they do have has much greater purchasing power.

Even more fundamental, however, is that unlike in poker, in the real economy, we make things. Every car that is made is a car that wasn't there yesterday, and that has value. All these things that are made, or services that are provided, have value even if they themselves are not money. It is for this reason that it is a mistake to think of our economic interactions as a zero sum game. The economy is productive, and so we must be in a positive sum game.¹²¹

It is important that we take ourselves to have established that the problem we face is positive sum. If we faced a zero- or negative-sum game, then a plural society would fail. In such games, the only way to gain is to take from others, and eliminate them from the game. Under such conditions, it would be psychologically impossible for agents to develop fellow feeling. Only positive sum games provide the appropriate incentives.¹²² It is a very fortunate fact that we do live in a positive sum world. Our first task is simply to recognize that we do.

Being in a positive sum game does not guarantee that everyone always benefits. After all, all economies are positive sum games, and there are plenty of

¹²¹ It is tempting to want to claim that the existence of a positive sum game should be one of the circumstances of justice, beyond what Hume specified. Positive sum games provide the opportunity for shared gains, rather than just merely agreeing to giving something up in an equitable manner.

¹²² This follows from the fact that in positive sum games, there is at least one cooperative option that increases social utility, at least in comparison to relevant alternatives. Even in the pre-industrial revolution period, when economic growth was safely under 1%, social utility would have decreased if something other than the cooperative option had been chosen. So too with the "stationary state" envisioned by Ricardo and Mill: though they imagine an upper limit of productive capacity, it can only be maintained through cooperation. If one party chooses to defect, social utility would decrease. This is unlike a zero sum game like poker – if I bluff and win a hand, no money is lost. It is just moved amongst the players.

people who clearly lose. But a positive sum game does guarantee that there is some way of arranging the benefits such that everyone does gain. This is why we can represent the question of distributive justice as a bargain. In doing so, what we are representing is the full range of possible allocation and distribution decisions that could be made in an economy. The particular bargaining solution we choose is then our statement on what kind of differential gains and losses we are willing to accept.

A downside of the bargaining representation is that it lumps together two important and distinct processes – production and distribution. Since we want to look at how these processes fare in a plural society, we must look at them individually and not as a single bargain. So let us now turn to the question of production.

Plural Production

Before we can determine how we ought to distribute wealth and income, we must first have some wealth and income to distribute. As such, we must first turn our attention to the production of wealth. While we are not interested in particular industries or even the distribution of industry, we are interested in the incentives that exist in the labor force to either be welcoming of or hostile to those who are different. This discussion follows the economic arguments for diversity that we examined in Chapter 1. In particular, we find that diversity leads to greater

economic productivity. The main reason for this is that the economy benefits from the division of labor, and a diverse society offers more opportunities for further division. Individuals can specialize at tasks best suited to their skill set, which allow them to reach their productive capacity. As an example, Alice, an expert mechanic, can focus on fixing cars all day, while Bob, a baker, can focus on baking bread. Since Alice is still interested in consuming bread, and Bob needs his car fixed, they can trade amongst themselves to make sure that they both get what they want. This is of course a simple illustration, but it is nevertheless useful for us to consider. What we see in this example is that diverse perspectives bring with them diverse problem-solving abilities. Whether it is on a micro-scale such as our example of a baker and a mechanic, or on a much larger scale where we have ethnic communities who take on certain labor functions¹²³, this principle of specialization has significant implications. This is important in two ways. First, job competition isn't direct. Second, the gains to trade are larger.

The first concern is about labor competition. Consider two societies, one that is homogenous, and the other that is diverse with respect to perspectives. Imagine that both Alice and Bob are looking for jobs, and are trying to determine which society they would rather be in. If they are in the more homogenous society, it is more likely that their skills are correlated with each other. In a society that has a

¹²³ It is not unusual to find immigrant communities specialize with certain businesses, like Korean groceries or Chinese laundries. Many immigrant communities carve out market niches for themselves in a similar manner. This standardization makes it easier for the communities to support business development efforts, since the business knowledge is shared in the community, and there are collective bargaining opportunities with suppliers to keep costs low.

shared culture, broadly shared values, and standardized education, we should expect that individuals would have a largely overlapping set of skills. Skill diversity needs to come from somewhere, and if individuals' upbringing is largely the same, there is less room for novelty.

We can see this case for diversity bringing novelty within the framework of perspectives. If two people share the same perspective, it must mean that many of those cultural or developmental factors were similar, for it is remarkably unlikely that two individuals who grew up in completely different environments will settle on identical ways of categorizing the world.¹²⁴ And so if two individuals do share the same perspective, they likely also share a huge cultural and educational background.¹²⁵ If they have different perspectives, on the other hand, it is likely that they have dissimilar backgrounds in at least some respect. This introduces an opportunity for a difference that is relevant to the sort of problem solving that is

¹²⁴ To press on how unlikely this is, remember that a categorization scheme is a full description of how a person organizes concepts. For, say, an Amish person with schooling through age 16 to have the same set of concepts as an African villager and a Ivy League educated hedge fund manager from a wealthy family seems radically implausible. At the very least, the Amish person would have more fine-grained distinctions of crafts and farming methods than the hedge fund manager, and likewise the hedge fund manager would have far more distinctions to make about modern technology.

