

The Problem of Sorting: A Dilemma Between Freedom of Association and Social Justice

by

Valerie-Jean Soon

Department of Philosophy  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Gopal Sreenivasan, Advisor

\_\_\_\_\_  
Wayne Norman

\_\_\_\_\_  
Allen Buchanan

\_\_\_\_\_  
Alex Rosenberg

\_\_\_\_\_  
Robert Brandon

Dissertation submitted in partial fulfillment of  
the requirements for the degree of Doctor  
of Philosophy in the Department of  
Philosophy in the Graduate School  
of Duke University

2021

ABSTRACT

The Problem of Sorting: A Dilemma Between Freedom of Association and Social Justice

by

Valerie-Jean Soon

Department of Philosophy  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Gopal Sreenivasan, Advisor

\_\_\_\_\_  
Wayne Norman

\_\_\_\_\_  
Allen Buchanan

\_\_\_\_\_  
Alex Rosenberg

\_\_\_\_\_  
Robert Brandon

An abstract of a dissertation submitted in partial  
fulfillment of the requirements for the degree  
of Doctor of Philosophy in the Department of  
Philosophy in the Graduate School of  
Duke University

2021

Copyright by  
Valerie-Jean Soon  
2021

## Abstract

What does justice require of us in the personal realm of intimate association?

Many seemingly innocent choices at the heart of private life, such as our tendency to sort ourselves into groups of friends, partners, and neighbors who share our social identities, play a significant role in perpetuating social segregation and its concomitant inequalities. And oftentimes, we rely on structures such as informal social norms, selective rules of club membership, or local land-use policies to protect these associational choices. I call this the *problem of sorting*. Because these choices are in the private realm, there is significant tension between personal liberty, especially the freedom of association, and the demands of justice. I begin this dissertation by using empirical social science to outline the problem of sorting. I argue that individuals have reasons to take political action in favor of shaping their search environments according to these criteria, and they also have reasons to offer social opportunities to the excluded. The state should take positive steps to create diverse search environments. However, the state ought not interfere in our actual associations, and individuals have no duties to form substantive associations with the excluded. I conclude that this understanding of freedom of association and its concomitant duties reduces the tension that we began with, though it does not eliminate it.

## **Dedication**

To my grandparents, Chang Sok Ing and Chia Meng Liat

# Contents

<b>Abstract</b> .....	<b>iv</b>
<b>Acknowledgements</b> .....	<b>ix</b>
<b>Introduction</b> .....	<b>1</b>
<b>1. The problem of sorting</b> .....	<b>5</b>
1.1 Introduction.....	5
1.2 Defining the problem of sorting.....	6
1.2.1 Persistent inequalities .....	7
1.2.2 Homophily, <i>aka</i> sorting.....	11
1.2.3 Egalitarian ideals .....	19
1.2.3.1 Formal equality of opportunity .....	19
1.2.3.2 Fair equality of opportunity .....	22
1.2.3.4 Luck egalitarianism .....	25
1.2.3.5 Takeaways.....	29
1.3 The scope of sorting .....	30
1.3.1 The problem of the family.....	30
1.3.2 The problem of sorting is distinct from the problem of the family.....	36
1.4 The depth of sorting.....	41
1.4.1 The inadequacy of “bypassing” solutions .....	41
1.4.2 The inescapable mathematics of segregation.....	46
1.4.2.1 Asymmetric homophily of the privileged.....	47

1.4.2.2 Symmetric homophily .....	48
1.4.2.3 Asymmetric homophily of the less-privileged .....	49
1.4.2.4 Segregation that no one seeks .....	52
1.4.2.5 Means to attain instrumental goods.....	53
1.4.2.6 Information transmission.....	57
1.5 Conclusion.....	59
<b>2 Exclusion, entitlements, and dilemmas of sociability .....</b>	<b>61</b>
2.1 Introduction.....	61
2.2 Private and public discrimination: parallels and a puzzle .....	67
2.2.1 Exclusion as a collective harm .....	67
2.2.2 Parallels between public and private exclusion.....	69
2.3 Justifications for private exclusion.....	72
2.3.1 Autonomy and association .....	73
2.3.2 The right to private inclusion .....	82
2.4 Dilemmas of sociability .....	85
2.5 Associational opportunities .....	96
2.6 Conclusion.....	100
<b>3 An intrapersonal, intertemporal solution to an interpersonal dilemma .....</b>	<b>103</b>
3.1. The structure of collective action problems.....	103
3.2. The inefficacy problem .....	108
3.3. An intrapersonal, intertemporal solution .....	118
3.4. Why stick to your commitments? .....	132

3.6. The content of reasons to cooperate.....	139
3.7. Conclusion.....	141
<b>4 The ecology of freedom of association .....</b>	<b>143</b>
4.1 Introduction.....	143
4.2 Background: homophily and racial segregation.....	146
4.3 The priority of freedom of association .....	150
4.4 An ecological account of freedom of association .....	154
4.4.1 Three conditions for freedom .....	158
4.4.2 The scope and upshots of the ecological account .....	164
4.5 Objections .....	166
4.6 Social structure as the site of intervention .....	171
4.7 Conclusion.....	179
<b>Conclusion.....</b>	<b>180</b>
<b>References .....</b>	<b>184</b>
<b>Biography .....</b>	<b>198</b>



## Acknowledgements

I am grateful to many people for supporting me during the course of writing this dissertation. To my committee members – Wayne Norman, Gopal Sreenivasan, Allen Buchanan, Alex Rosenberg, and Robert Brandon – who provided valuable discussion and feedback on numerous drafts. To Felipe De Brigard, Walter Sinnott-Armstrong, Owen Flanagan, Jennifer Jhun, Katherine Brading, Janelle Haynes, Lisa Olds, Stacey Meyers, Dechelle Harris, and Maria Wisdom for their invaluable support and encouragement.

I thank these wonderful fellow graduate students for their ideas and friendship: Angela Bischof, Kobi Finestone, Ewan Kingston, Tim Burkhardt, Sungwoo Um, Joe Nelson, H. Bondurant, Hannah Read, Tara Weese, Kelly Shi, and Lauren Bunch.

I thank my friends and family for their support, in particular: my ecology friends Alice Carter and Jackie Gerson; Wellesley philosophers Corinne Gartner and Claudia Yau; and my Durham family Jennifer Lamy, Tess Ruderman, Kristen Cowens, Emily Atkins, Katie Hargreaves, and Nora Schwaller. I am grateful to my parents, Chia Puay Eng and James Soon, for teaching me discipline and persistence, and to my sisters, Candice and Carissa Soon, for keeping me grounded.

And of course, I'm grateful to my four-legged friends, who made sure I got enough walk breaks and outdoor time to complete this dissertation: Gaia, Hera, Marty, Red, and Ribbons.

## Introduction

Our social lives consist of a multitude of associations. We choose to befriend, live among, join clubs with, and date and marry some individuals rather than others. Part of what it means to live an autonomous life, in which one is free to live according to one's conception of the good, is the ability to make these associational choices. For this reason, liberals consider the freedom of association – the right to associate with whomever we want, as long as they also choose to associate with us – to be sacrosanct. In John Stuart Mill's words, we have the "right to choose the society most acceptable to us" (Mill 1859). Neither the state nor third parties may interfere with our associations, so long as these associations do not harm others.

The freedom of association is an intuitively compelling negative right or basic liberty, and its protection is a minimal condition for any liberal state. Only an illiberal state, such as Oceania in Orwell's *Nineteen Eighty-Four*, would interfere with our associational choices. But there is a tension between freedom of association and ideals of justice. In particular, freedom of association can undermine equality of opportunity: the principle that individuals' life prospects should not be influenced by morally arbitrary factors unrelated to talent and effort. Because our associational lives are persistently segregated along the lines of race and class, associational life does influence individuals'

life prospects to a significant degree. Those who are excluded from certain social networks end up also being excluded from economic and social opportunities.

This dissertation explores this tension, which I call the *problem of sorting*. On the theory I develop, the problem of sorting is best addressed by substantive social integration. We ought to shape our social search environments so that they provide social opportunities and are adequately diverse. Individuals have reasons to take political action in favor of shaping their search environments according to these criteria, and they also have reasons to offer social opportunities to the excluded. The state should take positive steps to create diverse search environments. However, the state ought not interfere in our actual associations, and individuals have no duties to form substantive associations with the excluded. These conclusions are by no means radical, but they take us beyond prevailing thinking on these issues.

Chapters 1 and 2 outline the problem space. In Chapter 1, I argue that the problem of sorting is an irreducibly social problem. It cannot straightforwardly be solved by distributive or structural means. Both individual action and structural reforms are required to address the problem of sorting, but the former faces a collective action problem, and the latter might conflict with personal autonomy. Thus, the conflict

between the private realm of social life and the public realm of justice must be grappled with.

In Chapter 2, I address our inconsistent moral and legal treatment of public and private social exclusion. Public exclusion, also known as segregation, is seen as a serious harm to the excluded, and its victims are entitled to redress in the form of inclusion in public life. But although private exclusion has similar consequences, we do not see private exclusion as a moral wrong, and we are uncomfortable with the idea that its victims are entitled to redress in the form of any particular person's association. I argue that the concept of associational opportunities can justify this apparently inconsistent treatment. We wrong people when we discriminatorily fail to offer associational opportunities, but we do not wrong people by denying them our actual association.

Chapters 3 and 4 offer solutions to the problem of sorting. Chapter 3 addresses the collective action problem feature of sorting: if the state ought not interfere in associational life, it is up to individuals to choose their associates in ways that promote rather than undermine justice. But because the harms of sorting emerge only collectively, and no individual's action makes a difference to this collective harm, it does not seem rational for any individual to change their associative behavior. It also does not seem appropriate for individuals to form associations with others for the sake of justice,

as this undermines the purpose of social relationships. I argue that the framework of diachronic rationality gives individuals an instrumental reason to act well, and that acting well requires undertaking political action to change the social environment. The focus on instrumental reasons is in line with my general commitment to political liberalism's neutrality about the good.

Chapter 4 turns to the concept of freedom of association itself. I argue that freedom of association should be understood ecologically: it requires that the social environment be adequately diverse so as to provide individuals with substantive associational choices or opportunities. Social conditions can interfere with freedom of association just as much as external interference by agents can. I end with some implications for the design of social search environments such as neighborhoods and online social networks.

This is a normative theory, so I make no claims as to whether my theory will actually defuse the practical problem of sorting. In fact, I am skeptical that it will. It may end up being the case that the problem still persists, because ample social opportunities and diverse search environments are insufficient for promoting substantive integration. In any case, this dissertation presents an attempt to fully grapple with the conflicts presented by our private, social choices.

# 1. The problem of sorting

## 1.1 Introduction

Our social lives are composed of a multitude of associations. We choose to befriend, live among, form clubs with, and date and marry some individuals rather than others. However, our associational life is highly segregated – we tend to interact with people who are similar to us along certain socially significant dimensions, especially race and socioeconomic status. Pair social segregation with the fact that significant inequalities persist among social groups. The result is that social and economic opportunities remain concentrated within certain groups, rendering them inaccessible to outsiders. Thus, associational life can undermine several egalitarian ideals. Yet, the idea that we should be free to associate with whomever we wish lies at the heart of liberalism. Few things seem more illiberal than state intervention in our social lives.

I call this the *problem of sorting*: our associational choices can and do cause social segregation, which undermines several ideals of justice. Unlike coercively imposed segregation, as exemplified by regimes such as South African apartheid and American Jim Crow, sorting results from many individual choices. These choices are shaped both by preferences and by social structures such as norms and policies. This chapter explores

the problem of sorting and argues that it presents a wide, unique, and deep dilemma for liberalism. Section 1.2 defines the problem of sorting and argues that it is wide: it undermines various ideals of justice. Section 1.3 argues that the problem of sorting is unique: it is distinct from and just as troubling as the nearby problem of the family. Section 1.4 argues that the problem of sorting is deep: because social segregation can and does arise from complex processes, it is unlikely that the problem can be completely resolved on either a theoretical or a practical level by public policies that address stigma or historical injustice alone. Section 1.5 concludes that sorting is an irreducibly social dilemma between freedom of association and equality of opportunity.

## ***1.2 Defining the problem of sorting***

There are two notable facts about social life in the United States, and indeed, in most polities. First, inequalities persist among social groups, particularly among races and by definition, among classes. Second, our associations are primarily circumscribed by these categories; it is still relatively uncommon to associate outside the race or class to which we belong. The problem of sorting emerges at the confluence of these facts. In this section, I give an overview of the empirical literature concerning these facts and show how they generate the problem of sorting. Sorting functions as a mechanism of

opportunity hoarding: the concentration of opportunity within certain groups to the exclusion of others. As a result, it threatens several ideals of justice.

### **1.2.1 Persistent inequalities**

Despite an end to legal discrimination, social and economic inequalities remain a persistent feature of American life. The black-white wealth gap is staggering: the median wealth of black households is \$17,150, whereas the median wealth of white households is about ten times that (\$171,000). The median household income for Whites is \$76,057, compared to \$46,073 for Blacks, \$56,113 for Hispanics, and \$98,174 for Asians.

Inequalities in educational attainment roughly track economic inequalities: in 2019, while over 40% of non-Hispanic whites and 58.1% of Asians had a bachelor's degree or higher, just 26.1% of Blacks and 18.8% of Hispanics attained this level of education (Census Bureau 2020).

While these racial inequalities are trending in the direction of progress, they remain remarkably persistent. They are also taking place against the ominous trend of increasing economic inequality, which inhibits the closure of the racial wealth and income gaps. Income inequality has been on the rise in the U.S. Since the 1970s, the share of pre-tax income going to the top 1 percent of earners has almost doubled, while the share of pre-tax income held by the bottom 50 percent of earners has remained stagnant



(Saez et al 2016). While income inequality is not a uniquely American phenomenon, it is worse in the United States than in other comparable nations: as measured by the Gini coefficient, it is higher in the U.S. than in other OECD countries.<sup>1</sup>

These inequalities have deleterious effects. Inequality between low-income households and the rest of the population is associated with lower levels of economic growth, primarily because it depresses the ability of low-income individuals (especially children) to develop skills (Cingano 2014). “Economic anxiety” has been blamed for the rise of right-wing movements (Thompson 2016). Some measures of social cohesion are inversely correlated with economic inequality (Putnam 2000).

There are a multitude of explanations for these persistent inequalities. The most commonly cited explanations for racial inequality are straightforwardly *structural* – they pinpoint the causes of inequality as the aftermath, whether intentional or unintentional, of historically discriminatory federal, state, and local policy. The Black-white racial income and wealth gaps are rooted in a long legacy of institutionally sponsored

---

<sup>1</sup> The Gini coefficient measures income inequality on a scale of 0 to 1. 0 indicates a perfectly equal distribution of income, while 1 indicates a perfectly unequal distribution. The Gini coefficient is 0.434 in the U.S. compared to 0.392 in the U.K., which ranks second highest for income inequality in the OECD (Horowitz et al 2020).

discrimination. Slavery and its aftermath prevented Black families from amassing intergenerational wealth, and Jim Crow prevented Black Americans from enrolling in white educational institutions and taking well-paying jobs. Some unions, which were routes to well-paying skilled industrial jobs, prohibited Blacks from membership (Hill 1996). The United States Department of Agriculture (USDA) has systematically denied or ignored loan applications from Black farmers, resulting in many foreclosures on Black-owned farms. In addition to the federal government, state and local governments also bear significant responsibility for racial inequality. Both the government and market actors have worked in tandem to perpetuate inequality by employing methods such as: redlining, the denial of home loans to homeowners in Black neighborhoods; racial steering, whereby real estate agents deliberately “steer” Black prospective homebuyers to black neighborhoods; and the “reverse redlining” targeting of Black and Hispanic families with subprime mortgage loans (Rothstein 2017). Because of the way that public schools are zoned and paid for by property taxes, educational segregation follows on the heels of residential segregation, resulting in functionally separate and unequal public schools (Rothstein 2015).

On the structural picture, persistent inequality results from unjust policies. Much attention has rightly been paid to the structural picture, which draws attention to the

causes of inequality in the realm of public institutions such as markets and government. Yet there are also causes of inequality that lie in realm of private life: the domain of individual autonomy, associations, and family life, into which the state may not intrude. While the division between public and private is notoriously ill-defined, it seems uncontroversial that if anything belongs in the private realm, our personal and intimate associations do. But our private lives have public consequences. Intimate decisions, such as whom we decide to live among, befriend, and form families, with also perpetuate inequality. While these private causes may be less central to the story of persistent inequality, they are no less normatively important. As I will argue, they exemplify liberalism's unresolved tension between associational freedom and social justice.

The aim of this dissertation is to turn our attention to these private, associational causes of inequality. My central argument in this chapter is that our associations perpetuate injustice, and that this fact demands a unique normative analysis. The argument consists of an empirical claim and a normative claim. The empirical claim is that our tendency to associate with people like ourselves, particularly along the dimensions of race and class, perpetuates both social and economic inequalities. I go into detail about the causal role of associations in perpetuating inequality in the next section. I must emphasize that this claim is not an alternative to the well-documented structural

story laid out above; as a normative philosopher, my job is to synthesize social science to understand the empirical complexities of an issue, not to challenge empirical findings. Our associational tendencies should be understood as an amplifying and maintaining cause of these preexisting inequalities, which overlaps in important ways with straightforwardly structural causes. My argument is broad: I am not concerned solely with racial inequality or class inequality, but with *any* form of arbitrary group inequality. Race and class just happen to be the socially salient categories that structure inequality in our world.

The second component of my argument broadly consists of a normative claim: insofar as sorting perpetuates injustice, we ought to do something about it. But what that something *is* is not obviously within the scope of state intervention. As I will argue, the fundamentally associational nature of the problem of sorting creates difficulties for state intervention. Let's begin by defining some terms and the scope of inquiry.

### **1.2.2 Homophily, *aka* sorting**

The terms "segregation" or "sorting" just refer to the separation of groups. We are segregated, sorted, or separated along the lines of race, sex, class, interest in music, political ideology, and a plethora of other social categories. Segregation can be coercively imposed, as with Jim Crow, or chosen, as a glimpse into any public high

school cafeteria makes obvious. This project is not concerned with the kind of segregation caused by discriminatory legal regimes that force people into separate groups or physical spaces based on some social characteristic. It is clear that that kind of segregation is unjust, as Elizabeth Anderson (2010) has argued at length. Instead, this project is concerned with the kind of segregation that arises from seemingly uncoerced, individual associational choices that are robust across domains. While these choices might be influenced by unjust policies and institutions, historical and current, such policies are insufficient to explain the extent of sorting. Following economists, I will refer to this kind of segregation as “sorting”.<sup>2</sup> Whatever the underlying mechanism, the kind of sorting I am concerned with is the kind that takes place in private life, emerges from an interaction between individual choices and structural constraints, and is not imposed from the top-down by institutions. Whether it is motivated by racial bias, purely positive in-group preference, instrumental economic choices, or more likely, some mix of all the above, sorting perpetuates inequality by concentrating opportunity within certain social groups to the exclusion of others.