¹²⁵ To press the point further, let us look at how shared cultural heritage can lead to shared skills. The traditional jobs that European Jewish people take on revolve around fields that do not require property ownership, such as medicine and banking. This is in large part a result of the history of Jewish persecution in Europe. Jews were not allowed to own property, and Christians for a long time were not allowed to lend money with interest or study anatomy on cadavers. As such, Christians needed doctors and bankers, but were religiously prohibited from performing such tasks. Jews were politically restricted from property ownership, but did not have religious restrictions on medicine or banking. So Jewish people now have a long history of specialization in such fields, even though the restrictions on both sides no longer exist.

relevant to one's job skills. So if two people share the same perspective, then their skill sets are also likely to be shared. When this is the case, Alice and Bob have more to fear from each other as labor market competition. If they do not have unique skills, but instead are merely two among many who share largely the same range of talents and knowledge, then they have reason to attempt to erect barriers to enter the labor market. If they were in the more diverse society, however, it would be more likely that Alice and Bob have different perspectives, and thus different skill sets. This would reduce the chance that they are fungible with each other, and would mean that they are less likely to be in direct competition for the same class of jobs. In this case, there is no reason to fear new labor market entrants. A plumber doesn't care if a number of mechanics move to town, as they are not going to be driving down her prices. If anything, the plumber has reason to be excited, because it should lower her car repair costs, and free up more households to spend money on hiring plumbers. The more the labor market is diversified, the more reason we ought to have for cheering its expansion. When the market is homogenous, however, we have good reason to try and place as many restrictions on it as possible.

To illustrate this case, imagine that you are trying to land a job as a philosophy professor. The question at hand is whether you would rather be in a society that is famed for its production of philosophy professors, or if you would rather be in a society of the same size that produces some philosophy professors, but a lot more professionals in other, reasonably disparate, fields. While it may be

fun to have so many philosophy professors to talk to in the first instance, it also would mean that there would be a much higher level of job competition than would be found in the second instance. So, economic rationality would point one to wanting to be in the second society.

So far this has just been a purely negative case. We welcome diversity because it means we avoid unwanted competition. But this is not sufficient. We should also be able to articulate a positive reason for wanting a diverse society. And if we continue with our example, it may be fun to be a philosopher in a sea of other philosophers, but even philosophers need to eat something, live somewhere, and wear something. And to do that, we need farmers, grocers, chefs, architects, engineers, tailors and shopkeepers, amongst others. This is a component of being able to engage in a system of social cooperation. We need to have many different kinds of tasks performed, and for that we require some division of labor.

The primary reason for wanting as much diversity as possible in production is not for reasons of avoiding competition, but the huge gains to trade that are made possible by having many different kinds of specialists. As I have argued in detail in Chapter 1, Ricardian trade theory shows that trade amongst diverse specialists increases social product more than production by generalists or homogenous specialists. The average individual benefits in two important ways from this. First, for any given class of good, there is either more of it or it is of higher quality (or both), due to the specialization of production. So no matter what my taste is, whatever it is that I like there is more for me to consume. By definition, rational

agents must prefer this to any alternative. The second way is perhaps more interesting. There is also a larger class of goods to sample. This translates not just into thirty different kinds of cans of soup at the grocery store, and lots of new ethnic restaurants to try, but it also means new cinema, new artwork, new music, new literature, and new clothing. These latter choices are not just consumer goods that one can have a preference about: they also represent entry points into new ways of living and experiencing the world. While agents do not need to opt to experiment with how they live in the world, some do want such options, and they are only made possible in an environment where different perspectives can be mixed and reinterpreted.¹²⁶ This goes beyond preference satisfaction, and creates the possibility for discovering new tastes and new self-conceptions.¹²⁷

It should be noted that individuals may well remain stubbornly against diversity. They could simply have deep-seated fears or prejudices against other kinds of people, and this would lead them to do whatever they can to block others from entry into society. This is a very plausible attitude, and one that seems to manifest itself quite often in the real world. Though I cannot provide an argument that each of these people will come to decide that they are wrong, and begin to see the value of human differences, I can argue that the incentives are aligned against

¹²⁶ It is possible that an agent could have a preference for stasis, and find new options to be a disutility. While I register the possibility of such “curmudgeon” preferences, I do not seriously consider them in my analysis.

¹²⁷ This opportunity for smaller-scale experiments in living is essential for the overall procedure of social contract formation that has been developed in the past two chapters. In particular, new experiments in living provide societies with new evidence for our moral beliefs.

their views. If we think of tolerance as a measure for how much difference one is willing to accept in other people's views and customs, then it is easy to see that the more tolerant one is, the greater the potential for material rewards. The more different others are, the more gains to trade there are. If one is only willing to interact with others like her, then little economic growth is possible.¹²⁸ Even if these individuals never correct their intolerant ways, they will be economically marginalized.¹²⁹ Those that embrace tolerant attitudes thrive because of the greater economic opportunity open in front of them. As we saw in the previous chapter, this can help explain why trade-oriented cities are very tolerant. Tolerance allows an individual to take advantage of economic opportunities that would otherwise be closed off. Tolerance can be thought of here as ideological leeway – it is the range of positions that are different from one's own that one is willing to put up with. As one increases that range, more economic interactions become feasible.

Let us now take stock of what we have found in the case of plural production.

We have seen that it reduces labor market competition, since people who hold different perspectives are less likely to have the same skill sets and thus be

¹²⁸ Milton Friedman makes a similar case for how capitalism eliminates discrimination in chapter 7 of his *Capitalism and Freedom* (1982). While I do not endorse the argument insofar as it should be taken to convince every individual, I do think it should be taken seriously to illustrate the general push of incentives at the macro level.

¹²⁹ This argument relies on the assumption that not everyone in society is intolerant. If there is some variation of levels of tolerance in society, then there would be cases of some doing better than others, and incentives would be made clear. However, if society is uniformly intolerant, then there is no symmetry breaker to see that being more tolerant makes one better off, and so no one would become more tolerant without some external factor, such as the 1964 Civil Rights Act in the United States. Thanks to Samuel Freeman for pointing this out.

competing for the same kind of job. Because of this, there should be lower levels of complaint (and even happiness) about new entrants to the job market. It is easier to see that the new entrants are not direct competition, and so it is easier to tolerate their entry. More important, however, is the fact that diverse perspectives create more specialization, which in turn creates greater returns to trade.¹³⁰ In fact, the more diverse the society, the more social surplus we can create. And insofar as diversity creates opportunities for new innovations, it fosters even more diversity. In production, there are only incentives for encouraging more diversity. This is true from a social perspective, but it is also true from an individual perspective: each individual can only be made better off as production is made more diverse.

While we have shown that individuals always have reason to want more diversity when we consider production, we have not yet considered the crux of our challenge: the question of distribution. It is to this question that we now turn.

¹³⁰ One objection to this view is that there is an upper limit to how much specialization can be useful. Eventually, more specialization would reduce production, because it does not create any new efficiencies, and instead might reduce them. This objection is true if we consider a single manufacturing process, such as Adam Smith's pin factory, but is not true if we consider an economy as a whole. It is certainly the case that in any given factory process, there are only so many tasks to perform. But, part of what specialization enables in general is new market niches and new products and services. The bundle of goods and services available to consumers today is far more extensive and varied compared to consumers 20 years ago, let alone 100 years ago. Part of what enabled that is new areas of specialization.

Distributive Justice in a Diverse Society

Thus far we have seen that economies are positive sum games, and that diverse societies are the most productive societies. But what we have not yet addressed is that diverse societies are also going to face the most disagreement in how we distribute the gains in production. This is undoubtedly a challenge, but this challenge is muted to some degree by the fact that as we increase diversity, we also increase the resources to be distributed.¹³¹ This fact will help us with our distributive problem.

Even so, we face a challenge of how to distribute the products of social cooperation. As we have seen, Rawls proposes the difference principle as an appropriate mechanism to organize the basic structure of society. But we continue to face the problem that this principle is unlikely to survive the psychological dispositions of agents in a diverse society. While I will not aim to defend a particular distributional rule, I will try to argue that rules that follow two principles will be more stable in a plural society. In particular, the choice of distribution should be on the Pareto frontier of the set of possible outcomes, and the gains

¹³¹ The reason for this is straightforward: as we have seen in the previous section, as we increase diversity in a society, holding everything else constant, we also increase the productivity of the society. This ought to, in turn, lead to greater social surplus. We now consider the fact that, even though increased diversity creates more social surplus, it at the same time increases disagreements about how to divide that social surplus. So even though this society has more disagreements because of more diversity, it also has more surplus than it would have were less diverse, and thus less prone to disagreement.

should be distributed in proportion to contribution.¹³² These principles are meant to ensure that no individual or group is being taken advantage of by another group.

Let's consider the first, less controversial, principle. This principle states that we ought to ensure that we not only make Pareto moves, but that we make maximally efficient distributions. The Pareto frontier of a bargain defines the set of outcomes of the bargain at which there are no exchanges of goods possible that would benefit one party without harming the other. These outcomes are all *efficient* since there is no wealth or income that could have been distributed but was not. All resources are being consumed. This first principle, then, just assures us that distributions in society are not wasteful. It also assures us that *no one loses*. Every move to the Pareto frontier guarantees that every agent involved in the bargain is made no worse off for having participated in the bargain. It is possible that one or more agents would not benefit, but they will not be harmed. This principle is essential for stability. If it were possible for some arbitrary group to lose ground in their absolute wealth and income, they would have no reason for supporting the institutions that caused this change. In fact, they would have very good reason for opposing those institutions. They would be in essence subsidizing the gains of others in the population without receiving any benefit in kind. This would be

¹³² The idea of "contribution" needs further clarification. As I will explain in some detail later, I determine one's contribution counter-factually. One's contribution is the difference in production one makes by participating in the productive process. An alternative approach that is a little less abstract is that contribution is what the agent or agents could have produced absent the cooperation of others.

irrational to endorse. The agents who are taking losses must then be either in fact irrational, or being coerced into participation. In either case, the justification for this is untenable. As such, this first principle is meant to limit the space of possible distributional rules down to those that are incentive-compatible. This principle assures that any distribution that is chosen is one where no agent is worse off for having agreed to it.¹³³

The second principle we will consider is more substantive than the first. It states that the gains from social cooperation should be distributed in proportion to contributions to production. This principle is based on two ideas. First, that insofar as we are treating society as a system of social cooperation, we must ensure that everyone engaged in that cooperation benefits from the fruits of that effort. The second idea is that no group should be required to subsidize another group such that they would have been better off on their own. These two conditions together impose significant constraints on how wealth and income can be distributed, but both are importantly based on the production step.

¹³³ It should be noted that this principle would almost certainly eliminate the difference principle from consideration, or at least place significant restrictions on it. This is for several reasons. First, the difference principle, as it is not properly a bargaining solution (in that it lacks a disagreement point), cannot be evaluated for whether no agent is made worse off. It is possible that the best off in society would be made less well off under this distributional rule. Second, the difference principle does not speak to any efficiency questions, and in fact may run counter to them. If inequality can only be justified by making the worst off best off, it may well be the case that there are Pareto moves left to make to get to the Pareto frontier, but the difference principle ignores them because the gains could not be transferred to the worst-off in society. This can happen if there is significant deadweight loss in social transfers. It is possible to imagine a reformulated difference principle that avoids these problems, but it would not be as Rawls presents it.

I would like to focus on this second idea, which I call the no-subsidy principle. I aim to argue that the no-subsidy principle must be at the heart of any attempts at distributive justice in a diverse society if maintaining stability is an active goal. Unlike in more homogenous societies, we are unable to rely on a broadly shared culture or shared history to provide social cohesion. This might seem to be a small matter, but I take it to be of great importance. Societies, particularly those facing some level of disagreement in the public sphere, need some reason to motivate citizens to want to resolve the issues and be willing to compromise.

Shared culture is a powerful motivator. It provides a reason for why a society “should” stay whole, and not split into separate entities. It also often means that disagreements in the public sphere are limited in their severity—a shared culture often implies a baseline set of values that are also shared. When Democrats and Republicans argue in the United States, they may disagree over a few percentage points on the top marginal tax rate, or the generosity of the social safety net, but neither side questions more basic principles, like whether we endorse market capitalism as an efficient method of allocating resources. The instance in which the United States did face the potential of splitting into two countries was largely an area in which differences in culture played a major role. Though it is undoubtedly the case that the North had a stronger moral position in regards to slavery, it was made easier to adopt in that the North had become largely an industrial region, while the South was firmly agrarian. Shared culture does not

guarantee stability of course—the Revolutionary War in the United States showed that—but it does go a long way in dampening conflict.