---

<sup>2</sup> In Charles Tiebout’s seminal sorting model (1956), for instance, residential migration patterns are explained by competition among municipalities to provide bundles of public goods and services and “shopping” among individuals for the bundle that maximizes their utility.

Sociologists study this kind of sorting under the banner of “homophily”. Literally meaning “love of the same”, homophily describes the propensity to interact with individuals with similar characteristics at a rate higher than what a baseline model of random assortment would predict (McPherson et al 2001). Despite its name, homophily should not be taken to imply that this propensity has its roots in preference alone. Homophily can either be the result of preferences or the result of “the homogeneity of structural opportunities for interaction” – the fact that our social environments already tend to be homogenous, so we are simply more likely to encounter people who are similar to us (Kossinets and Watts 2009, 407). More likely, both preference and structural opportunities are at work (p.409).

The dynamics of homophily are best illustrated in the realm of intimate association: the choice of a partner. The realm of intimate association is as private a realm as there is. It seems to be the closest thing we have to a domain in which individual choice is directed *purely* by preference. Notably, these preferences are patterned. People overwhelmingly tend to marry those who are like them, especially along the lines of race, education, and income. Sociologists call this phenomenon “assortative mating”. Positive assortative mating describes the phenomenon of similar people marrying each other, while negative assortative mating describes the avoidance

of marrying people similar to oneself. In a review of the assortative mating literature, sociologists Christine Schwartz and Robert D. Mare (2005) find that the percentage of educationally homogamous pairs – couples with similar levels of educational attainment – increased from 45 percent in 1960 to 55 percent in 2003. This tendency is driven by the decreasing likelihood that individuals at both ends of the distribution will intermarry: those with the lowest and highest levels of education are least likely to marry outside their education level, and less so than in previous decades. Educational homogamy has increased despite the countervailing tendency among those in the middle of the income distribution to marry those with different levels of education. Schwartz and Mare take these results to be “consistent with a growing social divide between those with very low levels of education and those with more education in the United States” (p.621).

People also tend to marry primarily within their race, though multiracial marriages are becoming more common in the United States (Frey 2014). Because whites and Asians have higher educational attainment rates than Blacks and Latinos (Pew 2016), we might expect that racial homogamy interacts with educational homogamy to perpetuate racial income inequality. Moreover, multiracial marriages do not take place in equal degrees among racial groups. Blacks, and Black women in particular, are the racial group least likely to marry out of their race (Reeves & Rodrigue 2015). Given racist

attitudes against Blacks, the tendency of women not to marry men without jobs, and high incarceration and unemployment rates among Black men, Cohen and Pepin (2018) finds that Black women face more difficult marriage markets than white women. When Black college-educated women do marry, they tend to marry men less educated than they are.

According to Schwartz (2013), assortative mating has negative consequences for the distribution of resources:

Because spouses (and to a lesser extent cohabiting couples) share resources, assortative mating affects individuals' access to resources and the distribution of resources across families. Moreover, to the extent that parents influence their children's characteristics, assortative mating shapes the characteristics of future populations. Boundaries between social groups are maintained through assortative mating and weakened through intermarriage (p.452).

We will return to this key point later in Section 1.5.

Despite the seemingly preference-based nature of intimate associations, we cannot infer from the fact that people marry or befriend those like themselves that they have homophilous *preferences*, defined attitudinally rather than behaviorally. Preferences and structures interact in complex ways to produce patterns of homophily. The behavior might be better explained by the fact that physical environments comprise our social search environments, and following the "principle of least effort", it simply takes less effort to connect with people who are closer in physical space (Zipf 1949). Our



neighborhoods and cities are highly segregated along the lines of race and increasingly, along the lines of income, which may lead to unequal social outcomes (Reardon & Bischoff 2011). As such, “geography is the physical substrate on which homophily is built” (McPherson et al 2001, 431).

Residential segregation is not the emergent result of pure preferences (cf. Schelling 1976, Tiebout 1956). It is in large part the result of policies on the federal, state, and local levels that have officially endorsed and shaped exclusionary attitudes and preferences. Until the Fair Housing Act was passed in 1968, the Federal Housing Authority financed suburban developments that were closed to Blacks. Bank regulators knowingly approved “redlining” policies, by which banks denied loans to Black families in white suburbs and to Black neighborhoods (Rothstein 2015). Concurrently, “white flight” occurred (and continues to occur) because whites were reluctant to share public goods such as parks and schools with Blacks (Kruse 2013, Trounstein 2018). Segregation is now entrenched by exclusionary zoning, which consists of regulations that prevent dense and therefore cheaper housing from being built in certain neighborhoods. Exclusionary zoning, while facially neutral, prevents Blacks and low-income people from moving into middle-class neighborhoods (Silver 1997).

Once segregation is entrenched, it shapes people's information sets and attitudes. As Maria Krysan and Kyle Crowder (2017) show, when looking for a place to live, we are not omniscient actors who act with full knowledge of various residential options. Rather, we use heuristics such as "I wouldn't feel welcome there" or "That area's not safe" to make first cuts, then select among the options that survive the first cuts. These heuristics are based on the information that we glean from our social ties, such as family and friends. Whites tend to rule out minority-majority neighborhoods as a rule, whereas Blacks or Hispanics give all neighborhoods more consideration – but still end up selecting neighborhoods on the basis of social ties and their perception of belonging. Thus, explanations that rely solely on discrimination or on affordability ignore the role of segregated networks *themselves* in perpetuating residential segregation. Sometimes, the transmission of information within homogenous social networks is sufficient for segregation to beget segregation.

In sum, sorting has several features. First, the relatively abstract, network-based study of homophily and the concrete reality of segregation are inextricably intertwined; each influences the other in a complex feedback loop. To the extent that normative analysis makes causal assumptions, it is doubtful that we can provide a meaningful normative evaluation of one kind of segregation without bringing in the causal role of

the other. Second, the homogeneity of a social domain perpetuates itself through the transmission of information through social channels, regardless of the way in which that homogeneity first arose. Third, whatever sorting's causal origins, it ends up concentrating opportunity in some groups rather than others. Informal social relations are mechanisms of inclusion and exclusion that maintain group inequalities. Whether these relations emerge from innocuous preference or are constrained by the legacy of institutional discrimination is orthogonal to this point, and at any rate, the causes are empirically difficult to tease apart. The high-level concept of sorting intentionally abstracts away from the multifarious, complex causes of different types of social segregation. Instead, it directs our attention to the ways in which racially and socioeconomically homogeneous social clusters retain opportunity within some groups to the exclusion of others. Sociologist Charles Tilly (1998) calls this phenomenon *opportunity hoarding* (p.191).

While the term "hoarding" implies intentional exclusion, this need not be the case, nor is intention central to the problem. Just as an individual hoarder can unintentionally accumulate many physical items over time, so too can opportunity accumulate within social groups. To be sure, intentional opportunity hoarding is one aspect of the problem, as evidenced by residential segregation. But opportunity can be

hoarded without the intention to do so. Some people live in nice neighborhoods simply because they like the public goods that these neighborhoods provide.

### **1.2.3 Egalitarian ideals**

Sorting is a problem because it maintains group inequalities, thereby undermining several egalitarian ideals: formal equality of opportunity, Rawlsian fair equality of opportunity, and luck egalitarianism.

#### **1.2.3.1 Formal equality of opportunity**

In a slogan, formal equality of opportunity is the idea that “careers should be open to talents” (Rawls 1971, 72). Anyone who is qualified for a position should have a fair shot at attaining that position, meaning that the selection process should not be tainted by discrimination on the basis of morally arbitrary factors such as race, ethnicity, gender, or membership in other social categories. For example, to live up to the ideal of formal equality of opportunity, the Marines should accept, without regard to gender, any qualified individual who can pass their grueling physical requirements.

Formal equality of opportunity is an improvement over a caste system that selects people for social positions based on their group membership. Plato’s *Republic*, which assigned people to the roles of philosopher-king, warrior, or the brute majority

based on the makeup of their souls, is one such caste system. Formal equality of opportunity can be thought of as a pure meritocracy, in which the most qualified persons attain coveted positions. While it is not against hierarchy, and might even be construed as supportive of hierarchy, it gains its intuitive moral appeal from its commitment to non-discrimination on the basis of morally arbitrary characteristics.

However, formal equality of opportunity is widely criticized as insufficient as an ideal of equality. In Bernard Williams' (1962) re-imagining of the *Republic*, the warrior-guardian class is now open to all applicants, not just those born with silver in their soul. Anyone who can pass a rigorous battery of physical and mental tests can become a warrior-guardian. Because the position of warrior-guardian is prestigious, people from a broad swathe of society try out to become warrior-guardians. Yet, the overwhelming majority of people who succeed turn out to be those with silver in their souls. Is this because the silver-class are inherently better suited to becoming warrior-guardians? No. The persistence of the de facto caste system is better explained by developmental factors. Only those born into the former warrior class have access to the nutrition, schooling, and military training needed to develop into the kind of person that can pass the tests. Individuals born in other castes lack these resources, so it is no surprise that they cannot pass the test. However, from the perspective of formal equality of opportunity, this

result indicates no moral flaw in the system; a de facto caste system is permissible so long as that caste is justified by merit.

Williams argues that the so-called meritocratic Republic shows that formal equality of opportunity is far too weak an ideal. It only considers a “time-slice” view of opportunity – that at the moment of selection – but ignores the factors that shape people’s abilities up to that point. Recent critiques of elite universities, based on data that show that they overwhelmingly draw students from the top socioeconomic stratum of society<sup>3</sup>, parallel this old argument (Markovits 2019). A pure meritocracy is not a sufficient ideal of justice for society as a whole, even if it makes sense for individual institutions to select the best qualified applicants without regard to their developmental advantages or disadvantages.

As the tale of the meritocratic Republic and its recent iteration show, formal equality of opportunity is compatible with a caste system that is maintained by sorting. If the former warrior class decides only to socialize among, form families with, and circulate opportunities among themselves, then it will be difficult for resources such as

---

<sup>3</sup> A New York Times analysis shows that at Duke, 19.2 percent of students come from families in just the top 1 percent of the income distribution (making >\$630k a year). In comparison, only 16.5 percent of students hail from families in the bottom 60 percent (making <\$65k a year) (Aisch et al 2017).

nutrition, schooling, and military training to spread to the other classes. Thus, although sorting does not *directly* undermine formal equality of opportunity, it shows that formal equality of opportunity is an insufficient ideal. It is instituted in response to the evils of a caste system, but it effectively replicates that caste system under a different banner, and worse still, provides a justification for it. In that sense, formal equality of opportunity might be even more insidious than an explicit caste system.

For these reasons, egalitarians have moved toward more substantive ideals of equal opportunity, such as Rawls's fair equality of opportunity. But sorting also compromises fair equality of opportunity, for our social networks and personal relationships play a large role in shaping our developmental opportunities.

### **1.2.3.2 Fair equality of opportunity**

In response to the inadequacy of formal equality of opportunity, Rawls offers fair equality of opportunity (FEO). FEO is the idea that individuals with the same talent and ambition should have the same prospects of success in competitions that allocate positions of advantage. FEO requires that individuals have equal *developmental* opportunities; equality of opportunity does not just consist in the fair contest. Rawls writes:

The thought here is that positions are to be not only open in a formal sense, but that all should have a fair chance to attain them. Offhand it is not clear what is meant, but we might say that those with similar abilities and skills should have similar life chances. More specifically, assuming that there is a distribution of natural assets, those who are at the same level of talent and ability, and have the same willingness to use them, should have the same prospects of success regardless of their initial place in the social system. In all sectors of society there should be roughly equal prospects of culture and achievement for everyone similarly motivated and endowed. The expectations of those with the same abilities and aspirations should not be affected by their social class (Rawls 1971, 63).

There is much debate over several aspects of Rawls's FEO. Is there such a thing as a "distribution of natural assets"? Does merit really consist in natural talent and effort? Natural talent, insofar as it is innate, seems just as morally arbitrary as social characteristics such as race and gender – so why is it legitimate for the former, but not the latter, to influence one's life prospects? And is there even such a thing as natural talent that is separable in principle from developmental opportunities and effort (cf. Joseph Fishkin 2014, Ch. 2)? I set aside these important questions, which the luck egalitarian attempts to grapple with. Also, note that FEO is not meant to be a stand-alone principle; it plays one role in a broader scheme of distributive justice, which includes the difference principle to remedy any inequalities that might arise from the distribution of natural assets (Rawls 1971, 83). In spite of these qualifications, FEO has value in and of itself: it embodies the moral value of not having one's life path determined by morally arbitrary factors. As Rawls writes:



[I]f some places were not open on a fair basis to all, those kept out...would be justified in their complaint not only because they were excluded from certain external rewards of office but because they were debarred from experiencing the realization of self which comes from a skillful and devoted exercise of social duties. They would be deprived of one of the main forms of human good (Rawls 1971, 73).

What inequalities is FEO concerned with? FEO seeks to equalize, within the limits of liberty, the social primary goods. These goods are all-purpose means that are necessary for citizens to exercise the two moral powers: a conception of the good and a sense of justice. They include things such as the freedom of association, income and wealth, and the social bases of self-respect (Rawls 2001, 58-59).

To this list, Chiara Cordelli (2015) adds *relational resources* – “goods that are *distinctively produced* through and available within relationships or that are themselves *constitutive* of certain relationships” (p.90). For example, relational resources include trust and social influence (p.95-97). They also include social skills that help to secure stable relationships with others. Cordelli argues that “it is unfair if some people get more opportunities to access relationships than others for reasons that are arbitrary” (p.106). Because relational resources are a primary social good, they should be subject to the principle of FEO as well.

Sorting undermines FEO by concentrating both economic and social opportunities within some groups. Assortative mating and the intergenerational transmission of opportunity widen economic as well as childhood developmental gaps

between classes. Homophilous social networks and the transmission of opportunity within such networks exclude outsiders from accessing economic opportunities. And the geographical concentration of opportunity creates homogenous search environments, which in turn replicate assortative mating and homophilous social networks. Sorting thereby undermines FEO by silo-ing relational resources into some groups rather than others. For example, friendship segregation (Moody 1971) undermines Cordelli's fair distribution of relational resources by making it more difficult to cultivate cross-group relational talents, and by undermining minority children's self-confidence if they are continually excluded from majority cliques (Cordelli 2015, 102-103). As such, it affects the "expectations of those with the same abilities and aspirations" by their "social class" (Rawls 1971, 63).

#### **1.2.3.4 Luck egalitarianism**

Rawls and luck egalitarians both agree that there does not seem to be a principled distinction between demographic characteristics and endowments of talent; both are equally arbitrary. Where they differ is on the legitimacy of natural talent in shaping one's life prospects. While Rawls thinks that natural talents may legitimately shape one's life prospects as long as the resulting inequalities are remedied by the difference principle, luck egalitarians such as Ronald Dworkin disagree. Luck

egalitarians think that inequalities that result from the natural lottery are on a par with inequalities that result from unjust, arbitrary social factors. The only factor that may legitimately shape life chances is responsible choice.

Dworkin's (1981) theory of equality of resources distinguishes between "brute luck" and "option luck". Brute luck consists of circumstances that are out of one's control, such as circumstances of birth and native endowments such as talent. Native endowments need not be wholly genetic; they simply need to be more or less "fixed" in the individual after a certain point in the developmental trajectory due to factors out of her control, such as nutrition, upbringing, and the environment. Dworkin would compensate for these unequal native endowments, which are due to brute luck from the perspective of the individual (p.45). By contrast, option luck is the product of choice. We should not allow brute luck to shape our life chances, but we should allow option luck to do so. This distinction respects human choice and responsibility, while compensating for morally arbitrary factors. However, it is not easy to draw a principled line between brute and option luck, as Dworkin acknowledges. After all, brute luck (such as "natural" talent) shapes the kinds of preferences we develop and the life paths we choose to embark on.

Sorting expands the impact of brute luck factors on one's life prospects. One clearly cannot choose the family or social group into which one is born, and the effect of these circumstances of birth on one's life prospects are shaped by the social processes by which families are formed. Because of sorting, families tend to be racially and/or economically homogenous, and this homogeneity amplifies the effects of birth on one's life prospects. Rich families are richer than they would be in an unsorted world, and poor families are poorer. And because people sort geographically according to race and class, the geographic concentration of opportunity is amplified by these antecedent forms of social sorting. These developmental factors then influence the kinds of adulthood social relationships that one is in a position to form. This occurs in part by directing individuals to certain search environments such as schools and workplaces, and consequently, the kinds of opportunities that one has access to.

At this point, one might object that at least some of our associations are chosen after a certain point, and therefore belong in the realm of option luck. We cannot choose our family, but we can and usually do choose our friends, neighbors, and partners. Since these associations are chosen, it is legitimate for them to influence our life prospects. Thus, it cannot be said that sorting expands the impact of brute luck on one's life prospects.

A few things can be said in response. First, we shouldn't underestimate the extent to which upbringing and culture, both of which are unchosen, influence these later associations. This is not to say that upbringing or culture determine the associations we make in adulthood, but only that these unchosen factors exert a large degree of influence on apparently chosen factors. Someone who grew up in pre-gentrified East Durham will likely have a different perspective, norms, values, and ideology than someone who grew up in Trinity Park. It is simply more difficult for people to connect across such intangible barriers, even if they find themselves in the same search environment (suppose they both end up as Duke undergraduates).

Second, our associations don't seem to be the kind of thing that we should choose mainly or solely because of their influence on our life prospects. We should choose our friends and partners because we enjoy their companionship and value who they are. To be sure, one does have some responsibility to shape one's social environment. The force of the advice that children receive as caution to stay away from bad influences – "you are who your friends are" – nods to this responsibility. But it does not follow that we should choose associates, especially those in the personal realm, for the opportunities that they can provide to us. Doing so may instrumentalize these associations to a demeaning degree – hence the pejorative force of terms such as "gold

digger” and “schmooser”. These cultural, pseudo-moral prohibitions on certain motivations for association find precedence in Aristotle’s philosophical discussions on friendship. The friendship of utility (and pleasure) is deficient to the friendship of virtue (*Nicomachean Ethics*, Book VIII).

### **1.2.3.5 Takeaways**

There are several takeaway points from this discussion. First, equal opportunity is important not just for reasons of distributive justice, but it has moral value in and of itself. Equality of opportunity, broadly understood, is an attractive ideal because it attempts to eliminate the effects of morally arbitrary factors, such as the circumstances of one’s birth or other brute-luck factors, on the trajectory of one’s life. Second, however fuzzy the distinction between responsible choice and arbitrary circumstances are, circumstances of birth and social class clearly fall on the latter side of the line.

The problem of sorting, then, is this: Sorting undermines equality of opportunity by functioning as a mechanism of opportunity hoarding. Inequalities that influence individuals’ life prospects persist in part because sorting concentrates opportunities within certain social groups. These opportunities are economic, developmental, and social. But because sorting is the aggregate result of informal social relations, and these social relations are considered to lie within the realm of personal life, it is not clear what

can be done about sorting without infringing on personal liberty. I will address this aspect of the problem in later chapters.

### **1.3 The scope of sorting**

As the discussion above shows, the problem of sorting is intimately tied up with the problem of the family, which is thought to be the most difficult obstacle to the realization of equality of opportunity. In spite of these causal connections, they are fundamentally different and equally deep problems from a normative standpoint. Sorting therefore widens the scope of the private-sphere obstacles to egalitarian ideals, so that the prospect of realizing these ideals seems even more far-fetched than previously thought.

#### **1.3.1 The problem of the family**

Rawls acknowledges that the family places a limit on the realization of FEO:

the principle of fair opportunity can be only imperfectly carried out, at least as long as the institution of the family exists. The extent to which natural capacities develop and reach fruition is affected by all kinds of social conditions and class attitudes. Even the willingness to make an effort, to try, and so to be deserving in the ordinary sense is itself dependent upon happy family and social circumstances (Rawls 1971, 64).

Notice that Rawls runs together the institution of the family and “social conditions and class attitudes” – a point that I will return to shortly. The essence of the problem of the

family is this: parents transmit advantage or disadvantage to their children, thereby shaping children's life outcomes to a significant degree. James Fishkin (1983) fleshes out the problem of the family in the form of a trilemma: out of the principle of merit (the principle of the fair contest), equality of opportunity (the principle of fair life chances), and the autonomy of the family, we can only pick two. Maximizing all three simultaneously requires sacrificing one ideal. Joseph Fishkin (2014) argues that "In theory, we can imagine a society that eliminated the institution of the family – grossly violating the autonomy of parents – and therefore achieved both the principle of the fair contest and the principle of fair life chances" (p. 53). Both Fishkins' arguments assume that the problem of the family is the deepest problem for equality of opportunity. I think this assumption is mistaken: the problem of sorting is just as deep as the problem of the family, and it widens the scope of the challenge to equality of opportunity.

The problem of the family has both a normative and descriptive component. Descriptively, parents transmit advantage and disadvantage to their children simply in virtue of what they do (or do not do) for their children. Parents might cultivate their children's abilities by providing a supportive home environment or sending them to good schools, or they might stymie their children's growth by being abusive, neglectful, or simply lacking the resources to help their children thrive.



Many people also think that parents are morally obligated to do “what’s best” for their children, even if this gives them an unfair advantage over less fortunate children. This apparent normative obligation ramps up the problem of the family. In recent years, attention has been drawn to the role that upper-middle-class parents in diverse metropolitan areas play in segregating school systems (Williamson et al 2020). Parents with means routinely send their children to private tutoring or test-prep courses that help them gain a leg up over less advantaged children. While such behavior directly undermines equality of opportunity, some parents see this as morally required. This common view has philosophical backers: Thomas Nagel (1995) argues that parental partiality is both permissible and desirable, and that it would “make unacceptable demands on individuals” (p.5) to demand that parents act from impartial motives only. Partiality is thus a legitimate constraint on equality. Thus, the family is a problem for equality of opportunity because one’s circumstances of birth are morally arbitrary, yet they play a significant role in determining one’s life prospects. However, it is neither permissible nor desirable for the state to intervene in the private domain of family life for the sake of equalizing opportunity.

Where do we draw the line between permissible and impermissible actions that parents may take to fulfill their normative obligation? It is clearly impermissible for

parents to flout institutional rules, such as criteria for admission, in order to help their children gain advantage. Doing so violates both formal and fair equality of opportunity. In the 2019 “Operation Varsity Blues” case, rich and famous parents bribed officials at elite universities to admit their children under the false pretenses of college athletics. These parents were justly punished by federal law enforcement for circumventing the admissions process. Specifically, the underlying ethical principle that has been violated seems to be something like this: college admissions processes are supposed to be fair contests through which only qualified applicants should pass without unfair advantages, such as advantages given to them by their parents.<sup>4</sup>

But ordinary families act in ways that, while well within the boundaries of institutional norms, nevertheless confer advantage onto their children. Many wealthy and upper-middle-class parents pay for exclusive college admissions counselors who will help their children craft a perfect admissions package. Even parents with fairly modest means routinely send their children to enrichment activities, such as sports coaching or music lessons. And parents expend significant resources to place their

---

<sup>4</sup> It is worth noting that the principle of the fair contest does not capture what goes on in admissions, as students get admitted for all kinds of different, incommensurable things – athletic talent, musical talent, leadership ability, diversity, legacy connections, etc.

children in social environments that are advantageous in part because they are homogenous. Many parents move to the suburbs so that their children will be able to attend good public schools; still others send their children to exclusive private schools. These schools not only provide students with a good education, but can give them valuable social networks and inculcate in them certain social habits and mannerisms that will prove advantageous later in life.

Harry Brighouse and Adam Swift (2009) argue that even these innocuous behaviors are morally impermissible. They distinguish between familial “relationship goods” that are secured by partial parental behavior, and those behaviors that help children gain a competitive advantage without facilitating the attainment of relationship goods. For example, reading bedtime stories to one’s child facilitates intimacy between parent and child and the development of the child’s autonomy, so these are morally permissible activities. By contrast, activities that are not essential for securing relationship goods includes ones such as sending one’s child to a private school, investing in a trust fund for her, or engaging in nepotism in order to help her secure a competitive advantage. Accordingly, they argue that these activities are not legitimate expressions of parental partiality (p.56-57).

But drawing the line between permissible and impermissible partial behavior leads to the following objection: “a loving parent will be motivated quite generally to further her children’s interests, to care that her child’s life go better than it otherwise might. If so, then it looks as if any action by a parent to benefit her child realizes a familial relationship good, which might seem fatal to our project” (p.60). In response, Brighthouse and Swift say that:

although distinctively familial relationship goods can indeed be realized when parents act in ways intended generally to help their children’s lives go better than they otherwise would, *those* goods are *simply not weighty enough* [italics mine] to warrant the state’s being required to permit parents to engage in those interactions or transmissions if they conflict with other children’s interest in fair equality of opportunity (p.62).

This response is weak. Brighthouse and Swift do not offer independent grounds to evaluate the weightiness of familial relationship goods realized in ways that further a child’s advantage. Indeed it seems that they draw the line between weighty and non-weighty goods in such a way as to rule out goods that are antecedently deemed illegitimate. Moreover, the line between relationship goods and positional goods does not neatly track the transmission of advantage. Even what Brighthouse and Swift consider legitimate expressions of partiality, such as giving them sufficient nutrition and outdoor playtime, reading them bedtime stories, or even engaging them in conversation, give some children a competitive edge over others.

The upshot is that the family poses a deep problem for equality of opportunity. Parents transmit advantage or disadvantage to their children simply in virtue of being their parents, regardless of whether they intend to do so. Thus, the problem of the family poses a fundamental dilemma between the autonomy of the family and the imperatives of justice. It does not seem possible to simultaneously allow the autonomy of the family and realize any meaningful conception of equality of opportunity.

### **1.3.2 The problem of sorting is distinct from the problem of the family**

Because the problem of the family is so deep and so intertwined with other social circumstances, one might conceptually run the problem of sorting and the problem of the family together, as Rawls and both Fishkins seem to. On this line of thinking, assortative mating and the geographical sorting of families into high- and low-income areas are troubling only because they supercharge the intergenerational transmission of advantage. For example, two highly-educated, high-income parents will have more means to cultivate their child's development and help him secure positional goods than two parents with high school degrees who struggle to make ends meet. And through the normal course of raising their child, these privileged parents will pass on developmental advantages to him. As Joseph Fishkin writes:

Parents can choose to live in good neighborhoods that offer advantages in safety, special resources, and social opportunities that derive from interactions with privileged peers

and *their* families, thereby connecting children to networks that offer a wealth of opportunities. (Fishkin 2014, 49)

Sorting and the problem of the family clearly exacerbate each other in important ways. For example, for the sake of their children, parents are often motivated to move to relatively segregated areas with high-quality public goods such as schools. However, the normative focus on the problem of the family compared to the problem of sorting is misplaced. It is misplaced for two reasons: there is no reason to privilege the family as a form of private association, and intergenerational, familial transmission is not the only mechanism by which morally arbitrary advantages percolate through social networks. There are purely social mechanisms by which arbitrary advantage is transmitted.

First, consider the privileging of the family as a private association. As both Fishkins and others have argued, the autonomy of the family should not be infringed on. But neither should the autonomy of other associations. Other forms of private association, such as friendships, non-familial intimate relationships, clubs, and neighborhood associations are also private associations, and private associations are all protected by the freedom of association – “the right to choose the society most acceptable to us”, in John Stuart Mill’s canonical formulation (1859, Ch.1). This right is affirmed by Rawls in *Political Liberalism*. Rawls writes: “freedom of association is required to give effect to liberty of conscience; for unless we are at liberty to associate

with other like-minded citizens, the exercise of liberty of conscience is denied. These two basic liberties go in tandem" (Rawls 1993, 313). Rawls also says that the lexical priority of the basic liberties can never be denied "for reasons of public good or perfectionist values", including equality of opportunity (Rawls 1971, 74).

While the family is commonly thought of as the "nuclear" unit of society for contingent sociohistorical and biological reasons, there are no unique, non-associational liberal principles under which the family ought to be politically privileged. Such privileging might even violate liberal neutrality by prioritizing a certain conception of the good, thereby discriminating against other forms of caring relationships such as friendship (Brake 2012). If other forms of private association are no less private than the family from a liberal perspective, then the fact that sorting encompasses various non-familial associations widens the challenge to equality of opportunity. A visual analogy might help to illustrate the point. We can think of the family as a core association into which the state may not intrude, however much it undermines equality of opportunity. Surrounding that core is a penumbra of other, arguably less fundamental associations into which the state equally may not intrude, however much these associations undermine equality of opportunity. Thus the autonomy of associational life *in general*, not just the autonomy of the family, is the central challenge for equality of opportunity.

Nevertheless, one might think that the problem of the family is the deepest problem in practice because one's circumstances of birth have such profound effects on one's life trajectory, and that these effects are unmatched by any other kind of social circumstances. But this point is compatible with the problem of sorting. Sorting functions as a mechanism of opportunity hoarding at various points throughout the trajectory of one's life, by routes that are not directly related to the familial transmission of advantage: for example, by network effects and the geographic concentration of opportunity. Most individuals get their jobs through social contacts, especially through "weak ties" at the boundaries of one's social network (Granovetter 1973) – peripheral relationships such as "friend of a friend". When social networks are homophilous along socioeconomic dimensions, job opportunities will tend to be transmitted to those who are similar; upper-middle-class WASPs will put forth other WASPs for job referrals, thereby hoarding opportunity among themselves. By contrast, outsiders who are just as qualified, but who lack personal contacts, will find it more difficult to get jobs through these informal routes. While personal networks might seem like an illegitimate route by which to obtain jobs, such nepotism can be thought of as analogous to the Varsity Blues' case of parental partiality gone wrong. Social networks, like legitimate parental



partiality, hoard opportunity through some innocuous mechanisms such as information-sharing. The issue is that many social networks are both homogenous and closed.

In sum, sorting widens the scope of the private-sphere challenge problem for equality of opportunity. It undermines equality of opportunity via several pathways that are distinct from the familial transmission of advantage, so it broadens the protective sphere around which the concerns of justice may not intrude.

The fact that sorting tends to get conflated with the problem of the family highlights a puzzling blind spot in liberal political theory's conception of the trajectory of opportunity. We focus on circumstances of birth because these are most obviously out of an individual's control and set us on various life paths. Yet the presence or absence of opportunity is not confined to an early time-slice of a person's life, but occurs throughout our lives. There is no "moment of equal opportunity" (Chambers 2009). We have the potential to encounter or miss out on opportunities throughout our lives, including in adulthood, in large part because of the people that we associate with and live among. This can be due to our circumstances of birth, but just as well may be due to our later choices. If equality of opportunity is the goal, there does not seem to be a principled reason to ignore all factors that shape access to opportunity aside from the

earliest ones. From the perspective of liberalism, the family is but one association nested among others.

### **1.4 The depth of sorting**

The problem of sorting has only been recognized to a limited extent by philosophers, such as Elizabeth Anderson (2010) and Erin Kelly (2017). These philosophers recognize that associational life can perpetuate injustice (specifically, racial injustice). However, their diagnoses of the problem focus on stigma. I argue that this framing does not go deep enough and misleads us on the appropriate solutions. Sorting can and does arise from a variety of mechanisms that are unrelated to stigma. I present six models of sorting to show this.

#### **1.4.1 The inadequacy of “bypassing” solutions**

What I call “bypassing” solutions are appealing in situations in which: 1) one level or domain of possible intervention is blocked off from intervention for normative or practical reasons; 2) another, connected level or domain is open to intervention, and 3) there is reason to believe that intervening on the second level or domain might have desirable effects on the first. Because the problem of sorting is fundamentally a problem of associational life – a private-sphere issue – it does not seem permissible for the liberal

state to intervene on sorting directly, so we might want to address it using a bypassing solution. Elizabeth Anderson and Erin Kelly both offer such solutions.

Anderson (2010) argues that racial segregation is unjust because it is both caused by an unjust process and it has unjust consequences. Firstly, it is unjust because it is the result of racial stigmatization. Segregation is in large part the result of whites holding negative stereotypes about Blacks; for example, implicit bias studies show that whites are quicker to associate Blacks with violence and criminality. Secondly, segregation has unjust consequences because it prevents Blacks from having equal access to opportunities. In particular, segregation results in the closure of social networks, preventing social capital from transmitting across racial boundaries. Anderson's diagnosis of segregation is similar to my diagnosis of sorting: it is fundamentally a social problem, not obviously remediable by economic redistribution.

Anderson puts forth the *integrationist hypothesis* as a solution to segregation. The integrationist hypothesis is based on Gordon Allport's "intergroup contact hypothesis" (1954), which holds that contact between different racial or ethnic groups in controlled, institutional settings can reduce prejudice. According to the integrationist hypothesis, to ameliorate segregation, we should promote integration in settings of institutionalized support for interracial cooperation. Affirmative action is one such way of promoting

integration in institutional settings such as schools and workplaces. Although integration will initially increase both positive and negative interracial interaction, in the long run, integration will reduce racial stigmatization. People will learn to better manage interracial relationships and choose integrated lives in the private realm of neighborhoods, friendship, and family. This is an explicitly bypassing solution: over time, integration in the public realm will result in integration in the private realm, which is off-limits from public policy.

Similarly, Kelly (2017) focuses on the effects of racial stigma in associational life. Kelly argues that racial stigma in the informal spheres of life raises the “problem of historical injustice” for liberalism and writes that “open access to positions could be obstructed by ongoing patterns of racial partiality and bias” (p. 84). Kelly acknowledges the associational dilemma at the heart of the problem of sorting:

The challenge is that informal social relations and networks, which are mechanisms of inclusion and exclusion, lie largely in the domain of personal choice and freedom - a domain that is and should be protected by law. An effective approach to countering historically rooted injustice threatens to run up against reasonable personal liberty. (p. 90)

Fair equality of opportunity, Kelly argues, would not automatically be realized by the implementation of the two race-neutral principles of justice – fair equality of opportunity and the difference principle. This is because “Racial identity, regardless of formal commitment to a color-blind society, might still be tied to stigma: in hiring and

promotion decisions, the development of intellectual potential, and the recognition of creative accomplishment, for instance" (p.84). Consequently, Kelly argues that an expanded principle of fair equality of opportunity should include a *principle of historical redress*. This principle would "check the tendency among class-conscious liberals to ignore the background culture dimensions of racial justice or to reduce the significance of culture to socioeconomic indicators" (p. 91). A principle of historical redress would require policies that directly address historical injustice, such as funding for school integration and affirmative action. Kelly's solution, which is pitched at the level of ideal theory, is compatible with Anderson's solution, which is pitched at the level of non-ideal theory.

Both Anderson and Kelly are leery of direct state intervention on associational freedom, which insulates social relations from state intervention whether or not they result from racial stigma or a more innocuous cause. Thus, their solutions attempt to bypass the issue of associational freedom. Instead, they focus on public-sphere solutions that may either have an indirect effect on associational life or compensate for the effects of associational life on social justice. Kelly's principle of historical redress intends to manage the problem of sorting, rather than solve it, whereas Anderson's integrationist hypothesis is a suggestion for directly solving the problem of sorting.

But the burden on each of these solutions is very high. To be sufficient solutions to the problem of sorting, each of these bypassing solutions (independently or jointly) needs to *completely* mitigate the effects of associational freedom on equality of opportunity. Assuming these solutions are fully implemented, but associational freedom continues to undermine equality of opportunity in some way, then there still remains a dilemma between freedom of association and equality of opportunity. A solution to the problem of sorting needs to address this dilemma head-on.

There are good reasons to think that these bypassing solutions are insufficient for remedying the injustices related to segregation. The central issue is that they focus on one mechanism by which sorting arises: eradicating or managing exclusion that results from stigma. However, sorting can and does arise even without stigma. It can arise in two major ways: from the result of pure homophilous preference, or as long as homophilous associations are effective means of attaining other goods. This is not to say that stigma plays no role in creating or maintaining sorting, only that it is not necessary, and that stigma cannot explain the full degree to which sorting occurs. We should not hinge our normative solutions on a single causal story.