Without any assurance of shared culture, we must find some alternative social glue. Rational self-interest is a powerful social glue, particularly as the incentives for production all point toward greater levels of social cooperation with more diverse populations. While it may be the case that, over time, a more robust sense of country and citizenship could develop, this is not something that can be assumed. Some individuals, whose perspective involves some kind of nationalist identity, may well recognize others as citizens, and seek to treat them preferentially on those grounds. Others may be more cosmopolitan in nature, and just look to embrace other citizens as fellow moral agents who are worthy of concern. But this does not change the fact that many others, who may sort people based on race or some other identifier, will seek to treat many fellow citizens as members of out-groups. Because of this, we cannot rely on an assumption of goodwill between citizens. While we do not necessarily have to suppose that there would be an adversarial relationship between citizens, we cannot suppose that the relationship would be altruistic.

The main concern that a diverse society must face is some form of factionalization. If the various parties in a society have no natural alignment with each other, the onus is on the political apparatus of the state to show why they should tie their fates to each other. If the state is unable to do so, those parties may find it to be in their best interests to split with each other and form smaller unions.

Under these conditions, the reason parties have to stay together is that they are more able to provide themselves with material resources as a larger and more diverse economic unit. But this is compelling to any given party only so long as each party gets more than they could have achieved on their own. It is here that we can see the importance of a contribution-based account of distribution. If we think about contributions in terms of what an agent or agents could have produced on their own, we get a baseline measure for what they must get in order to be rationally motivated to remain in society. If this goes unsatisfied, it suggests that their productive capacity is subsidizing other parts of society – an effective “taking” of their wealth and income. This is not to express a Nozickean notion of absolute property rights: it is not that in practice, this principle implies no system of taxation. It just implies that rates of taxation or other forms of transfer must be limited by this baseline account of contribution. Individuals must have a rational reason to remain in their particular society rather than some other society, and if we are to suppose no cultural connections, we must then suppose that it is because the society is providing them with benefits greater than they can find with some other social arrangement (that is held to the same set of constraints).¹³⁴ If we adhere to a no-subsidy approach, we assure that this is the case.

¹³⁴ This last condition is just meant to prevent objections that the wealthy could always just claim that their alternative society is some kind of slave-owning kleptocracy where they have unlimited control of resources and political power. The counter-factual alternatives to be considered must be held to the same constraints that the actual society is held to, otherwise the counterfactual societies would have the same stability problems that we are aiming to resolve here.

We can now see that disagreements about how to divide resources, whether it is in actual distribution of wealth and income, or the choices over which public goods¹³⁵ to provide, must be resolved in such a way so as to respect the condition that each party gets more of what they want than what they could have provided for themselves. This is made easier by the earlier-noted fact that the more diversity-caused disagreement there is, the more there is to distribute. But this fact can only help us so much. What we must be most careful of is that, no matter what set of distributive goals we might have¹³⁶, we do not cause any one group to have reason to break off from society to provision their own goods.

As an example, consider Iraq. We can simplify the example by assuming three coalitions: Kurds, Sunni and Shia. Let's further assume that Kurds control most of the oil wealth, but Sunnis and Shi'ites provide useful labor in getting that oil

¹³⁵ Recall that public goods are defined as those goods that are both non-excludable and non-rivalrous. A good is non-excludable if, in providing it to some people, you are unable to prevent others from accessing it as well. An example of this is a lighthouse. Even if one shipping company builds a lighthouse for the benefit of their own ships, they cannot prevent the light from being seen by other boats. A good is non-rivalrous if it does not cost anything to provide it to additional people. While, say, providing chairs to people costs a fixed amount per new chair, operating the lighthouse does not cost extra if more boats pass by.

¹³⁶ Recall that we are being agnostic to particular goals of distributional justice. Some factions of society may want to distribute money towards the poor, some might want to distribute it towards the most productive members of society, or some might want to distribute it towards religious organizations. The question I am attempting to address here is largely independent of these concerns. I contend that a significant component of my framework is that societies can experiment with different distributional goals to see which ones best achieve the ends that society is interested in. As such, I do not wish to make any arguments for which distributional scheme is the best. However, I am interested in apportioning wealth and resources between competing factions in society. These factions can do with the money what they will, but the claim that I am pushing is that each faction must receive its fair portion of the resources in order for stability to be maintained.

to market. Kurds may be willing to accept distributing wealth in the direction of the other two groups, but only insofar as their absolute share of wealth remains larger than what it would be if they were autonomous and received no help. If their wealth share ever drops below this level, then they would be subsidizing the other coalitions and receiving nothing in return. There is no rational reason for going along with this subsidy, and so Kurds would simply exit from the social contract, and form a smaller one amongst themselves.

The goal of proportional distribution is to purchase a united citizenry. In the absence of a strong cultural bond, we are left only with economic ties. These ties can help to facilitate cultural connections in the future, but they must at least in the beginning serve as the primary means by which a society stays together. By distributing the gains of a union in such a way so as to make sure that every faction is made better off than they would have been otherwise, each faction has a reason for remaining in society. When proportional distribution based on production is violated, various factions might find their interests better served by forming a separate union.¹³⁷

¹³⁷ In some ways, similar issues sparked the Revolutionary War. Britain taxed imports at high rates, demanded that colonies buy only British goods, and instituted a number of other taxes that were excessive. Economic self-interest demanded that colonists have some governmental representation to resolve these taxation issues, and it was denied. Early instances of rebellion, like the Boston Tea Party, were explicitly economic complaints.

The Distribution of Public Goods

Up until this point, we have discussed economic distribution without regard to whether we mean the distribution of wealth and income, or the procurement of public goods. I would like to now turn to a short discussion of public goods as they raise some considerations that do not come up in the distribution of wealth and income.