## 1.4.2 The inescapable mathematics of segregation

As Thomas Schelling argued, we cannot infer prejudice or malintent from the fact of segregation. Schelling calls this the “inescapable mathematics of segregation” (1978): the outcome of segregation is compatible with multiple causal stories, each of which has its own normative valence. In addition to exclusion that is motivated by stigma, there are at least six possible ways in which sorting can (and does) arise from a distribution of preferences among the privileged and less-privileged. Let me clarify that privilege is a relative term; Alice is privileged relative to Bob insofar as she has more opportunities and resources available to her, regardless of how this disparity came about. The relative nature of that term is important for capturing the intricacies of opportunity hoarding. By “less privileged”, I do not mean *unprivileged*. For example, the top 20 to 2 percent (exclusive of the top 1 percent) of households in the income distribution are less privileged relative to the top 1 percent, but privileged compared to the bottom 80 percent.

Below I present six ways in which sorting can and does arise or persist: from the asymmetric homophily of the privileged, symmetric homophily, asymmetric homophily of the less privileged, and most disturbingly, from agents’ preference for being in the minority. Sorting can also arise because associations are often a means to attain

instrumental goods. Only the first of these models straightforwardly matches Anderson and Kelly's stigma model of exclusion. These are not just "how-possibly" stories (Bokulich 2014) that detail a pathway by which sorting can emerge, but to some extent, each is a "how-actually" story of how sorting does emerge in the real world.

#### **1.4.2.1 Asymmetric homophily of the privileged**

Imagine that the privileged have more homophilous preferences than the less-privileged. "Homophilous preferences" here do not just describe actual choices, but mental states such as desires. The privileged prefer to associate with people like themselves, whereas the less-privileged either prefer to associate with the privileged or are indifferent between groups. The privileged sort themselves into a closed network through which opportunities circulate, excluding the less-privileged from access to these opportunities. This model most closely matches Anderson and Kelly's stigma model of exclusion.

Consider "white flight" as one example of this model in real life. White flight occurs when the number of blacks in a given neighborhood reaches a "tipping point", inducing whites to flee to whiter pastures, and explained a significant portion of segregation in the 1910s and 1920s (Shertzer & Walsh 2019). We don't even need to assume that racial stigma is at work here, even though it surely is in the real world; mere



preference for whites is enough to generate white flight (Schelling 1971). Given that whites are more socioeconomically privileged than blacks, whites can and do hoard opportunity via their strong preference for homophily.

#### **1.4.2.2 Symmetric homophily**

A second way in which sorting can arise is via symmetric homophily: both the privileged and the less-privileged equally prefer to associate with their own kind. Schelling illustrates this in his minimal model of residential segregation (Schelling 1971). On this model, agents are divided into two groups, and each agent occupies a random spot on a checkerboard. All agents have a small preference to live next to at least two members of their own group, but no aversion to living with members of the other group. This sole behavioral condition governs their movement. Over time, agents sort themselves from a state of integration into an equilibrium state of total segregation. While the Schelling model is perhaps more of theoretical rather than of practical interest – it does not explain facts about real world residential segregation, which is complicated by institutional discrimination (Elliot-Graves & Weisberg 2014) – we can draw one important lesson from it. We cannot infer ill will or exclusionary intent from the fact of segregation, as segregation can arise from perfectly innocuous homophilous preferences.

But ill will or not, opportunity hoarding can arise from this innocuous situation, given an inequality of resources between two groups and little interaction between them.

Symmetric homophily does not fit Anderson and Kelly's model of stigma-based exclusion. Indeed, symmetric homophily does not seem unjust at all, but understandable and even desirable in some situations, especially for oppressed minorities. There doesn't seem to be anything inherently wrong with the fact that people prefer being around those with whom they share similar backgrounds. In fact, there may need to be a certain degree of homophily in order to maintain the group-differentiation that is necessary for diversity; otherwise, society may blend into a culturally homogenous mass. The challenge is how to cultivate diversity without veering into the toxic aspects of segregation (Muldoon 2018).

#### **1.4.2.3 Asymmetric homophily of the less-privileged**

A third way in which sorting can arise is through the asymmetric homophily of the less-privileged. Suppose the less-privileged prefer to associate only or mostly among themselves, whereas the privileged either prefer to associate with the less-privileged or have neutral tendencies. This situation is the mirror image of the asymmetric homophily of the privileged, and some degree of sorting will arise. In the real world, this phenomenon is known as "self-segregation" (Schelling 1978), where the prefix "self"

differentiates this form of segregation from other-imposed segregation. It is commonly used to refer to minorities who create closed-off communities of their own accord. This can create two forms of simultaneous opportunity hoarding. The privileged are able to hoard opportunity among themselves simply because the less-privileged do not want to associate with them. Conversely, the less-privileged may hoard their own forms of opportunities.

For example, many Black students attend historically black colleges and universities (HBCUs) such as Howard, Morehouse, and Spelman, rather than elite predominantly white institutions (PWIs). This is often motivated by strong cultural affinity (Freeman 1999). But PWIs remain the gateway to elite networks and opportunities, simply because elite networks and opportunities are dominated by whites. So, by attending HBCUs, Black students may inadvertently lock themselves out of opportunities.

The less-privileged may also engage in their own forms of opportunity hoarding. Consider ethnic communities such as Korean immigrants, who are economically successful in part because they circulate economic opportunities only within their ethnic network. Korean grocers, laundromats, beauty salons, and other small businesses thrive on their connections to other links in the supply chain that are owned by their

compatriots (Tilly 1998, Ch.5). This form of opportunity hoarding has attracted federal attention. In *EEOC v. Consolidated Services Systems*, the Equal Employment Opportunity Commission sued a Korean-owned janitorial services company for “illegally discriminating against non-Koreans as a class in both the recruitment and hiring of persons to perform janitorial services”. Or consider Orthodox Jews involved in the diamond trade in Williamsburg, Brooklyn, who also hire and deal with largely among themselves. Barak Richman (2006) argues that Orthodox Jews have historically dominated the diamond industry because they are able to rely on informal social mechanisms such reputation mechanisms and intergenerational legacies enforced by community institutions. Because of the nature of the diamonds – diamonds are valuable and easy to steal and transport – trading within an insular community allows merchants to diminish an otherwise high risk of theft or flight by relying on the trust that stems from close social relations. In both the Korean and Jewish cases, social relations effectively anchor business relations by providing a check on avarice and betrayal. But by favoring their own compatriots, these small-business owners may hoard opportunities from those in less-privileged groups – hence, for example, there has long existed a degree of animosity between Blacks and Asian businessowners in inner cities, which reached a flashpoint during the 1992 Los Angeles riots (Yi & Hoston 2020).

How should we evaluate these forms of opportunity hoarding? One approach is to see how they function against a broader social backdrop. Both Korean immigrants and Orthodox Jews face significant barriers to integration in mainstream society – their form of opportunity hoarding might be a response to being locked out of wider opportunities. In the HBCU case, cultural self-segregation can inadvertently lock students out of important social networks. The asymmetric homophily of the less-privileged can thus lead to both opportunity hoarding within less-privileged communities and the inadvertent exclusion of the less-privileged from networks within which opportunity circulates.

#### **1.4.2.4 Segregation that no one seeks**

The previous three models show that various combinations of homophilous preferences can result in sorting and opportunity hoarding. But strikingly, homophilous preferences are not even *necessary* for sorting to occur. In a series of models, Ryan Muldoon et al (2012) show that segregation can occur even when agents have “maximally heterogenous preferences” (MHP) – when they want to be surrounded by agents that look different from themselves, and even when they want to be in the smallest minority. This result can be explained by the interaction between information sets and interests. As agents of the same type move closer together in space, their

information sets (about the characteristics of their neighbors) become increasingly correlated. Because they have similar preferences and information sets, they will make similar movement decisions. The correlation between information sets and movement decisions leads to further information convergence, and thus increasingly correlated movement decisions over time. Thus, “similarity of information sets becomes its own attractor” (p.24).

MHP models are far too simple and minimal to actually explain real-world sorting. So what is their epistemic value? MHP models show that merely taking a social characteristic into consideration can lead to individuals to sort themselves into a segregated equilibrium. Individual preferences for diversity can have unintentional aggregate consequences due to the dynamics of interaction; as with the Schelling model, there is no direct mapping between individual preferences and collective outcome. Thus, contra Anderson and Kelly, we cannot conclude that an obvious, necessary, or sufficient solution to the problem of sorting is to erase stigma, or to inculcate preferences for diversity such as being in the minority.

#### **1.4.2.5 Means to attain instrumental goods**

Sorting might not result from any kind of group-based preference. It can arise or persist as long as associations are means for securing certain instrumental goods, and

associations are already socially sorted according to race or class. Associational choices are mediated by factors such as status-seeking (Frank 1985) and the pursuit of high-quality public goods – so again, we cannot read problematic preferences such as exclusionary desires off of associative choices.

Country clubs are a typical example of the status-seeking. Although these have fallen out of favor with younger generations, they have traditionally been bastions of privilege and prestige. These clubs have attractive amenities such as golf courses and swimming pools, but the composition of that club's membership is just as much a club good as the physical amenities. The club can give one access to high-prestige social contacts and networking opportunities (Janiak 2003). Thus, current members have an incentive to attract prospective members who would increase or at least maintain the social cachet of the membership pool. If people with social cachet are mostly rich, white men, then the club will have an incentive to attract such people, because such people can help members of the club attain instrumental goods. This will be the case regardless of whether people in the club *truly* prefer to associate with rich white men, so long as their membership in the club is driven by the desire to increase their social cachet.

The same dynamics are also at play in student clubs such as fraternities and sororities – students often join these not just for social reasons, but also for professional

networking purposes. But Greek organizations tend to be exclusive and more easily accessible to wealthier students. Students are coming to recognize that Greek life perpetuates inequality of opportunity by facilitating unequal access to alumni networks, among other social ills, and some members within these organizations are agitating to “Abolish Greek Life” (Nguyen 2020).

Consider residential policy as a “closer to home” example of this kind of dynamic. Neighborhood-level residential segregation is an endemic feature of American life. While segregation was caused by federal, state and local policy driven by anti-Black racial stigma (Rothstein 2017), it is maintained by land-use regulation. The kind of land-use regulation most relevant to segregation is anti-density zoning, which places restrictions on the quantity of housing that can be built on a given plot of land. For example, many localities<sup>5</sup> reserve a large percentage of land zoned for residential use specifically for single-family homes, which are houses that are not adjoined by any other units on all four sides. A locality might also institute building permit caps, prohibit more than 8 or so dwellings to be built per acre, or require a certain number of parking lots to

---

<sup>5</sup> A catch-all term to describe all local-level entities, including cities, towns, unincorporated areas, counties, and Census “places”. This term is used because local land-use regulation is extremely piecemeal and messy, and the specific kind of local jurisdiction in charge of local land-use policy varies widely.



be allocated per housing unit. These anti-density regulations are otherwise known as “exclusionary zoning” because their effect is to exclude Black and low-income people from living in certain areas. The underlying dynamic is simple: fewer housing units results in higher housing prices, which are more difficult to afford for people of color and low-income people.

Exclusionary zoning arose as a response to the Supreme Court’s *Buchanan v. Warley* (1917) ruling, which prohibited local ordinances that explicitly excluded individuals on the basis of race. Such zoning provides a facially race-neutral means of maintaining the homogeneity of neighborhoods, achieving much of the same racially exclusive effects by means of exclusion on the basis of economic means (Rothstein 2017). Recent work by social scientists finds that exclusionary zoning is associated with Black residential segregation in the U.S. and may inhibit desegregation measures (Rothwell & Massey 2010); exclusionary zoning is associated with higher levels of segregation, specifically, the segregation of the affluent (Lens & Monkkonen 2016); and that it limits the number of Blacks and Hispanics who move to exclusionary jurisdictions by limiting the number of rental housing units (Pendall 2000).

Despite exclusionary zoning’s discriminatory roots, it is not only, or solely maintained by racial stigma. It is in large part maintained by the interests of

homeowners in solving collective action problems. As Jessica Trounstine (2018, 2020) argues, homeowners need to act in coordinated ways to maintain the quality of certain local public goods, school quality, safety, and environmental quality, as well as the maintenance of their home values. If one homeowner in a single-family neighborhood decides to sell his property to a developer who intends to build an affordable housing apartment building, this can affect the quality of local public goods and everyone else's property values. Using this public-goods framework, Trounstine shows that "more stringent land-use regulations are supported by whiter communities and that they help preserve racial homogeneity" (2020, 443). This framework is compatible with Anderson and Kelly's stigma model of exclusion. But the dynamic can generate segregation even if we set aside the causal role of stigma. As Trounstine explains:

Once racist policies are in place, individual beliefs (e.g., racism) among individual beneficiaries of the system become largely irrelevant. Obviously, the level of racism among whites is both variable and impactful for political and economic outcomes...But, in the end, because government policy generates segregation through land use, the consequences of this variation are reduced. The choices of the racially resentful and the less racially resentful can become indistinguishable. Whites tend to make decisions that reinforce their privilege without thinking too deeply about it because they want stable property values, good schools, nice parks, and low-crime neighborhoods, and they have the financial opportunity to pursue those goals. (Trounstine 2018, 208)

#### **1.4.2.6 Information transmission**

A final and underestimated way in which sorting occurs is through the transmission of information within insular social networks. Maria Krysan and Kyle

Crowder (2017) argue that neither economic factors nor residential policy suffice to explain the degree and persistence of neighborhood residential segregation. Using a combination of qualitative and quantitative methods, they show that past neighborhood experiences and social networks affect mobility patterns in ways that solidify segregation. Social networks inform the consideration sets of those who are searching for a place to live. While economic models presume that searchers have equal access to information and are hampered only by affordability and discrimination, or are influenced by homophilous preferences, Krysan and Crowder show that these assumptions are misaligned with reality. Rather, according to their “social structural sorting” theory, residential segregation is facilitated by the transmission of information through social networks. Searchers gain information about the neighborhoods in which to search, and specific housing units to look for, through word-of-mouth within their social networks. Searchers also tend to want to live near their friends and family. Thus, social networks play an important role in sorting individuals into racially homogenous neighborhoods (p.78-79).

To sum up, sorting is a deep problem. By this I mean that sorting is robust across a multiplicity of preferences, motivations, and structural constraints that drive people to sort themselves into segregated groups. It is a mistake to focus only on supposed

linchpin causes that also happen to be clearly normatively troubling, such as racial stigma and exclusion by the majority. Theories of sorting that focus only on stigma or majority exclusion are liable to be incomplete at best, and misleading at worst. The problem of sorting demands more attention to various mechanisms, including those that are apparently innocuous.

### **1.5 Conclusion**

In this chapter, I have laid out the contours of the problem of the sorting. Sorting has three main facets. First, sorting undermines several plausible egalitarian ideals, including fair equality of opportunity, by enabling various forms of opportunity hoarding. Second, because our social relations are protected by the freedom of association, sorting widens the challenge to equality of opportunity beyond the scope of the well-known problem of the family. Although sorting interacts with the problem of the family in important ways, it is normatively distinct and just as troubling. Third, sorting is a deep problem: it can and does result from complex mechanisms driving associational choice, but these choices nonetheless result in opportunity hoarding. Several of these mechanisms cannot straightforwardly be traced to a clearly objectionable “linchpin cause” such as racial stigma. In any case, whether the underlying causes of sorting are malicious or innocuous, these associational choices are protected by

the basic liberty of freedom of association. Together, sorting's empirical complexity and this normative barrier to state intervention create a difficult normative problem.

Crucially, sorting is a problem that is due to closed, homogenous social network structures. It is not an individual-level problem, but one that is most visible on the collective level. Thus, the problem of sorting is a fundamentally social dilemma between freedom of association and equality of opportunity. The rest of this dissertation is devoted to fleshing out and exploring this dilemma.

## 2 Exclusion, entitlements, and dilemmas of sociability

### 2.1 Introduction

The freedom to associate entails the freedom to dissociate. After all, if we are not free to dissociate from some people, then we are compelled to associate with them. Thus, the freedom of association implies the freedom to exclude: we are permitted to decline social invitations, leave clubs and religious associations, ignore people for no reason at all, and shun people even if this shunning is borne from ill motivations such as racial animus. Exclusion may occur formally; we are permitted to enact rules that keep certain people, or certain types of people, out of our associations. Or exclusion may occur de facto; it may result from the fact that we simply choose some people over others, and the unchosen are therefore excluded from our associations. While exclusion occurs in both public and private life, they are treated quite differently in each domain.

**Public exclusion:** It is morally wrong to exclude a class of people from public life on the basis of their membership in some socially salient group, or on the basis of actual or perceived but morally irrelevant traits. It is wrong for instrumental reasons, because such exclusion excludes people from access to public goods such as quality education and a safe environment. It is also wrong for intrinsic reasons, because such exclusion harms the dignity of the excluded people. Because public exclusion is morally wrong, the state has a reason to prohibit it. Furthermore, those excluded from public life are entitled to redress.

**Private exclusion:** It is sometimes morally wrong to exclude people from private life, either on the basis of their group membership, or on the basis of actual or perceived but morally irrelevant traits. But such exclusion might also be morally

neutral. Private exclusion sometimes prevents people from accessing social and economic resources, and sometimes harms their dignity. Regardless of these consequences and whether private exclusion is morally wrong or neutral, the state should not interfere in private life, so it should not intervene on such exclusion. It follows that those who are excluded from private life are not entitled to redress.

The central purposes of this chapter are to explore the inconsistency between our treatment of public and private exclusion and between various forms of private exclusion, to draw out some parallels between public and private exclusion, and to offer a justification for this inconsistent treatment.<sup>6</sup>

It is curious that we employ a relatively laissez-faire moral approach to private compared to public exclusion. One possible explanation is that it seems inapt at best, and morally ugly at worst, to talk of an entitlement to redress in the realm of interpersonal relations. The claims of “incels” – men who are involuntarily celibate – show us just how dark these claims of entitlement can get. These men rage online at women who “deny” them sex and claim that they are “entitled” to have sex with women – perhaps even as a form of “emotional reparations” for racial injustice. This sense of

---

<sup>6</sup> For ease of reference, I’ll call them “public exclusion” and “private exclusion” respectively, with the qualification that the terms “public” and “private” are controversial and that their usage doesn’t always track principled political boundaries.

entitlement can have real consequences (Manne 2020). In 2014, Elliot Rodger, a young man who held misogynistic beliefs and also had mental health issues, murdered several women at a sorority house in a killing spree. He felt enraged that white, blonde women were depriving him of their affections. If the language of entitlement seems jarring here, it is because we have a strong intuition that *no one* is entitled to our company, especially not to our most intimate associations.

This intuition finds its expression in John Stuart Mill's (1859) famous words: that we have the "right to choose the society most acceptable to us" (Ch.4). In liberal theory, this right is known as the freedom of association. The freedom of association is a negative claim-right against interference on our associations, which also implies the right to exclude. After all, if you can't exclude some people from your associational life, you are not really free to choose your associations at all. Mill's words invite further questions about the scope of freedom of association, especially when it butts up against ideals of justice. Our exercises of freedom of association appear especially troubling when they are shaped by seemingly problematic patterns of preference and create unjust distributions of resources. In these situations, we are torn between two competing impulses: boldly asserting the Millian right, and the sneaking suspicion that, in doing so, we are acting unjustly. As Amia Srinivasan (2018) writes:



Suppose your child came home from primary school and told you that the other children share their sandwiches with each other, but not with her. And suppose further that your child is brown, or fat, or disabled, or doesn't speak English very well, and that you suspect that this is the reason for her exclusion from the sandwich-sharing. Suddenly it hardly seems sufficient to say that none of the other children is obligated to share with your child, true as that might be.

As Srinivisan forcefully asserts, "Sex is not a sandwich"; it is not a resource that it is wrong to withhold from some people, nor is it a resource that we should share with others out of pity, nor should the state should encourage us to share sex equally. While Srinivisan's main concern is with the entitlement to sex, her key point generalizes to other forms of interpersonal associations: the idea of justice in interpersonal relations is not something we can simply brush aside, but it is extremely fraught to frame these relations in moral terms. While we may not sympathize much with the incel's sense of entitlement to women, we might sympathize much more with people who are excluded from certain kinds of interpersonal relationships on the basis of their race, gender, or the intersection of both. We might even say that the latter suffer an injustice, and we certainly do say that the victims of public exclusion suffer an injustice. But why?

The goal of this chapter is to explore these issues. What explains our differential treatment of public and private exclusion? And within the realm of private exclusion, what justifies our sympathy for some of the excluded, but not others? The scope of this inquiry is wide, so let me clarify my aims and methodology. First and foremost, I am

trying to make sense of how we do think about these issues, so I seek theoretical consistency rather than foundational justification. Because we have such deep-seated, biased intuitions about social life, I am skeptical that a foundational inquiry is possible here. As such, I seek a theory that will accommodate the following data. It should explain why public exclusion, such as government-sanctioned segregation, is wrong, but it should exculpate certain kinds of private exclusion. Specifically, Rodger and those of his ilk are not wronged when women do not show them affection, but people who are excluded because of discrimination against their race or class are wronged.

I defend the following negative claims: there is a real inconsistency between our treatment of exclusion in the public and private realms. The special importance of autonomy cannot justify this inconsistency, but neither do individuals have entitlements to association. I offer an alternative justification for this differential treatment: the concept of the *opportunity* to associate, as distinct from the freedom of association itself. The state should offer opportunities to form associations by creating suitable search environments, but there are no individual duties or entitlements to associate with particular persons.

The chapter proceeds as follows. In Section 2.2, I elaborate on the nature of the problem by drawing out the parallels between public and private exclusion. In spite of

these parallels, public and private exclusion are treated very differently in theory and practice. The former is subject to redress, while the latter is not. This inconsistency is, on its face, puzzling. In Section 2.3, I give an overview of existing justifications for this difference. I defend Matt Zwolinski's (2006) view that the special value of autonomy does not justify this difference, so for the sake of consistency, we should either deregulate public associations or regulate private associations. But deregulating public exclusion would result in the harmful consequences that motivated this inquiry, and in Section 2.4, I argue that we do not have the entitlement to association that could justify regulating private social life. So it seems that we are still left in need of a justification for this distinction. In Section 2.5, I argue that the concept of the *opportunity to associate* provides such a justification. It allows us to retain the view that individuals are not entitled to others' association, while implying that the state should do much more to foster opportunities for association. Section 2.6 concludes. The upshot is that equality of opportunity is irreducibly social but does not entail that others are entitled to our company.

## ***2.2 Private and public discrimination: parallels and a puzzle***

### **2.2.1 Exclusion as a collective harm**

In various social domains, groups of people find themselves on the sidelines either because they cannot gain access to associations where resources are concentrated, or because they cannot gain access to any associations at all. Poor people and people of color find themselves shut out of middle-class neighborhoods that grant access to good-quality public goods. The elderly, the indigent, the formerly incarcerated, and those with mental illness are often forgotten or shunned. In the realm of intimate relationships, Black women and Asian men have a much more difficult time dating than do Black men and Asian women (Rudder 2014) – these patterns of desirability are two sides of a coin weighted by sexual stereotypes and racialized norms of attractiveness. While many of the individuals within these marginalized groups are able to form some meaningful associations, they are excluded from others.

This form of exclusion does not seem idiosyncratic. It is not like the exclusion of grumpy, misanthropic, or creepy people (Valentini 2016). Exclusion on the basis of individual character traits, while possibly problematic (cf. Brownlee 2020), differs from broad patterns of exclusion based on social characteristics. Of course, these types of exclusion may be more similar than they appear; it is possible that individuals who

belong to different social groups tend to have different personality characteristics due to the effects of culture and experience, so that on an individual level, one person may exclude another based on a mismatch between these individual characteristics. But it would be too hasty to revert to this as the explanation for patterns of exclusion, for these patterns often occur on the basis of little substantive personal interaction – for example, gay men on Grindr, a dating app, report receiving different qualities of interaction based on race. In particular, Asian men report being discriminated against, just as they are in the heterosexual world (Lim et al 2020).

The fact that there are patterns of exclusion, not apparently based on substantial personal interaction, involving certain groups is evidence that our associational life is cleaved by *discrimination*. In the broadest sense, “to discriminate against someone is to treat her disadvantageously relative to others because she has or is believed to have some particular feature that those others do not have. In short, discrimination against someone simply is disadvantageous differential treatment” (Lippert-Rasmussen 2013, 15). This sense of “discrimination” is agnostic about the underlying mechanism driving the disadvantageous differential treatment. The mechanism could be prejudice against members of a certain group, merely positive preference for members of a certain group,

or the filtering of a third variable that happens to correlate well with group membership, such as income. In this case, discrimination involves social exclusion.

### **2.2.2 Parallels between public and private exclusion**

There are two overlapping interpretations of the term “public exclusion”. On the first interpretation, public exclusion is exclusion by the government. For example, public exclusion occurred when the Federal Housing Authority enabled loans to white families in white suburbs, but not to Black families (Rothstein 2015). This was a form of discriminatory public action that resulted in the social exclusion of Blacks from certain physical spaces and public goods. On the second, broader interpretation, public exclusion refers to any kind of exclusion that occurs in the public realm, which concerns employment, housing, and education, regardless of the agent who carries it out. Private actors can exclude in public, too: businesses are private actors that can exclude certain individuals in employment, hiring, and relations with customers. Because the inconsistency between our moral treatment of public and private exclusion does not specify an agent, I adopt the second interpretation for the purposes of this chapter.

Private exclusion occurs in the so-called private realm of family, friendship, the home, and intimate interpersonal relations. It occurs when people choose, whether consciously or unconsciously, certain people to befriend, date, and live among. Our

intimate decisions are shaped by social categories such as race, religion, political ideology, and physical attractiveness. For example, marriages primarily take place within races (Reeves and Rodrigue 2015), and Black women and Asian men, who aren't taken to conform to aesthetic ideals, also tend to have a harder time finding partners (Rudd 2014).

Public and private exclusion parallel each other in two crucial ways: both are instrumentally harmful because they restrict access to opportunity, and both are also intrinsically harmful because they send a message that the excluded are not worthy of being included in social life, which can harm the psychological wellbeing of the excluded. Access to both social and economic opportunities are clearly restricted by public discrimination. When the state prohibits Blacks from living in the same neighborhoods or attending the same schools as whites, that form of segregation is a clear restriction of access to both these types of opportunity. When businesses refuse to serve transgender customers or requires them to use the bathroom that conforms to their biological sex<sup>7</sup>, that both limits their social opportunities and is intrinsically demeaning.

---

<sup>7</sup> North Carolina's HB2 law, repealed in 2017, prohibited transgender people from using the bathroom or locker room that aligned with their gender identity. Thus, people who were "assigned male at birth" were legally required to use the male bathroom even if they currently identified as women.

When building codes do not consider physical disabilities, and therefore prevent people with disabilities from accessing them, that limits these people's opportunities to work and socialize. That much is uncontroversial.

But private exclusion by individual actors can also have the similar effects. When wealthy individuals decide not to associate with those who are low-income, perhaps by sending their children to private schools and leaving public schools segregated with only low-income students, the effect is that the low-income individuals are excluded from important developmental and economic opportunities. A de facto separate and unequal school system still sends the message to students that they are not worth as much as the students whom they are separated from, "even though the physical facilities and other 'tangible' aspects of white and Negro schools may be equal" (*Brown v. Board of Education*, 1954). When Asian men or Black women find it hard to date because of gendered, European-oriented norms of attractiveness and sexual stereotypes, they may suffer harms to their psychological wellbeing. They may also suffer from the economic and social effects of remaining single in a world that is built around the assumption that adults will marry and form two-parent families at a certain age. For example, rates of single motherhood are high among Black women for complex reasons, creating obstacles to economic stability (Damaske 2017 et al). Transgender women face



difficulties trying to date compared to trans men (Blair & Hoskin 2018), and these challenges can be emotionally harmful. Political theorist Sonu Bedi (2015) calls the first of these “sexual racism” and argues that sexual racism is wrong because it “limits an opportunity that is a social primary good and a capability central to human dignity” (p.998). We will return to the idea that interpersonal relationships are social primary goods shortly.

Whether the agent of exclusion is an institution or the exclusion is an emergent effect of many individual decisions to exclude, the effects are similar. Social exclusion, particularly exclusion that is the patterned result of discrimination on the basis of race, gender, class, or other socially salient characteristics, is harmful.

### ***2.3 Justifications for private exclusion***

Despite the parallels between the consequences of public and private exclusion, they are treated quite differently. Public exclusion is seen as a moral wrong for which victims are entitled to redress. Such redress usually takes the form of state intervention that seeks to equalize opportunities in the public sphere, such as access to integrated public schools. But private exclusion is often not seen as wrong, and whether or not it is seen as wrong, its victims are not entitled to redress or other state intervention. What justifies this differential treatment? One plausible candidate is the idea that autonomy in

the private sphere is more important than autonomy in the public sphere. I argue that the autonomy justification is unconvincing, leaving us with a dilemma put forth by Matt Zwolinski (2006): for the sake of consistency, either deregulate public associational life or regulate private associational life. But neither of these options are palatable forms of state intervention. Public deregulation leaves the excluded vulnerable to harms, whereas private regulation relies on a controversial understanding of the right to access to social resources. I defend Laura Valentini's (2016) argument that access to social resources is not within the purview of justice, because we cannot be entitled to the association of others.

### **2.3.1 Autonomy and association**

At its core, autonomy involves the ideal of self-governance: the idea that an agent should exercise sole authority over her actions, and that her actions should not be subject to coercion or subversive influence. While there are many varying conceptions of autonomy, one thread common to all is that some degree of negative freedom is required for an agent to be autonomous (Mill 1859, Berlin 1969). This idea is expressed most famously in Mill's *On Liberty* (1859). Here, Mill argues that each individual has an interest in charting the course of her own life, free from coercive state regulations as well

as from oppressive informal social norms (Ch.3-4). Since our associations are core elements of life, Millian liberalism strenuously protects associations from interference. Thus, freedom of association has traditionally been interpreted as a negative claim-right that prevents our associations from being interfered with by the government or third parties. Freedom of association in the Millian sense also implies the freedom to exclude, for if we cannot exclude others from our associations, then we are not free to choose our associates at all.

But autonomy alone is too broad a concept for justifying the difference between public and private exclusion, for our associations are not confined to private life. Schools and public facilities are also sites of association. In mid-century Atlanta, white parents and students claimed that racial integration in schools violated their children's "freedom of association". Sandra Melkild, a white high school student in Atlanta, cited this freedom in seeking to transfer from a desegregating high school to an all-white one, citing "freedom of association". A year earlier, Thomas J. Wesley Jr., the leader of the Metropolitan Association for Segregated Education, had also invoked the right to freedom of association to resist the desegregation of schools (Kruse 2013, 161). While one might object that the freedom of association was being invoked in bad faith here, and

that it is not a legitimate invocation of this right, some further explanation must be offered for this objection.

Those who object to the use of freedom of association in cases of public exclusion do not take the route of explaining why freedom of association does not apply to the public realm. Instead, they argue that freedom of association is *overridden* by justice considerations in the public realm, but that it should be protected in the private realm for reasons of autonomy. Stuart White (1997) argues that exclusion rules, which determine who may or may not join an association, are presumptively legitimate if they protect an interest, such a religion's ability to worship its god. However, exclusion rules are not legitimate if they are opportunity-depriving or affront an individual's dignity (p.381). This principle would render school segregation on the basis of race an illegitimate exclusion rule. But there is an exemption for intimate relationships, such as friends and lovers. These, White argues, should be protected by the "special right to intimate exclusion":

an exclusion rule is legitimate if it serves to protect freedom of intimate association. And, once again, our interest in freedom of intimate association, so understood, is arguably so basic to the free development and expression of ethical personality that exclusion rules which protect this interest have a strong presumption of legitimacy even if they also damage others' opportunity interests. (White 1997, 386)

White justifies the special right to intimate exclusion by reference to its “basic” role in securing autonomy, here conceived of as the “free development and expression of ethical personality”. This justification assumes that intimate associations are more fundamental to autonomy than non-intimate ones, and that a clean division between intimate and non-intimate associations can be made.

Andres Moles (2014) makes a similar move. According to Moles, the “public ecology” of private associations has spillover effects into the social environment, which is composed, in large part, of many different associations. Thus, “Racist and sexist associations might change the social environment in which all citizens exercise their sense of justice and form, revise, and pursue their conception of the good” (p.91). They do so through a process of mental contamination, which creates automatic, unconscious biases that lie outside the agent’s control<sup>8</sup>. Psychology does not respect the normative line dividing the public and private sphere. Moles argues that because racist and sexist associations contaminate the social environment, they compromise autonomy, understood as the ability to exercise our ethical agency: “it is more difficult to exercise

---

<sup>8</sup> See Madva (2016) for a review of the implicit bias literature.

our moral powers in a society in which racist associations are protected, and this gives us weighty reasons to interfere with associations so that their existence has positive rather than negative effects" (p.94). Since these unreasonable associations compromise our autonomy, there is reason to regulate them. But like White, Moles carves out a special exemption for unreasonable *intimate* associations, such as friendships or relationships that discriminate along racial lines. Unreasonable intimate associations can exist within "brackets of privacy" due to their "spontaneous and informal nature" (p.99). Moreover, Moles argues, we cannot regulate intimate associations because no one has a claim-right to be befriended or married. Even if unreasonable associations are morally wrong, they do not violate anyone's moral rights, because no one has a moral right to be associated with in an intimate way.

White and Moles share one key assumption. They both assume (despite some equivocation) that there is a relatively clear normative boundary between the public and the private realms. This boundary is carved by the relatively greater importance of autonomy in the private than in the public sphere. In the private sphere, the value of autonomy overrides the potential harms of unreasonable or otherwise discriminatory associations. Thus, even if intimate associations are morally wrong, the moral wrongness of state interference on these associations is of greater concern.

Can this distinction be sustained? Matt Zwolinski (2006) argues that the autonomy justification is inconsistent, so it cannot do the boundary-policing work required of it. Specifically, autonomy cannot justify the right of private individuals to discriminate and deny that right to commercial actors. Zwolinski argues that the autonomy justification “is reliant upon a sharp distinction between individuals in their private capacity and individuals in their capacity as economic agents” (p.1058). On this distinction, individuals require maximum negative freedom in their private lives in order to fulfill their projects and aspirations, but there is no such need for negative freedom in commercial life. This difference is justified by the assumption that the economic actor has the sole goal of increasing profit. Therefore, “the state may legitimately prohibit or regulate any activity that is not required for this goal, if it can serve some legitimate state interest by doing so” (p.1059).

But this distinction, Zwolinski argues, cannot hold up under scrutiny. Some individuals use economic activity as a means to realize their private projects: “For small family-run businesses, the primary goal might often be something intangible, such as the continuance of a family tradition or the provision of meaningful work for people whom one cares deeply. Such businesses might have a multitude of interests other than profit maximization, some of which might constitute an interest in wrongful discrimination”

(p.1059). Federal law carves out an antidiscrimination exception for these small businesses, in recognition that discrimination by these actors might have a legitimate function. But even large for-profit corporations should not be assumed to have profit as their sole purpose. So, Zwolinski concludes, autonomy is an inadequate justification for our regulation of the public sphere, but not the private sphere. Our personal goals and interests do not fall neatly into the private sphere alone, but bleed over into the public sphere. To remedy this inconsistency, “we could regulate commercial discrimination less, or we could regulate private discrimination more” (p. 061).

I am sympathetic to Zwolinski’s argument that our differential treatment of discrimination in the private and public spheres is unjustified by autonomy considerations. But Zwolinski’s focus is on commercial versus private discrimination, so some additional reasons are needed to show that a version of this dilemma applies to public versus private social exclusion. I argue that the realization of autonomy in the private sphere sometimes *requires* public scaffolding; the private realm is not simply one that is insulated from state activity, but one that the state actively sculpts and protects in order to preserve some realizations of autonomy. So, autonomy cannot justify the insulation of private life from public concern. Autonomy is therefore a weak justification for a laissez-faire approach to private exclusion.



Consider marriage as a form of private association. Broadly, marriage is defined as a two-person intimate relationship that comes with a bundle of rights and exemptions. If two people are married to each other, all others are excluded from entering into marriage with either of them. For many people, marriage is part of their conception of the good life, so being married would help realize this interest. The state acknowledges the centrality of this interest to many of its citizens by recognizing marriage as a civil contract, apart from its religious bells and whistles, and by conferring certain legal and economic benefits onto married people. Without such state recognition of marriage, all intimate relationships – romantic, sexual, platonic, monogamous, polygamous, polyamorous, etc. – would be on a par in the eyes of the state<sup>9</sup>. So it is false to think that respect for autonomy in the private sphere requires that the private sphere be insulated from state activity. As the case of marriage shows, sometimes, it is quite the opposite.