When we distribute wealth and income, the effects of these distributions can be largely ignored: if people whom you consider undeserving receive money, you can at least stop paying attention. Money, after all is a private good that (at the margin) does not have any externalities. And once one chooses to ignore it, their lives can largely go on without having to think about it again. This is not necessarily the case with public goods. Public parks, trains, and abortion clinics¹³⁸ can all be hard to ignore if one has a deep-seated objection to the state providing these services. When one simply distributes money, it is easy to imagine that much of the money will be spent largely on fairly innocuous purchases: food, clothing and shelter tend to take up a large portion of people's expenses. But this does not happen when public goods are provided. You can either be happy with the good or not. If the state provides free abortion clinics, many people may be glad to know that such services are available for those that need them, but others might decide

¹³⁸ Trains and abortion clinics do not meet the precise definition of public goods, but are relevant in that providing them at a public scale involves collective action, and they have externalities.

that the state is forcing them to endorse something that they would otherwise vociferously object to.

The provisioning of public goods is where we can see the limits of toleration, and the most likely source of factionalization. If we assume that we are in a positive-sum game, there is always a distribution of wealth and income such that everyone is made better off than before. But with goods that have externalities, it is very feasible that agents can be negatively impacted by having them provided to others. This can be seen even with relatively “harmless” goods, like public schools. If one region or group decides to provision extremely high-quality public education, in part by paying high salaries to get the best teachers, their investment effectively crowds out the teacher investment made by other regions or groups with lower levels of expenditures. So all the best teachers will flow to the high-salary schools, and the other schools are left with lower-quality teachers. This creates a negative externality on the students in other school districts, by lowering the quality of their education.¹³⁹ Many goods, however, can become even more contentious than simple cases involving public schools. Conservative Christians are not going to be happy with abortion clinics, even if the rest of society wants them available. Similarly, if Christians decide to put up crosses or the Ten Commandments in the

¹³⁹ This sort of case is best captured by the theory of “local public goods.” Local public goods are close to public goods in that they exhibit non-excludability, but they are partially rivalrous. That is, they face a congestion problem: town pools and public schools can both become crowded. So this encourages localities providing local public goods for their own communities. This was first developed in Tiebout, Charles M. “A Pure Theory of Local Expenditures.” *Journal of Political Economy*. 64 (October 1956), 416-424. It was extended to a model of clubs in: Buchanan, J.M. “An Economic Theory of Clubs.” *Economica*. 32 (1965), 1-14.

public park, non-Christians would have every reason to become upset at the apparent evangelism. More generally, every faction is likely to have some public good that it would like to provide for its members that other factions would find objectionable. The set of public goods that everyone would like is most likely going to be fairly narrow.

There are two potential solutions to this problem. The first is to limit public goods to the set of goods that everyone can come to agree to. This would ensure that no one is upset with any perceptions of being forced to endorse something that they would otherwise object to. But it would have the consequence of drastically limiting the ability of the state to provision goods that many people want. This attempt to limit harms might in fact create more harm by denying so many people access to the public goods that they want. Additionally, it could well encourage factionalization on the basis of agreement over public goods. If people feel constrained by the perspectives of others, they will find a smaller group that shares their views on the desirability of particular public goods, and form a social contract with them.

For these reasons, we must turn to the second possible avenue, which is to allow the development of any goods that people are willing to produce, so long as the goods are paid for only by those that seek them.¹⁴⁰ As we saw in our

¹⁴⁰ As money is a fungible good, it may be the case that the best a state can do is assure that public expenditure on given goods is proportionate to interests. It would be nearly impossible to guarantee that a particular dollar meant for parks was spent on parks as opposed to the military.

examination of rights in the previous chapter, it is possible to bargain over these goods, so long as everyone abides by the principle that they pay for the goods that they seek, and do not attempt to force others to pay for them. It should be noted that this would likely constrain spending in a number of arenas. For example, it is unlikely that something like military spending could be maintained at the current US level, if people could indicate how their money is to be spent.¹⁴¹ This bargaining approach would encourage either goods that can be universal and thus garner reasonably high funding levels, or those goods that are of limited appeal, which would secure commensurately lower levels of financial support.

A consequence of this system of providing public goods is that those goods that are seen largely as wealth transfers would either have to be limited in size, or universal in nature. A system of welfare payments, then, would either mean fairly small checks to poor people, or larger checks to everyone.¹⁴² On this approach, we should expect to see a number of smaller goods be the dominant mode of public production. This is purely a function of the nature of collective action problems.

¹⁴¹ While it is clear that the state has an interest in requiring a certain funding level for a defensive force, offensive force is less clearly essential. It may be that, in times like WWII, the public would be willing to provide money for an offensive action. But this approach to appropriating money would make it very difficult to fund unpopular wars, like Vietnam. In terms of the system we are outlining, unpopular wars must be seen as an unwanted subsidy. The productive capacity of the state could have been redirected towards other goods that more people want. Popular wars, on the other hand, would not suffer from this problem.

¹⁴² This assumes that some portion of the population is not interested in wealth transfers, like those that are compelled by the “Cadillac driving Welfare Queens” argument that we examined earlier. If everyone finds wealth transfers appealing, then the circumstances change, and it would be possible for larger, targeted welfare systems. But this is a fairly unlikely scenario in a diverse society.

Less expensive goods require fewer agents to pay for them, and as it is easier to arrange a coalition of smaller sizes than larger sizes, smaller goods are more likely to be paid for than larger goods, even if the larger goods are even more desirable.