Consider a second example, the home. The home is considered the inner sanctum of private life, but it is physically located amongst other homes. Any unowned plot of

---

<sup>9</sup> Cf. Elizabeth Brake (2011) and Clare Chambers (2017), who argue that state recognition of marriage unjustly privileges one form of intimate association over others, and that a liberal state should minimize (Brake) or abolish (Chambers) marriage.

land is located in the commons until it is carved up and doled out according to certain rules and regulations related to zoning and property rights. Then, where any particular home is located depends on a complex nest of residential land-use regulations that permit only certain types of dwellings, and therefore exclude certain types of people, to congregate in shared space. For example, the single-family home is prominent in the cultural imagination as a physical symbol of private life in the United States. But even this physical representation of privacy is enabled by land-use regulations, many of which prohibit denser dwellings such as apartment buildings and multiplexes from being built in the same neighborhoods as single-family homes (Pendall et al 2006, Rothstein 2017, Trounstine 2018). These anti-density land-use regulations often have the effect of pricing out low-income people and people of color, increasing neighborhood segregation (Rothwell & Massey 2009). Again, the realization of autonomy in the private sphere, including the exclusion of some from associational life, depends on public interference.

To sum up, private and public exclusion share parallels in terms of their consequences, but the state regulates them very differently, and we also think of them very differently in moral terms. Some philosophers argue that the autonomy of the private sphere justifies this differential political and moral treatment. However, this

autonomy defense is inadequate for justifying the inconsistency between public and private exclusion. If the realization of autonomy in private life sometimes requires significant state intrusion, as the cases of marriage and housing illustrate, then it is ad hoc to use autonomy in the private sphere as a barrier against state regulation of social exclusion.

### **2.3.2 The right to private inclusion**

At this point, we are left with Zwolinski's dilemma: either deregulate public exclusion or regulate private exclusion. Taking the first horn of the dilemma, deregulating public exclusion, is morally unappealing. Doing so would permit harms such as school and neighborhood segregation. While there may be a libertarian case for such deregulation (Epstein 2008, Lomasky 2008), addressing that case would take us too far afield from the starting point of this inquiry: the fact that social exclusion causes serious harms.

Since that is our moral starting point, let's turn our attention instead to the second horn of the dilemma, regulating private exclusion. Such regulation would require the state to enforce prohibitions on certain relationships and promote others. This sounds Orwellian, but instinctive repugnance is not a sufficient reason to avoid grasping the second horn. Aside from autonomy, what reason do we have to not

regulate private exclusion? Let's now turn to Laura Valentini's (2016) argument that access to social relationships is not a fundamental concern of justice. The excluded do not have a right to inclusion, because the framework of rights renders such an entitlement both incoherent and demeaning, in addition to intrusive. Briefly, justice concerns Hohfeldian claim-rights (1917). A claim-right is a directed duty owed by an agent to a particular other: "A has a claim that B  $\varphi$  if and only if B has a duty to A to  $\varphi$ " (Wenar 2020). In other words, claim-rights generate entitlements of some individuals to goods or actions furnished by others. Further, a claim-right must also be rightly enforceable by some agent, such as the state, at least in principle. What would the enforcement of this claim-right look like? Valentini argues that the state can either force or incentivize others to associate with the excluded, but both of these options incur "unacceptable moral costs" (p.56-57). These enforcement actions would be demeaning to the intended beneficiaries (Anderson 1999) or impossible, because meaningful social relationships simply cannot be forced through coercion or incentivization (Pettit 2015, Collins 2013).

While there are no duties of justice to associate with the excluded, Valentini argues that in some cases, there are at most duties of *beneficence* to do so. An agent acts beneficently when they promote the good of others. So, a duty of beneficence is a duty to

promote the good of others, and one acts wrongly when one fails to promote the good of others. But no one in particular has any right to one's good action (p.53). For example, I may have a duty of beneficence to watch Angela's cats when she is out of town and no one else can do so. But I would not wrong her if I declined to watch her cats, nor may she demand that I watch them.

In sum, the distinction between duties of justice and duties of beneficence allows Valentini to argue the following. No one has the *right* to the association of others, because the framework of rights is inapt when applied to associations. If it is inapt to speak of rights to association, then there is no basis for the state to enforce such rights, and it would be morally problematic to enforce such rights in any case. However, we do sometimes wrong people by refusing to associate with them. We do *not* wrong people when we exclude them on the basis of individual traits that make it burdensome to associate with them – perhaps they are creepy, frightening, or make us uncomfortable, and these are characteristics that are morally relevant to relationships (p.63). But we *do* wrong people when we exclude them on the basis of social characteristics that are morally irrelevant to relationships, such as race, gender, or class. The white women who ignored Elliot Rodger, the incel, may have wronged him if they refused to associate with him on the basis of his race, but they certainly did not wrong him if they refused to

associate with him on the basis of his creepy behavior and outbursts. Nevertheless, because this is a wrong that results from a failure of beneficence, Rodger is not entitled to anyone's association.

## ***2.4 Dilemmas of sociability***

Valentini's arguments bring us firmly into the realm of personal morality. If there is no right to association, then we shouldn't grasp the second horn of Zwolinski's dilemma – to regulate private exclusion. What follows is that our differential treatment of public and private exclusion is arbitrary.

A final option is available to remedy this inconsistency: those who are excluded from private life are entitled to moral redress, even if state regulation of private associations is neither desirable nor justified. Admitting that those who are excluded from private life are so entitled would not theoretically justify the inconsistency, but practically remedy it by bringing our differential treatment of public and private exclusion in line with each other. Kimberley Brownlee (2015, 2020) offers such an argument. Brownlee rejects Valentini's view that we do not have duties of justice to associate with the excluded, and that we do not wrong people when we refuse to associate with them because they are unpleasant to be around. Brownlee proposes instead that:

Associational freedom is at most 1) a highly restricted moral permission to associate, 2) a conditional moral permission not to associate provided our associative contributions are not required, 3) A highly constrained, content-sensitive moral claim-right that protects from interference only those of our wrongful associations that honour other legitimate concerns such as consent, need, harm, and respect (Brownlee 2020, 135-36).

The third condition, which I will focus on, goes surprisingly far. Brownlee argues that we do not have an unconditional freedom of association. We have individual moral duties<sup>10</sup> to associate with others, conditional on their consent, if their social needs are great, if the harms of not associating are serious, and if association is respectful (Ch.6). We do not have the moral permission to dissociate from people, even if they are unpleasant. The argument for this is negative: if we accept that we do have this moral permission, several collective dilemmas of negative sociability follow. Therefore, there cannot be an unconditional moral permission to dissociate (p.116-17).

These collective dilemmas of sociability are similar to the “tragedy of the commons” (Hardin 1968): they arise when each person, acting in a way that is better for each, generates an outcome that is worse for all in the same way (Parfit 1984). Brownlee calls these *distributive each-we dilemmas of sociability*, which “arise when we use our social resources (abilities, opportunities, and connections) negatively or positively in ways

---

<sup>10</sup> Brownlee does not distinguish between duties of beneficence and duties of justice.

that, if everyone does the same, lead to indefensible allocations of social resources”  
(p.124).

Foremost among these is the *rejected person dilemma*:

If Andy shuns Jo, then he will feel safe, contented, and well. But if everyone decides to shun Jo, then the result is not only bad for Jo, but worse for everyone’s safety, contentment, and well-being than if no one (or not everyone) shuns Jo. (Brownlee 2020, 125)

Maybe Jo is a creep, and that is why Andy shuns him. Or maybe Jo belongs to a stigmatized minority, and Andy is racist. In these each-we dilemmas, the explanation for Andy’s rejection of Jo does not matter; the focus is on the harms that would result if Andy reject Jo Brownlee argues that each Andy faces costs if everyone or most people shun Jo. When everyone shuns Jo, each Andy loses out on Jo’s contribution to the pool of social resources. If Jo’s isolation is persistent, then the probability of harm to each Andy increases: Andy might become the target of Jo’s anti-social behavior, violence, or psychosis (p.125-26). Thus, Andy does not have an unconditional right to dissociate from Jo.

A variant of the *rejected person dilemma* is the *rejected people dilemma*:

If Andy shuns all people like Jo, then he will feel safe and contented. But, if everyone who is not like Jo shuns all people like Jo, then over time this will be worse for the safety and contentment of all, and it poses the risk that the pool of available social resources will be smaller than the pool of needed social resources. (This dilemma turns on the assumption that the people like Jo are unable to secure social connections from each other.) (Brownlee 2020, 126)



Brownlee concludes from these each-we dilemmas that even if there is no positive right to associations per se (cf. Valentini 2016), we do not have the right to exclude others from social *opportunities*: “since chronic rejection can put us at risk of dependency, others’ freedom to refuse association is checked by the requirement that our minimally adequate social opportunities be meaningful and not render us dependent” (p.128).

What follows from the lack of our freedom to associate as we wish are certain duties to offer associational opportunities to those who need us, if the harms of exclusion are severe enough (Ch.6). Notice that Brownlee’s dilemmas are meant to show that even in an extreme case, where Andy’s safety and contentment are *threatened* by Jo, Andy has duties to offer social opportunities to Jo. If such an extreme case can generate directed duties from Andy to Jo in these cases, then individuals who do not threaten Andy’s safety and contentment are also entitled to social opportunities from Andy.

I argue that Brownlee’s dilemmas do not succeed in undermining the Millian interpretation of freedom of association as a negative claim-right. My first objection is that the dilemmas can arise even if each Andy gives Jo opportunities to interact, so Brownlee cannot merely conclude that the dilemmas show that we lack the freedom to dissociate as we wish, and that Andy ought to give Jo social opportunities. Consider a variant of the rejected person dilemma, which I call the *many opportunities dilemma*:

Suppose Andy is a kind, forgiving person, but Jo is an unstable, creepy person. Although Andy gives Jo many social opportunities, and employs significant charity in interpreting Jo's actions, Jo repeatedly makes Andy feel unsafe. At some point, Andy decides to give Jo a wide berth. Once Andy shuns Jo, Andy feels safe, contented, and well. But if everyone like Andy decides to give Jo a wide berth, after many social opportunities, then Jo still ends up excluded. Over time this will be worse for the safety and contentment of all, because Jo will be more likely to increase the extent of his anti-social behavior.

And here is a variant of the many-chances dilemma. Call this the *many people, many opportunities dilemma*:

Suppose that all Andys give all Jos many social opportunities. But all Jos repeatedly make all Andys feel unsafe. At some point, all Andys decide to give all Jos a wide berth. All Andys then feel safe, contented, and well. But if everyone like Andy decides to give everyone like Jo a wide berth, after many social opportunities, this still results in the exclusion of all Jos. This exclusion, if persistent enough, will be worse for the safety and contentment of all, because Jo will be more likely to increase the extent of his anti-social behavior.

These many-opportunities dilemmas show that Brownlee's conclusion is too weak. The conclusion that Brownlee should draw is that Andy must *actually* associate with Jo, not merely provide social opportunities. Falling short of this conclusion leaves open the possibility that the exact same harms will result in the many-opportunities dilemmas as in the rejected-people dilemmas. But another possible lesson to draw from this dilemma is that Andy acts wrongly in excluding Jo, with finality, at the end of the many opportunities. Is Andy wrong to do so? Should Andy continue to give Jo further opportunities for association? Brownlee might respond that Andy should indeed

continue to leave the door open for Jo to interact with him. She concludes her book: “Even when we have acted egregiously or we frighten others just by being who we are, we need to know that we can nonetheless be sure of a few specific people and sure of our acceptance within the wider social world” (p.193). Brownlee’s vision is a deeply humane, compassionate one. But it overlooks the reciprocal nature of social interaction, as well as the burdens that associations with Jo puts on Andy. In empathizing with Jo, Brownlee ignores the plight of Andy.

The many-opportunities dilemmas show that social opportunities do not guarantee that associations will be formed. This is due to the reciprocal nature of social interaction – call this the reciprocity objection. All that Andy can do is offer Jo opportunities for social interaction. But Jo must hold up his end of the bargain. Jo cannot repeatedly fail to properly take up these social opportunities and reasonably expect Andy to leave the door open for him. If someone continually fails to support and engage with her friend, she is at the very least being a bad friend; she might not even be a friend at all, and it is not unreasonable for her friend to disengage with her. And even relationships that are not between equals, such as parent-child relationships of care, can be construed as reciprocal in some sense. For example, the Confucian ethic of filial piety is an intergenerational form of reciprocity in which the child eventually takes care of an

aging parent as reciprocity for the parents' care. In continuing to act in ways that make Andy feel unsafe, Jo fails to hold up his end of a reciprocal social interaction. But if Andy decides that enough is enough from Jo, the same harmful consequences may result. For the rejected-people dilemmas to show that there is no unqualified right to freedom of association, however, Brownlee must conclude that Andy should associate with Jo, in spite of the burdens that Jo places on him, because that is the only way for Andy to avoid harm

Here, we must weigh the burdens that Andy faces by continuing to leave the door open for Jo compared to the costs of Jo's social exclusion. If Jo indeed makes Andy feel unsafe and unhappy, then this fact alone gives Andy a good reason to exclude Jo. Continually offering Jo social opportunities would be too burdensome on Andy, who is not a "social resource" that Jo may tap without contributing. But if Jo is only slightly irritating and not unduly burdensome, then perhaps Andy should be open to further interaction with Jo. In either case, Jo's behavior is a morally relevant factor for evaluating Andy's obligations to include him. Brownlee ignores this factor.

A second reason to reject the rejected-people dilemmas is that they lead to repugnant conclusions in practice, because they ignore the fact that Jo does not need

everyone's associations equally. Consider this variant of the rejected-people dilemmas, which I call the *selected person dilemma*:

Jo only seeks Bella's company, not Andy's. If Andy decides to shun Jo, Andy will feel safe and well, and Jo will be unaffected. If Bella decides to shun Jo, Bella will feel safe and well, but Jo will be very upset. Bella's rejection is not only bad for Jo, but worse for everyone's safety, contentment, and well-being than if no one (or not everyone) shuns Jo.

The selected-person dilemma's collective variant is the *selected people dilemma*:

People like Jo only seeks the company of people like Bella, not of people like Andy. If all Andys decide to shun all Jos, all Andys will feel safe and well, and all Jos will be unaffected. If all Bellas decides to shun all Jos, all Bella will feel safe and well. However, the latter result is not only bad for all Jos, but worse for everyone's safety, contentment, and well-being than if no one (or not everyone) shuns Jo, and especially worse for all Bellas.

The selected-people dilemmas are more realistic versions of the rejected-people dilemmas. We usually do not seek company indiscriminately; our associational needs are highly specific to certain *kinds* of people, or even to particular people. Elliot Rodger, for example, was not upset that men or Black women declined to associate with him. He was particularly aggrieved that white, conventionally attractive women – Bella, in the selected-people dilemma – did not want to associate with him.

The selected-people dilemmas show that if the harms of exclusion generate duties to offer social opportunities to Jo, then these duties will fall unevenly on certain people. In some cases, it is likely that such associational duties will fall

disproportionately on women, and that these duties will take particular forms – perhaps Jo does not merely want general social opportunities with Bella, but sexual and romantic opportunities. While such a case may seem extreme, it is unfortunately quite common. So, if Brownlee’s theory implies that women do have such duties, then this is a serious reason to think that it is false.<sup>11</sup> Its conclusion is repugnant. One may, of course, bite the bullet and accept such a repugnant conclusion, and maintain that Bella does have duties to offer social opportunities to Jo because of the harms to everyone that may result if she rejects him. One might go even further and maintain that Bella has continued duties to offer *particular* kinds of social opportunities to Jo. I contend that it is more humane to Bella, who is unfairly burdened both by Jo’s noxious attention and the weight of the continued duty to associate with him, to take this conclusion as a reductio of Brownlee’s argument.

A third reason to reject the implication of the each-we dilemmas is that its purported implication simply does not follow. Brownlee’s perspective is explicitly Kantian (p.132-33), so she ignores a consequentialist interpretation of the dilemmas. The

---

<sup>11</sup> Surprisingly, Brownlee does not address such a case.

non-universalizability of Andy or Bella's actions leads Brownlee to conclude that Andy and Bella act wrongly in dissociating from Jo. But from a consequentialist perspective, Andy and Bella act rationally. By withdrawing from Jo, Andy and Bella secure their own interests while removing only a small part of the social resource pool. There is no instrumental reason for them to associate with Jo, given the significant costs of doing so, and their withdrawal does not "make a difference" to Jo's exclusion, assuming that a large enough number of people exist within Jo's orbit. As some consequentialist philosophers have argued, if an action does not make a morally relevant difference to an outcome, then there is no reason to do it (Sinnott-Armstrong 2005, Sinnott-Armstrong and Kingston 2018, Budolfson 2018, Nefsky 2019). Therefore, an equally plausible interpretation of the each-we dilemmas is that they are collective action problems that do not intrinsically generate duties for Andy and Bella to offer Jo social opportunities. This collective action problem cannot be dismissed by Kantian fiat, but must be dealt with on its own terms.

To sum up, I have argued that Brownlee's argument against an unconditional freedom of association does not succeed. There are four related reasons for this. First, Brownlee's conclusion is too weak: from the each-we dilemmas, she cannot just conclude that we lack freedom to dissociate, and therefore need to offer social

*opportunities* to the excluded if the harms of exclusion are severe enough. She must conclude that we should *actually* associate with the excluded in substantive ways. But this proper conclusion leads to the second objection: that an association should be reciprocal, and if an association continually burdens one party, it is not reciprocal. A third objection is that the dilemmas lead to a repugnant conclusion in practice. Our associational needs are typically directed towards certain people, and these needs take particular forms. If we have moral duties to offer associational opportunities to those who need them, then some people will suffer from unfair burdens of association. This is a reductio of the argument. Fourth, each-we dilemmas cannot be resolved by simply invoking the Kantian categorical imperative, but should be addressed on consequentialist terms as well.

My arguments constitute a defense of Valentini: no one has a right to the association of others. At most, we should give others associational opportunities if doing so is not unduly burdensome. Unfortunately, the arguments so far appear to leave us in a dismal spot. Because no one has an entitlement to anyone else's association, the state has no basis for regulating private exclusion, but neither should we remedy the differential treatment of private and public exclusion by morally requiring some



individuals to associate with others. At this point, we might conclude that the differential treatment of public and private exclusion is unjustifiably inconsistent.

## **2.5 Associational opportunities**

This resignation would be premature. I argue that the concept of *associational opportunities* actually provides us with a principled way of justifying the differential treatment of public and private exclusion, even if there is no right to association. Chiara Cordelli (2015) argues that associational opportunities<sup>12</sup> – opportunities to access social or relational resources, such as relationships – are a Rawlsian primary social good. They are a primary good because they are necessary for exercising the two moral powers: a sense of justice and a conception of the good (Rawls 1993, 19). Associational opportunities are therefore subject to the principle of fair equality of opportunity. I argue that the concept of associational opportunities can also help us accommodate the following data (from Section 1): 1) public exclusion is wrong, 2) private exclusion of the unpleasant is not wrong if it is due to individuals' character traits, 3) but private exclusion is wrong if it is discriminatory.

---

<sup>12</sup> She calls them “relational opportunities”.

Public exclusion, such as residential or school segregation, is wrong because it excludes people from associational opportunities. Associational opportunities are often enacted in physical space. When people are excluded from public schools, public pools, or certain neighborhoods because of their race or class, they are excluded from associational opportunities. Thus, minorities are entitled to redress by the state in the form of increased associational opportunities. In practice, this will take the form of policies that promote the substantive school and residential integration, such as eliminating restrictive land-use regulations and severing the funding link between property taxes and public schools in the neighborhood. One might object that perhaps excluded people would rather not associate with those who stigmatize them. However, that is not the point. Excluded people are well within their associational rights to choose whomever they want to associate with. They should simply be given the opportunity to make this choice.

By contrast, private exclusion is not wrong if it is because of individuals' character traits or behavior, such as unpleasantness, violence, or toxicity. If these individuals have been given associational opportunities, but those individuals have not used their associational opportunities appropriately, then they are not wronged when they are excluded. But perhaps this result is too harsh. Some people just have difficulties

forming meaningful associations through no fault of their own. Perhaps they have social difficulties, suffer from mental illness, or perhaps they were raised in an environment that was not conducive to the healthy development of social skills. For the purposes of private exclusion, it does not matter. Regardless of the genesis of these individual traits or behaviors, they are not wronged when particular individuals exclude them, because these traits are morally relevant to relationships. While the causes of such exclusion might be unfortunate and sad, they should be addressed at a structural level prior to that of a particular relationship. As Cordelli (2015) argues, people who lack “relational talents” due to mental illness or disability should be given associational opportunities through access to psychiatric care, social workers, and adequate schooling (p.107). But they are not wronged when they are excluded from private associations due to morally relevant facts about them, however unfortunate the causal story behind those facts is. Return to the example of Elliot Rodger, who suffered from an unspecified form of autism spectrum disorder: Rodger was entitled to have his associational abilities nurtured by adequate psychiatric care, but he was not thereby entitled to the association of any particular person. Neither was he wronged when women shunned him on the basis of his inappropriate behavior.

However, private exclusion *is* wrong if it is based on individuals' social rather than individual characteristics. The reason is that these individuals are denied the opportunity to associate in the private realm. To exclude someone on the basis of their race or class is to make a preemptive ruling about their suitability for association, based on some morally arbitrary characteristic. If women excluded Rodger because of his race, then they did wrong him by denying him an associational opportunity based on some arbitrary factor. For a less extreme, more common example, consider the racialized patterns of online dating that I began with. Black women and Asian men are preemptively ignored as romantic or sexual partners are denied associational opportunities on the basis of stereotypes (Rudder 2014). These people are wronged by those who ignore them, and they are wronged because they are denied associational opportunities on the basis of a morally arbitrary factor – not because they are denied particular associations. To say that these people are wronged by their form of exclusion is compatible with the excluders being *entitled* to have specific physical preferences in the intimate realm. Those who exclude on the basis of race are both within their rights to do so, because no particular person is entitled to their particular association, and also wrong to do so, because they are excluding certain people from associational opportunities. Unlike unpleasant people, who are excluded on the basis of who they

really are, those who are excluded on the basis of social characteristics are not even given the opportunity to fail at forming an association. Thus, people may have the duty to provide social opportunities (if they are not unduly burdensome) to those who would otherwise be excluded due to morally relevant factors. The state and other institutions, such as the firms who create online dating apps, may have an obligation to cultivate such associational opportunities. For example, the designers of platforms such as Grindr and Tinder should eliminate the use of race/ethnicity filters (Hutson et al 2018), because this predictably wrongs people of certain races by constraining their associational opportunities. Additionally, because opportunities to form private associations are constrained by the structure of public space, the state may have a reason to shape public space in ways that promote private associational opportunities.

## ***2.6 Conclusion***

In this chapter, I have attempted to grapple with one central question: what, if anything, justifies our differential moral treatment of public and private exclusion? A related question follows from this inquiry: what justifies our differential moral treatment of private exclusion that is based on individual characteristics compared to private exclusion that is based on social characteristics? We think that public exclusion is morally wrong, and that its victims are entitled to redress by the state. But private

exclusion is commonly seen to be either morally neutral, and even if it is morally wrong, it does not warrant state intervention on our associations. Further, private exclusion that is based on discrimination seems wrong, even if private exclusion that is based on individual character traits is not. There has not yet been a principled way of justifying these three common intuitions.

Answering this question requires delving into questions related to the right to freedom of association, such as whether anyone is ever entitled to anyone else's association. I have argued that public and private exclusion both harm individuals in similar ways, raising a puzzle about why private exclusion is exempt from morality or state regulation. The autonomy of the private sphere provides a poor justification for this differential treatment, so for the sake of consistency, we should either deregulate public exclusion or regulate private exclusion. Neither of these options are palatable. I then defended the argument that no one has a right to particular associations, even if exclusion has harmful consequences.

The concept of associational opportunities gives us a way to justify the difference between public and private exclusion. The denial of associational opportunities is a wrong in certain cases, but it does not follow that anyone is entitled to anyone else's association. In this respect, this chapter hopes to enlarge the scope of moral and political

concern for private, associational life, while protecting the individual prerogative to actual associations.

## 3 An intrapersonal, intertemporal solution to an interpersonal dilemma

### 3.1 *The structure of collective action problems*

Consider a variant of the *rejected people dilemma* from Chapter 2. Call this the *excluded people dilemma*:

If Andy excludes all people like Jo, then he will feel happy. But, if everyone who is not like Jo excludes all people like Jo, then over time this will be worse for the happiness of all, as it poses the risk that the pool of available social and economic resources will be smaller than the pool of otherwise available social and economic resources. Jo's ability to contribute to this pool will diminish as his exclusion persists, and Andy's ability to empathize with and consider the interests of Jo will also diminish. (Brownlee 2020, 126)

Brownlee argues that we cannot will all Andys to exclude all Jos. Because Andy's exclusion of all Jos is non-universalizable, this shows that Andy should not exclude all Jos. But Brownlee's perspective is staunchly Kantian, so it is vulnerable to the consequentialist objection that Andy may, and indeed is rationally required to, exclude all Jos.

The consequentialist argues that in collective dilemmas, many individual actions create a harmful outcome in the aggregate, but each individual action seems to make no difference to the outcome. Therefore, there seems to be no reason for the individual to "cooperate": to take the action that would be part of the set of actions that bring about the desired outcome. The consequentialist perspective is not simply abstract, but does



seem to characterize the way people actually think about collective social dilemmas and explain why various forms of social sorting persist. To make this vivid, consider a variant of the excluded people dilemma – call this the *segregated people dilemma*:

If all people like Andy live in a separate neighborhood from all people like Jo, then all Andys will feel safe, content, and happy. All Andys will have access to a high level of public goods and services. But as a result, all Jos will only have access to substandard levels of public goods and services; these conditions will breed resentment toward Andys and various social pathologies. While Andy sympathizes with Jo and feels strongly that it is unjust that Jo is excluded from his nice neighborhood, Andy thinks there is nothing that he can or should do about Jo's exclusion. If all or many people like Jo move into Andy's neighborhood, then Andy's property value will go down because few other Andys want to live among Jos, and it is likely that over time, his neighbors will sell their properties to Jos. And if Andy is going to be honest, he doesn't want to live among Jos either.

Andy cannot, by himself, bring about justice for Jos. He only has control over his property in the neighborhood, which represents a large proportion of his assets. If he were to sell that property and move out of his neighborhood to live among Jos, there's no guarantee that he'd be welcome among them. Plus, Andy simply feels more comfortable living among Andys, and he knows that an influx of Jos into this neighborhood would undermine that comfort. And there are other Andy-neighborhoods in the area, so there's no reason why Andys in *this* neighborhood have to take on the burden of undermining their comfort and property values so as to make it possible for Jos to live there. So, Andy reasons, there is no reason for him to challenge the status quo. As a result, segregation between Andys and Jos persists.

As the Kantian predicts, collective harm is a result of this reasoning and could be avoided if enough Andys decided to cooperate. While Andy's reasoning might seem cynical and self-serving, it also rests on good consequentialist grounds. Why should Andy cooperate, when no individual act of cooperation makes no difference to the outcome? Julia Nefsky (2019) calls this the *inefficacy problem*. As Nefsky and other moral philosophers have argued, the inefficacy problem makes it extremely difficult to find a moral reason for action in collective action problems. This debate parallels a longstanding debate in rational choice theory, where the consensus view is that no instrumental reasons to cooperate can be derived from the preference to make a difference in collective action problems. If this is the only preference under consideration, the dominant strategy is to defect rather than to cooperate. Therefore, rational choice theorists have sought to explain cooperation by appealing to preferences that are only indirectly related to the outcome, such as expressive preferences.

My strategy is a bit different from both these camps. The goal of this paper is to excavate an instrumental, outcome-based reason for action in collective impact cases by appealing to diachronic rationality – rationality over time, especially with respect to our commitments and values. The core idea behind my argument is this: in daily life, each of us has commitments that require many individually ineffective actions to be realized. It

is often rational to “defect” at a time, or to put off taking one individual action. If we do this often enough, we will fail to realize the relevant commitment. However, commitments bind us to “cooperate” even when this seems irrational at a time, because we know that defecting too often will undermine our ability to follow through on our commitments. In this paper, I argue that this strategy, quotidian as it is in daily life, can shed light on our reasons to cooperate in collective situations. Throughout this paper, I will remain agnostic as to whether moral reasons can be provided for action in collective impact situations.

The rest of this paper proceeds as follows. In Section 3.2, I critically sketch out different approaches to the inefficacy problem and the collective action problem that it derives from; I argue that calls toward purely structural change or political advocacy merely give rise to another iteration of the inefficacy problem. In Section 3.3, I draw on Michael Bratman’s planning framework to argue that switching to a diachronic perspective can give us outcome-based, instrumental reasons to cooperate in collective action problems, in spite of the pull of the inefficacy problem. This allows us to bypass the expected utility debate. Section 3.4 argues that these reasons are relative to each individual’s commitment to the good. Section 3.5 concludes.

Before launching into the main argument, it will be good to make clear my assumptions, methodological commitments, and the scope of my target. I will assume that diachronic rationality is a distinct concept from synchronic rationality, which concerns rationality at a time, and that there is such a thing as a subjective reason. I assume, along with much of the collective action literature, that collectives are constituted by individuals, so collective action requires individual action;<sup>13</sup> hence the focus on individual reasons for action. I also assume, more controversially, that intentions, particularly commitments, carry with them certain norms of coherence, consistency, and stability.

Finally, my arguments are targeted toward a class of individuals with a particular psychology, not to everyone. The Ordinary Person is someone who is rationally compelled by the inefficacy problem towards inaction. Unlike someone who does not care at all about collective harm, the Ordinary Person is committed to the good. But she is not someone with an activist-like mindset who acts on deontological commitments, nor does she have an overwhelming desire to disregard instrumental

---

<sup>13</sup> This ontological position does not entail methodological individualism, which holds that all human activity is best *explained* in terms of the activity of individuals (Elster 1989). Methodological holists usually also accept the claim that all human activity is *ontologically* constituted by the activity of human individuals.

calculations in favor of virtuous action. The inefficacy problem is truly a *problem* for this person; she would like to act but finds no outcome-based reason to do so. Thus, she is the person whom we need to provide reasons for. Even though my arguments are limited in scope, I think that the Ordinary Person is a challenging enough target. She is familiar to many of us.

### **3.2. *The inefficacy problem***

The inefficacy problem lies at the heart of collective action problems. According to the inefficacy problem, “if acting in the relevant way won’t make a difference, it’s unclear why it would be wrong. Each individual can argue, ‘things will be just as bad whether or not I act in this way, so there’s no point in doing otherwise’” (Nefsky 2019, 2). Acting in the relevant way (“cooperating”, for short) doesn’t make a difference because it is neither necessary nor sufficient for bringing about the desired outcome.

The inefficacy problem is starkest in cases where an individual action does not cause any small-scale, localized harm by itself, so we cannot bypass the inefficacy problem by appealing to these small-scale harms to deliver moral reasons against action. For example, consider segregation: if a white, wealthy family moves from the city to the suburbs, they do not harm anyone. But when enough of these families do so, this is

known as “white flight”, which drains cities of economic resources and perpetuates racial segregation (Massey and Denton 1993, Kruse 2013).

Let’s distinguish between two versions of the inefficacy problem. On the weak version of the inefficacy problem, individuals are not morally required to cooperate on a first-order level if doing so won’t make a difference. No one is required to stop driving if doing so won’t make a difference to climate change. Assuming that there is such a thing as supererogatory action, it could be morally good or virtuous of individuals to cooperate, but the key point is that we are not *required* to (Sinnott-Armstrong 2005, with Kingston 2018). Rather, what we are required to do is to cooperate on a second-order level, by engaging in political advocacy to push our government to act against collective harms (Sinnott-Armstrong 2005, 2018 p.185). For example, we ought to push our governments to regulate gas-guzzling SUVs, but we need not refrain from joyriding in these SUVs. Call this solution *Political-Not-Personal (PNP)*.

However, PNP is vulnerable to the inefficacy problem as well.<sup>14</sup> It is not clear why we are morally required to take political action if individual political action doesn’t

---

<sup>14</sup> Nefsky (2019) raises this point in footnote 8.

make a difference and is also personally costly. It is not realistic to expect any individual political action to make a difference to an outcome such as climate change; for this to happen takes a great deal of luck and is exceedingly rare, if it ever happens.<sup>15</sup> Given the difficulty of this problem, various strategies to make sense of voting, protesting, and organizing have appealed to fact that political action aims to secure some form of private benefit to the individual, rather than to its direct efficacy in securing a public good. For example, political action may help activists gain personal status (Tullock 1971), help those who are unhappy with a social norm satisfy a direct preference for conditional cooperation (Bicchieri 2005), or function as a means for expressing one's political preferences (Lomasky & Brennan 1993). Or political action may simply be a way to act in accordance with one's moral beliefs, whether these be grounded in virtue or deontological principles. None of these strategies appeal to efficacy itself.

Moreover, the inefficacy problem is heightened by the costs of political action.

Take voting as a simple instance of political action. The costs of voting can range from

---

<sup>15</sup> By "luck", I mean that the individual must take the right sort of action, have the right sort of personality, and ultimately, being in the right place at the right time, among other factors. For example, Greta Thunberg is one person who has had an outsize impact as a climate change activist, but even with her prominence, carbon emissions continue unabated.

the relatively trivial, such as standing in line to cast a vote; to the more expensive, such as deeply researching candidates' policy positions and the relevant social science so as to be able to cast an informed vote; to extremely serious ones such as dangers to one's physical security in situations where voter intimidation is a problem. In addition to these direct costs of political action, there are also significant opportunity costs, such as the time and effort one could spend on one's leisure pursuits, personal relationships, or career. If the inefficacy problem renders even an action as simple as voting irrational, then it is not clear what efficacy-based reason could be provided for more aggressive forms of political advocacy, such as direct action against environmental injustice. Therefore, for those who are focused on the efficacy of action, it is not clear what the justification is for a moral or political obligation to undertake political action. PNP cannot be justified by the efficacy of political action.

PNP also fails on a strong version of the inefficacy problem. On the strong version, you not only have no moral requirement to act; you also have no moral reason to, supererogatory reasons included. As Nefsky writes, "When one says 'but it won't make any difference,' more than just saying 'it doesn't seem that I am obligated to act in that way', one is saying 'there doesn't seem to be any point in acting that way.' Doing so, in this light, looks like a mere waste" (Nefsky 2017, 2744-2745).



The inefficacy problem applies in both threshold and non-threshold cases. Let's first look at threshold cases. In threshold cases, once a threshold is reached and exceeded, a significant harm results. To avoid crossing the threshold, a certain number of people are required to cooperate. For example, food supply chains are thought to fall under the umbrella of threshold cases (Kagan 2011, Budolfson 2018). We might abhor the suffering caused to animals in factory farms, but there is no clear efficacy-based reason to cooperate (i.e. refrain from buying factory-farmed meat). If not enough people cooperate, then your cooperation will be futile. If enough other people cooperate, then your cooperation is superfluous, so there is also no reason to cooperate (Nefsky 2019). Therefore, the dominant strategy is non-cooperation.

In many collective action problems, we will not be operating under anything like the degree of certainty implied in the claim that acting "*won't* [italics mine] make any difference". Most of the time, we are confronted with situations in which we are simply not sure whether doing so will or won't make a difference. As such, Shelley Kagan (2011) and other consequentialists, such as Alastair Norcross (2004), Peter Singer (1980), and John Broome (2019), have appealed to expected utility in threshold cases. As a normative theory of rational choice, expected utility theory tells us to choose the act with the highest expected utility. Kagan argues that in threshold cases, a small probability of

being the triggering act, multiplied by the consequences of the outcome, will always deliver negative expected utility, so consequentialism can straightforwardly condemn the act on the grounds of its negative expected utility (Kagan 2011, 119-120). Even if one's action is unlikely to make a difference, its negative expected utility tells us not to do it.

Nefsky (2011) and Mark Budolfson (2018) have come out with arguments against expected utility that are, to my mind, convincing. Both argue that we cannot a priori conclude that expected utility will always come out negative in every threshold case. Nefsky says that "Whether it does or not depends on the probabilities and on the goodness or badness of the relevant consequences" (2011, 369). Budolfson makes an even stronger claim. He argues that due to empirical facts about supply chains in industrial agriculture, it will *rarely* be the case that expected utility shakes out in favor of action in that sort of case. The mistake that Singer, Norcross, and Kagan make is that they assume that expected utility of an action is equal to the average effects of all similar actions. However, it is more likely that the expected utility of cooperating approaches zero, so expected utility provides weak consequentialist grounds for cooperating.

Once we also consider the instrumental costs of cooperation, Nefsky and Budolfson's skepticism about the ability of expected utility to ground reasons to

cooperate is strengthened. Assuming an individual has a preference to help bring about the desired outcome, she can expect that her action will most likely not do so, and that its utility will be outweighed by the costs of action. In sum, expected utility cannot provide a general solution to the inefficacy problem, especially when we expand the scope of the relevant consequences to personal costs.