Though this bias toward the production of smaller-scale goods may seem problematic, it in fact has a number of virtues. First among them is that a society can compare the effectiveness of a number of smaller projects with similar goals, and over time increase support for the project or projects that have proven to be the most successful. And for those projects that may be met with skepticism initially by the public at large, it provides an opportunity to demonstrate their usefulness on a small scale to build larger support. In a society filled with diverse perspectives, we must assume that many different groups will advocate alternative approaches. What this bias towards smaller projects does is allow all groups to have the opportunity to test their ideas, and try to prove to others that their approach has the most merit. It allows all groups to be included in ongoing social experimentation, and can organically represent mounting support for a given approach when that good is able to appropriate more funding than other approaches. Rather than try to resolve these differences with a priori debates and the imposition of a universal solution, this bias allows actual experience to dictate what approaches should win the debate. Following this principle, the state can be seen in the role of a clearinghouse of sorts for the provisioning of rights and public goods. Individuals

can state their interests, and their willingness to pay, and once a coalition exists to support the production of a good, the state sanctions its creation.¹⁴³

There is still one remaining challenge in how we treat the provision of public goods. Even if Christians pay for their crosses, Muslims pay for their mosques, and women's rights activists pay for abortion clinics, they are imposing externalities on others simply because of the existence of these goods. People who do not subscribe to Christianity or Islam could easily be offended by blatant religiosity if it violates their own views. Anti-abortion advocates could feel harm by the presence of abortion clinics if they see them as sites of murder. Third parties can suffer harms from the transaction or creation of goods, even if they have no part in them. This is easily seen in the case of second-hand smoke: breathing in smoke harms people even when they did not choose to smoke a cigarette. Likewise, forced exposure to goods that represent values that are seen as undermining one's own is potentially harmful.

Local public goods are not the only goods that have this problem with externalities. Another potential source of externalities that must be accounted for are positional goods. Positional goods are those goods that derive their value from their relation to each other. In particular, positional goods suppose some ranking of status or desirability. People use positional goods to compete for status with each other. Because of this quality, competitions over positional goods are zero-sum in

¹⁴³ This view has several implications about the broader function of the state, as well as the dynamics of goods production. However, I will leave developing this for future work.

nature. Gaining position must mean that someone else has lost in position. This zero-sum characteristic of positional goods implies that they have negative externalities. Examples of positional goods are having an “exclusive” address, or a bigger-than-needed house, or a large diamond on one’s engagement ring. These are sought after not for their intrinsic qualities, but to demonstrate status. Likewise, people spend money on cars that go well above any speed limits to compete over who has the fastest car, even though that speed can never be legally attained on public roads.¹⁴⁴ These goods exist primarily for their external effects. As such, they ought to be subject to the same principles that govern public goods.

The previous chapter set up a mechanism for dealing with balancing out the positive and negative externalities created by rights allocations. This mechanism can also be applied to public and private property, but here we will see some additional caveats. In particular, in the previous procedure, all rights could be priced in terms of the limitations on other rights that they created. When we have goods that have a monetary value, this pricing mechanism is less straightforward. The root of the problem is that we have two different types of goods that we are attempting to distribute. First, we have the wealth and resources that we have already committed to distribute in proportion to contribution, so as to not violate a

¹⁴⁴ Some positional goods are also Veblen goods, which are commodities for which the demand curve is the inverse of what it is for normal goods. That is, as the price increases, demand increases. This is because the price is part of the value of the good, insofar as it can be used as a signaling device for status and quality. Expensive wine, chocolates, cars and jewelry can fall into this category.

no-subsidy principle. But when we factor in externalities as well, we have to rely on the language of harms, as economic externalities are by their nature not priced.¹⁴⁵

Assuring that the initial distribution of wealth and income satisfies the no-subsidy principle and simultaneously that the harms caused by the externalities from public and positional goods are equally distributed across society would require very special circumstances. It is in effect a mutual maximization problem, which cannot be solved except in very unusual and constrained circumstances. Because of this, one of these principles will have to be treated as subordinate to the other. However, no matter which principle is chosen to be subordinate, we will necessarily weaken the case for stability. If we choose the no-subsidy principle, then it means that some group will be subsidizing other groups, which is not in their rational self-interest. If we choose the principle of equal burden sharing, then some group would have to take on more burdens than others, which they might consider to be unjustifiable. The important question is then not *whether* one of these principles has to be violated, because one will, but *how much* it is violated. This will depend on the particular circumstances: the mix of goods that people are interested in, the range of perspectives and utility functions, and the productive capacity of various groups. Small deviations from these principles might be acceptable, but as the violations become larger, social stability will become more and more tenuous, as

¹⁴⁵ If externalities were priced, that would mean that the social costs of the good would be included in the price of the good. But then the outside harms would not be externalities: they would just be costs. And those costs would be included in the price, and just made “internal.” So there would no longer be any externalities.

the social glue of rational self-interest will weaken and eventually disappear, at least for some portion of society.

What this means is that there will be cases where diverse societies will simply not be stable enough to maintain themselves as a single political unit. In the case where disagreement outstrips the gains in social surplus, societies will—and should—break apart. However, given that diversity is able to generate so much social surplus, and our distributive approach is designed to promote a social experimental approach to resolving disagreements about public goods provisioning, while attempting to assure that no group is suffering in comparison to realistic alternative social arrangements, social stability should be the norm. However, it is important to realize that, absent extensive social-cultural bonds, we must expect that not every collection of people can successfully form and maintain a political union. Though there are many incentives for promoting diversity in society, there are also limits on how far those incentives can take us.

Conclusion: The Limits of Diversity

The traditional liberal account of tolerance, as exemplified in Locke's *Essay On Toleration*, is one of duties. We ought to tolerate (religious) differences in others, and not force our views on them in part because we would not want them to be able to do the same to us if the tables ever turn. But this view does not present any case for why one might want a more diverse society. Instead, the assumption is that

diversity is a problem, and it needs to be managed. This may be a satisfactory account in relatively homogenous societies, but as societies become more diverse, we cannot depend on people following obligations that they perceive to be only burdens and restrictions on their activity. As the Reagan example from the beginning of this chapter demonstrates, people are quick to see people different from themselves as a burden, and from there, people seek to relieve themselves of the perceived burden.

Part of what I have argued for in this chapter is an alternative conception of tolerance. Instead of relying on a notion of obligation, I have argued that individuals have a self-interested reason for wanting to promote diversity in their society. Diversity creates social surplus. Diversity also creates new opportunities for social innovation, opening up possibilities for living one's life that were not previously available. This positive case for diversity will typically overwhelm the perception of burdens imposed on individuals from being in a diverse society. While this self-interested embrace of diversity will not always ultimately win out, it will help to promote and sustain societies much more diverse than what a Lockean account of toleration could.

Though it is a fact that societies are becoming more diverse, one could potentially argue that rather than abandon the liberal account of tolerance, along with other aspects of the liberal framework, in order to adapt to these more diverse societies, we should aim to pull back and return to more homogenous societies. However, in doing so, we would lose something that only really becomes possible in

a diverse society. With an embrace of diversity comes recognition of the importance of experiments in living. These experiments provide the evidence needed to further refine our moral concepts, our ideas about what rights we ought to guarantee, and how we ought to organize our lives as constrained by our moral beliefs and our rights and privileges. Different ideas of social importance, like provisioning public goods, can be tried and successful ideas will slowly become more dominant as people learn about their success.

As the benefits of diversity must be weighed against its costs, I have argued for the constraints necessary for the promotion of both diversity and the framework of social experimentation. Pareto optimality, the no-subsidy principle, and equal burden sharing aim to assure that every individual is made better off in their current society, as compared to possible relevant alternatives. These constraints aim to assure the stability necessary for the benefits to diversity to manifest themselves. If people feel like they are being treated fairly, and not being taken advantage of by others, they become more willing to embrace diversity, and the constant flow of changes and opportunities it brings.

Conclusion

Let us now consider what has been shown, and then look to what work lies ahead. What we have seen is that if diversity is taken seriously, much of social contract theory is subject to revision. In particular, I have argued for the central role of perspectives in the design of the procedures to generate a social contract. This central role of perspectives leads us to abandon the model of public reason for political discourse, and replace it with a bargaining model. This is further supplemented by a no-subsidy principle for the determination of the distribution of wealth and income in society.

Though these components have been motivated from an interest in developing an account of social contract theory for diverse states, they also offer a distinct epistemological approach to social contract theory. This is most explicit in the account of the View from Everywhere, in which I argue that a moral stance must properly be understood as a social stance. But this approach can be seen throughout. The bargaining model articulates a method of political negotiation that does not unduly burden either side with having to take on perspectives that they are not able to take on. And this bargaining model is importantly checked by the real-world results that it produces. In this contract framework, individuals in society are asked to try out particular rights allocations, and then try out different ways of living within the confines of that rights framework. As they better understand what options for life they have within a particular rights regime, individuals can then

assess whether or not they would like to keep, refine, or abandon the current contract and follow the procedures again armed with new information. It is this ability to return to the procedures for forming a social contract that is starkly opposed to other accounts of social contract theory. However, I take this to be a key improvement. We can only assess the wisdom of particular political decisions once we have some information about them. Often, this means that we must see what it is like to live with them. But if we are to take a more experimental approach to the choices that frame the basic structure of society, we must not be forced to stick with choices that have proven to be mistakes. Revising our contract allows us to be able to take more risks without having to be permanently stuck with the negative consequences if they arise.

Though I take this project to have developed the key components of a social contract, through the articulation of a initial moral standpoint, a procedure for the distribution of rights, and an account of the distribution of wealth and income, there is much that I wish to develop in future work. This project has led me to begin thinking about an alternative account of the state and its role in public decision-making. As I briefly noted in the final chapter, I have the beginnings of a view of the state as a clearinghouse for rights and goods. This can also be seen in the View From Everywhere and in the bargaining model. The state's main function is to serve as a facilitator for the views and interests of the people. On this view, legislators would not have clear powers beyond being proxies for their constituents.

I believe that this has wider implications for an account of the state. The political philosophy literature generally takes a country's borders as fixed, and does not often address the role of states or regions versus a central government. In large part, this is because political disagreement is addressed at a higher level, and regional differences are abstracted away. If we extend the account of stability developed in the final chapter, however, we can find that more diverse states can remain more stable if we allow for greater regional autonomy. That is, rather than requiring everyone to agree to all of the same rules, we can allow a smaller subset of agreement to establish a central government, and then each region can better focus on larger levels of agreement on more matters of justice. This parallels a confederation model of government as opposed to a federation model, which requires much more central control.

I believe that adopting a confederation model leads to two important ideas. First, this analysis pulls us away from the model of the nation-state as an appropriate political unit, and towards a model of federations of city-states. This increase in regionalism allows for larger-scale social cooperation where it can be agreed upon, but reduces conflict by allowing for greater local differentiation. Second, it will allow me to extend this analysis to develop a more robust account of Millian experiments in living. On this account, not only do individuals have the freedom to try small-scale experiments in living, but states are then encouraged to conduct larger-scale social experimentation. On this view, city-states must in effect compete for citizens, and are thus driven to experiment in what public goods they

provide, and what life plans they can best support. Successful city-states draw in more members, which in turn pushes other city-states to adopt similar policies as they attempt to innovate on their own.

This competitive model of regional confederations allows me to present the entire political framework as a dynamic process. While my previous work has established the epistemological benefits of leveraging diverse perspectives, it has not considered long-run social dynamics. These dynamics not only allow us to develop a richer account of stability, but also deepen our understanding of the true role of social experimentation in shaping liberal societies. In the competitive model that I look to develop, the evolutionary dynamics of the social contract are brought into the forefront. This is essential to understanding the effects of a procedural account of justice. In this sense, I can claim that justice should be understood as a trajectory, and not a state of being. Rather than claiming to know what the ideals of justice are, this view articulates the procedure by which we can find them.

Bibliography

- Kenneth Arrow, "A Difficulty in the Concept of Social Welfare", *Journal of Political Economy* 58(4) (August, 1950), pp. 328-346
- Kenneth Arrow, "Some Ordinalist-Utilitarian Notes on Rawls' Theory of Justice", *Journal of Philosophy*, 1973.
- Cristina Bicchieri, "Cooperation and Communication: Group identity or Social Norms?" in N. Gold [ed], *Teamwork: Multi-Disciplinary Perspectives* (London: Palgrave-Macmillan, 2005).
- Cristina Bicchieri, *The Grammar of Society* (Cambridge: Cambridge University Press, 2006)
- Marilynn Brewer, "The Social Self: On Being the Same and Different at the Same Time." *Personality and Social Psychology Bulletin*, 17, (1991) 475-482.
- James M. Buchanan, "An Economic Theory of Clubs." *Economica*. 32 (1965), 1-14.
- Thomas Carlyle, "Occasional Discourse on the Negro Question" *Fraser's Magazine for Town and Country*, London, Vol. XL., (February 1849).
- Gregory Clark, *A Farewell to Alms*. (Princeton: Princeton University Press, 2007)
- Ronald Coase, "The Problem of Social Cost", *Journal of Law and Economics* (October 1960)
- Ronald Coase, *The Firm, The Market, and The Law* (Chicago: University of Chicago Press, 1990)
- Donald Davidson, "On The Very Idea of a Conceptual Scheme" *Proceedings and Addresses of the American Philosophical Association*, Vol. 47, (1973 - 1974), pp. 5-20.
- Jared Diamond, *Guns, Germs and Steel* (New York: W.W. Norton & Co., 1997)
- Marc Ereshefsky, "Species Pluralism and Anti-Realism", *Philosophy of Science*, 65 (March 1998) 103-120
- Milton Friedman, *Capitalism and Freedom* (Chicago: University of Chicago Press, 1982).
- David Gauthier, *Morals by Agreement* (New York: Oxford University Press, 1999)
- Allan Gibbard, "Manipulation of voting schemes: a general result", *Econometrica*, Vol. 41, No. 4 (1973), pp. 587-60

Jurgen Habermas, 'Reconciliation through the Public Use of Reason: Remarks on John Rawls's Political Liberalism', *The Journal of Philosophy* XCII(3), 1995. 109–31.
Russell Hardin, *Collective Action* (New York: RFF Press, 1982)

John Harsanyi, "Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking", *Journal of Political Economy* 61(5) (1953) 434-5

Friedrich Hayek, "The Use of Knowledge in Society," *American Economic Review*. XXXV, No. 4. (1945) pp. 519-30.

Albert Hirschman, *The Passions and the Interests*. (Princeton: Princeton University Press, 1997)

Thomas Hobbes, *Leviathan*, (New York: Hackett, 1994)

Saumitra Jha, "Trade, Institutions and Religious Tolerance: Evidence from India" (January 10, 2008). Stanford University Graduate School of Business Research Paper No. 2004 Available at SSRN: <http://ssrn.com/abstract=948734>.

John Locke, *Locke: Political Essays*. Mark Goldie [ed] (Cambridge: Cambridge University Press, 1997)

Lu Hong and Scott Page, "Groups of diverse problem solvers can outperform groups of high-ability problem solvers" *Proceedings of the National Academy of Sciences*, November 16, 2004

Hong and Page, "Problem Solving by Heterogeneous Teams" *Journal of Economic Theory* 2001

Lloyd, Lisa. *The Case of the Female Orgasm: Bias in the Science of Evolution*. (Boston: Harvard University Press, 2005)

J.S. Mill, "The Negro Question", *Fraser's Magazine for Town and Country*, London, Vol. XL., (February 1849).

Mishler, B. and Donoghue, M., 1982, "Species Concepts: A Case for Pluralism", *Systematic Zoology*, 31: 491-503

Thomas Nagel, *The View From Nowhere* (Oxford: Oxford University Press, 1986)

Thomas Nagel, "Rawls on Justice." *The Philosophical Review*, Vol. 82, No. 2. (Apr., 1973), pp. 220-234

Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974)

Mark A. Satterthwaite, "Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions", *Journal of Economic Theory* 10 (April 1975), 187-217

- Mancur Olsen, *The Logic of Collective Action* (Boston: Harvard University Press, 1971)
- Scott E. Page, book *The Difference* (Princeton, Princeton University Press, 2007)
- John Rawls, *A Theory of Justice* (New York, Belknap Press, 1971)
- John Rawls, "The Idea of an Overlapping Consensus" *Oxford Journal for Legal Studies* (Spring 1987), 7 (1): 1-25
- John Rawls, "Kantian Constructivism in Moral Theory", *Journal of Philosophy* (September 1980), 77 (9): 515-572
- John Rawls, *Lectures on the History of Political Philosophy*, Samuel Freeman, ed. (New York, Belknap Press, 2007)
- John Rawls, *Political Liberalism* (New York, Columbia University Press, 1993)
- John Rawls, "Some Reasons for the Maximin Criterion", *American Economic Review* (1974)
- Rousseau, *Discourse on Political Economy and The Social Contract*. Christopher Betts, [ed] (Oxford: Oxford University Press, 2009)
- Justin Scheck, "Mackerel Economics in Prison Leads to Appreciation for Oily Fillets", <http://online.wsj.com/article/SB122290720439096481.html>
- Sen, *Inequality Reexamined* (Boston: Harvard University Press, 1992)
- J Sidanius, F Pratto, M Mitchell, "In-Group Identification, Social Dominance Orientation, and Differential Intergroup Social Allocation." *Journal of Social Psychology*, 134(2), 151-167, 1994
- Adam Smith, *The Theory of Moral Sentiments* (New York, Liberty Fund Inc, 1984)
- Adam Smith, *The Wealth of Nations* (New York, Bantam Classics, 2003)
- Tajfel, H., Billig, M., Bundy, R. P., & Flament, C. "Social categorization and intergroup behavior." *European Journal of Social Psychology*, 1, 149-177. (1971)
- Charles M. Tiebout, "A Pure Theory of Local Expenditures." *Journal of Political Economy*. 64 (October 1956), 416-424.
- William Wimsatt, "Robustness, Reliability and Overdetermination", in *Re-Engineering Philosophy for Limited Beings* (Boston, Harvard University Press, 2007)
- Wolpert and Macready, "No Free Lunch Theorems for Search" (Technical Report SFI-TR-95-02-010). *Santa Fe Institute*, 1995

M. E Yaari and M. Bar-Hillel, "On Dividing Justly." *Social Choice and Welfare* (1984) 1:1-24

H. Peyton Young's, *Equity: In Theory and Practice*, (Princeton: Princeton University Press, 1995)