Expected utility has even more difficulty in handling non-threshold cases, which are cases where there is no tipping point at which an outcome decisively occurs. Non-threshold cases resemble Sorites paradoxes, in which predicates with vague boundaries generate difference-making problems. Consider the harmless torturer (Parfit 1984) and the puzzle of the self-torturer (Quinn 1990). In the harmless torturer case, a thousand torturers turn a switch on some instrument once. Each turn increases a victim's pain imperceptibly, but the aggregate effect of a thousand turns is that each victim ends up in severe pain. However, "none of the torturers makes any victim's pain perceptibly worse" (Parfit 1984, 80). At no point does expected utility tell each of the torturers that they should not turn the dial, so none of them act wrongly.

Quinn's puzzle of the self-torturer is similar, but its point is to show that an agent can make a series of rational decisions that land her in an unwanted outcome due to the intransitive structure of her preferences. The reasons involved in this situation are

entirely instrumental. The self-torturer is hooked up to a medical device that increases her pain level by one imperceptible increment at each turn of the dial. At each turn of the dial, she receives \$10,000. Because any individual increase in pain is imperceptible, it is rational for her to turn the dial to receive the \$10,000. But the self-torturer eventually ends up in a state of severe pain, which she disprefers to the total amount of money she's gained, due to her intransitive preferences. Orthodox rational choice theory (RCT) tells us that the self-torturer is irrational because her preferences are intransitive, so it is no surprise that expected utility leads her to an unwanted outcome. However, the self-torturer can also be seen as presenting a difficulty for orthodox RCT, since it does not seem irrational for her to trade an imperceptible increase in pain for a perceptible increase in money (Andreou 2006). These non-threshold cases are more difficult than threshold cases, for expected utility issues a verdict in favor of the action (i.e. non-cooperation) that, when iterated enough times, will turn out to lead to a harmful aggregate outcome.

To summarize so far, expected utility is not a promising route for providing individual reasons for cooperation, as Nefsky and Budolfson have argued. But a moral or political obligation to advocate for governmental action (in the cases where this would be an effective intervention) is vulnerable to the inefficacy problem as well, so

Political-Not-Personal, as articulated by Sinnott-Armstrong and Kingston, is at least inconsistent. Moreover, both moral and instrumental reasons appear to counsel us against cooperating. Thus, the inefficacy problem is a serious problem on multiple dimensions.

Much of the debate on the inefficacy problem has focused on moral rather than instrumental reasons. The theoretical reason for this is that the inefficacy problem poses a problem for consequentialism. In addition, there are also practical explanations as to why we should lean on moral rather than instrumental reasons. As the discussions of the torturer cases show, it is instrumental reasoning that seems responsible for collective harm. Each person can argue that “it’s not my fault” (Sinnott-Armstrong 2005) because their individual contribution makes no difference to the outcome, and that it would be personally costly for them to change their behavior. In fact, we often hear Ordinary People make a version of this argument. Since instrumental reasons create the problem, we might think that instrumental reasons cannot get us out of it. That is the negative explanation – why *not* instrumental reasons. Therefore, focusing on instrumental reasons sharpens the problem.

The positive explanation – *why* moral reasons are at the focus of this debate – requires making some conjectures about the metaethical assumptions framing the

debate. One possible explanation is that moral reasons are assumed to have universal normative force; they are supposed to swamp instrumental reasons. If agent-relative, instrumental reasons create the problem, we need agent-neutral reasons, such as those provided by morality, to get us out of it. But if the above interlocutors are right about the inefficacy problem, as I think they are, morality doesn't seem to be fertile ground to harvest reasons for cooperating in collective action problems. Because the inefficacy problem is so difficult, Nefsky (2017) has turned her attention to undermining the importance of difference-making for grounding reasons for action.

The aim of this paper is to explore territory that has been left within the space of reasons. So in what follows, I will take a different strategy, focusing instead on temporally extended instrumental reasons. I will argue that the Ordinary Person has instrumental reasons, given by diachronic rationality and self-governance, for cooperating in collective action problems. This strategy allows us to bypass the debates about expected utility to reach a conclusion that applies even to the difficult non-threshold cases.

### **3.3. An intrapersonal, intertemporal solution**

The argument that follows builds on the core intuition that collective action failures and intrapersonal, intertemporal failures share the same overall structure: no one act makes a difference, but many acts together do, and enough instrumentally rational failures to cooperate will result in the failure to achieve the overall goal. But we have a good solution to this in intrapersonal cases, such as long-term projects. In such cases, we adopt commitments as a device to get over the inefficacy problem. I will argue that commitment can also help us get over the inefficacy problem in the collective action case. Making sense of commitment requires taking a diachronic perspective on rationality.

First, let's introduce two forms of instrumental rationality, synchronic and diachronic. An agent is instrumentally rational if and only if she intends the necessary means to her intended ends. Synchronic rationality concerns rationality at a time, whereas diachronic rationality concerns rationality over time (Bratman 2007), or constraints on intertemporal combinations of intentional states. Both forms of rationality require at least *means-end coherence* and *consistency*. To be coherent, the agent should intend what they believe to be the necessary means to their intended end. To be consistent, the agent should not intend A and B if they believe that A and B are not co-

possible. In addition to the norms of coherence and consistency, diachronic rationality also includes the norm of *stability of intention*: your commitments shouldn't change for arbitrary reasons (Bratman 2012). The connection between belief and action will become important later.

Synchronic and diachronic rationality can conflict, though not necessarily. Sometimes, we make choices that are rational at each point in time, but a series of such synchronically rational decisions leads to a diachronically irrational outcome. The conflict arises because of what each temporal perspective on rationality tells us to do. Synchronic rationality tells us that we should maximize expected utility at a given time. This is what the self-torturer does by opting to turn up the dial. Diachronic rationality rejects this claim; rather, we should do whatever is necessary for us to realize our long-term commitments, even if this sometimes requires not maximizing expected utility at a time. For an example of diachronic rationality in operation, consider Ulysses contracts. These precommitment devices help us "bind ourselves to the mast" by overriding present desires, thereby preventing us from succumbing to temptations that would lead us astray from our commitments. They can be as quotidian as self-imposed 25-minute "Pomodoro" time blocks, during which you are not supposed to do anything but your



prescribed task no matter what, or they can be as weighty as advance directives that allow people to spell out end-of-life decisions in case of dementia.

The conflict between synchronic and diachronic rationality is apparent in both threshold and non-threshold cases. Consider the following non-threshold case, *Fitness*. Suppose Hera sets herself the goal of being fit by the end of the year. This is a vague goal: there is clearly a difference between someone who is fit and someone who is not, but there is no precise threshold that we can point to to identify this difference.<sup>16</sup> To be fit, she must work out with some intensity most days of the week. But also suppose that occasionally missing a workout or two won't undermine her goal; it won't make a difference to her fitness level. And perhaps something more pressing comes up, or she simply doesn't have the desire to work out that day. If, at every choice point, Hera thinks, "I can miss this workout, because doing so won't make a difference to my fitness goal," she will not achieve her goal. There is something irrational about this chain of reasoning—after all, she is not taking any steps towards my long-term goal. There is also

---

<sup>16</sup> There is some debate about the possibility of a vague goal. I assume that we can aim for vague goals: to be fitter, morally better, etc. See Tenenbaum and Raffman (2012) for an argument that they are possible—we aim for them all the time—and that a theory of rational choice needs to make sense of this.

something extremely tempting and reasonable about it;<sup>17</sup> after all, no synchronically rational choice undermines her goal. Hence the conflict between the synchronic and diachronic.

Now consider a threshold case, *Running*. Hera sets a goal of running a 6-minute mile in one year's time, to be measured during a race. She either achieves it or she does not. She currently runs a 10-minute mile. In order to achieve her goal, she needs to run at least 70 times during the year.<sup>18</sup> But Hera is a busy and somewhat distractable person, and oftentimes, she would rather skip a run to do something else. Can synchronic rationality tell Hera to run enough to achieve her goal?<sup>19</sup> It may or it may not. Earlier in the year, it will often be synchronically rational for her to skip runs whenever she is busy, assuming that the value of the alternative activity is sufficiently high at the time, as long as she is still able to complete at least 70 runs. Perhaps she is involved with travel, or taking care of her grandparents, or giving a series of talks. It doesn't really matter for achieving her goal that she skips her run on these busy days. It will maximize

---

<sup>17</sup> The apparent reasonableness of this chain of reasoning underlies the temptation of procrastination (Andreou and White 2010).

<sup>18</sup> Obviously, this case greatly simplifies how exercise works.

<sup>19</sup> I thank an anonymous reviewer for pressing this objection and suggesting this example.

her expected utility on each day to do so, as long as the expected costs of skipping a run on a given day earlier in the year are lower than the expected costs of skipping a run on at least one day late in the year.

Predictably, however, synchronic rationality can sometimes lead Hera into a last-minute push during the last 70 days of the year to make up for lost time. Suppose Hera has to finish an important project in 9 months; but she need only reach her 6-minute mile goal in 12 months. The opportunity costs of running before the project is due are higher than after, so synchronic rationality can lead Hera to that last-minute scramble.<sup>20</sup> At that point, synchronic rationality will tell her that she must run each day in order to achieve her goal, even if opportunity costs are high, as long as the value of achieving her goal is higher than the opportunity costs of running. This is a bad situation to be in. During these last 70 days, there may be a day in which the opportunity costs of working out are so high (e.g. there is a personal emergency that must be attended to, she falls sick) that she is forced to skip a run, thereby eliminating the possibility of achieving her

---

<sup>20</sup> Thanks to an anonymous reviewer for suggesting this case.

goal.<sup>21</sup> Therefore, synchronic rationality can lead Hera to fail to achieve her goal through a series of individually expected-utility-maximizing decisions.

If Hera is diachronically rational, she will avoid putting herself in such a situation. She will look forward in time and plan accordingly, incorporating into any present decision the expectation that opportunity costs will likely be higher in the future if she skips today's run. Even if it maximizes her expected utility at one point in time to skip a run, she will not always allow that calculation to justify skipping. Her commitment structures her decisions. Thus, she will run more consistently throughout the entire year rather than leaving everything to the last 70 days.

In both the threshold and non-threshold cases, the conflict between synchronic and diachronic rationality means that expected utility will not always tell us to "cooperate" in service of the overarching goal. It may plump in favor of cooperation, but

---

<sup>21</sup> I am simplifying the cases such that the law of diminishing returns does not crop up, since that is not a feature of the inefficacy problem. Nevertheless, it's worth addressing what happens when we do take diminishing returns into account. One might object that most activity tends to follow the law of diminishing returns, so synchronic rationality will tell Hera to run early on when the marginal benefits are high. Assume that the activity in question does follow this law. Here, synchronic rationality is still insufficient — there is no guarantee that it will tell Hera to run when marginal benefits begin to level off. At that point in the curve, it's possible that the marginal costs of running will exceed the marginal benefits at each point, even if continuing to run is necessary to achieve her overarching goal. I thank an anonymous reviewer for prompting me to address this point.

it just as easily may not, depending on the contingent structure of preferences at each decision point. And what we are seeking is a theory that tells us to cooperate at a necessary number of decision points. Diachronic rationality, which tells us to structure our actions based on our commitments, doesn't allow us to fall into the trap of potentially self-defeating synchronic utility-maximization.

These individual cases of temporal conflict share the same structure as collective action problems. In both types of cases, one individual action makes no difference, but noncooperation for that reason will predictably lead to a suboptimal outcome. But commitment tells us to cooperate in the service of our long-term interests in the individual case. As I will show, commitment can perform the same role in collective cases. In fact, it already does in some fairly quotidian, artificial collective action problems, in which many of us ignore the pull of the inefficacy problem. Specifically, we don't buy this argument in the case of organized collective activities such as team sports; we think that we have reasons for action in such situations, despite the fact that the inefficacy problem applies. I will use this observation as an anchoring point to build an argument from coherence in favor of an instrumental reason for action.

Consider an intrapersonal, intertemporal dilemma built into an interpersonal dilemma, *Rowing* and *Gymnastics*. *Rowing* is a threshold case:

You are in a boat race with a 9-person crew. The race has a highly unequal prize structure: you either win and receive \$10,000, or you don't and receive nothing. You strongly desire to win. You are slightly behind the next crew, and from your limited vantage point, they seem to be pulling ahead with 500 meters to go. There is still a small chance that you might win. You are exhausted and in pain, and your strongest desire at this moment is to pull just a tiny bit less hard to alleviate your pain. After all, there are seven other rowers in your boat. Dialing down your effort would barely be perceptible, if at all. You have no assurance that everyone else is pulling as hard as possible. You know that maximum effort over the length of the race by at least four other rowers will be required for you to win, and if at least six other rowers are putting out maximum effort, then you will definitely win.

The inefficacy problem applies here. Even if it is synchronically rational for the boat as a whole to put in maximal effort at each stroke, it doesn't follow that it is synchronically rational for each individual to do so. As with the *Running* and *Fitness* cases, the cost of putting in maximum effort at a time is high, and the probability that that particular stroke will lead to the win is exceedingly low. The inefficacy of each stroke is magnified by the necessity of others also working hard: there are eight rowers in the boat, and each typically takes hundreds of strokes over the course of a race. If fewer than four or six rowers are cooperating, then your effort will be wasted relative to either of the two possible good outcomes. If more than four or six rowers are cooperating, then your effort is not required. So, in either of these scenarios, it seems rational for you to not put in maximal effort. That is, the chance that the expected utility of a stroke will turn out negative is higher than not in either scenario, so one should not exert maximal effort. Synchronic rationality therefore can tell each rower that it is fine to keep the pressure off

on each stroke. Noncooperation seems like the dominant strategy, giving rise to an interpersonal dilemma. Yet following this prescription will put everyone in a last-ditch situation similar to the one Hera faces in *Running* and will significantly raise the likelihood of a loss.

Next, consider a non-threshold case, *Gymnastics*:

You are on a team of eight gymnasts competing at a meet. For your team to win, each gymnast needs to acquire a certain number of points, and to do so, each of you will have to execute your routine successfully. But what counts as success is vague; no particular action results in success. Success is evaluated holistically and requires that the judges believe that the routine as a whole has been done well. To do well, you're required to execute a series of strenuous, potentially dangerous moves. The danger of a move increases with the effort, or power, put into it. You have no assurance that anyone else on your team will put in maximum effort to perform well. Maximum effort by at least four other gymnasts will be required for you to win, and if at least six other gymnasts are putting out maximum effort, then you will definitely win.<sup>22</sup>

Compared to *Rowing*, the inefficacy problem is heightened in *Gymnastics*. While in *Rowing*, it is merely possible or likely that expected utility will turn out negative, the vague structure of success in *Gymnastics* ensures that no move will make a difference to the outcome. But this narrowly synchronic form of reasoning seems inappropriate in both *Rowing* and *Gymnastics*. In collective activities such as sports and other group

---

<sup>22</sup> This case is constructed with a layperson's understanding of gymnastics, with apologies to those who know better.

activities, most people don't act in the way licensed by the inefficacy problem.<sup>23</sup> Most people will put in maximum effort at every choice point. We do so because we take commitments seriously. Participation in a collective endeavor entails committing to a common goal, and commitment overrides the synchronic, instrumentalist reasoning that would otherwise give rise to a collective action problem.

In these collective cases, even if an individual is certain or nearly certain that her individual cooperative action (or even all of her individual cooperative actions) will not make a difference to the outcome, commitment still requires her to cooperate. To see why, consider the epistemic situation of the agent. The agent has committed to a collective goal: a goal that is achievable only with collective action. And she must believe that it is achievable only with collective action, in order for the inefficacy problem to get a hold on her. If she did not believe this, then she would not be able to say that her individual cooperation doesn't matter, because it doesn't matter only if others' cooperation is necessary or sufficient to achieve the goal. So, if the agent believes that the goal can only be achieved through collective action, she knows that the expected utility

---

<sup>23</sup> See Nguyen (2019) for an overview of the implications of game-playing for practical rationality.



of her action will either be negative or zero (in the non-threshold case) or be very low to negative (in the threshold case). Either way, expected utility tells her that she shouldn't cooperate, because doing so will certainly or nearly certainly make no difference to the outcome. Should she follow expected utility's dictates?

She should not. If she is committed to her goal, she must take the actions necessary to achieve it. Here, the relevant necessary action is the *set* of individual actions that together bring about that outcome. The agent's action may not *itself* be necessary within that set, but because commitment implies action, she is required to participate in that set of actions, even if her action itself ends up being inefficacious. And because she has committed, the question "why should I, in particular, cooperate, when doing so doesn't make a difference?" cannot arise for this agent. This question is puzzling, for it is like saying: you are committed to your team winning, you know that this depends on the cooperation of everyone else, your action by itself is certain or not to make a difference, or is very unlikely to make a difference, and your action alone doesn't matter—therefore you'll just skulk around on the sidelines. Clearly, this line of reasoning is incoherent. An agent who employs this line of reasoning does not seem genuinely committed, and if she is committed, she has not thought through what commitment requires.

There are a few disanalogies to tidy up between the artificial cases and the real collective action problems. First, the cases illustrate outcomes that require a series of actions to achieve, but real cases are sometimes one-shot. It is a short step to generalize the argument to one-shot situations. In one-shot situations, it is imperative that one cooperate, or one has failed to participate in the set of actions necessary to achieve the goal.

Second, there is an importance difference between the real-world cases and artificial sports cases. The real cases are not organized, artificial activities with arbitrary rules. We find ourselves in collective action problems not of our own volition, and not because we are deliberately acting in concert with others, but due to social, political, and economic structures. This disanalogy naturally raises the question: even if commitment gives us a reason for action in artificial collective action problems, why does this reason apply in general collective action problems?

Return to the Ordinary Person, whose psychology I described in Section 1. The Ordinary Person is the target of these arguments. She is sincerely committed towards bringing about the good. It's just that the inefficacy problem blocks her from taking action. When the agent makes a commitment to bringing about the good, her situation is similar to that of athletes who have committed to winning as part of a team. Thus, she

enters into a *Rowing*- or *Gymnastics*-like situation with actors who have made a similar commitment, even if action is diffuse across time and space, and even if she never forms actual relationships with the other actors. In these situations, we ought to act in spite of the inefficacy problem.

In what cases does commitment require cooperation? My account is similar to Nefsky's (2017, 2760-2764), who argues that an individually inefficacious action is required as long as the outcome that it can help to bring about is possible. There is a distinction between two types of situations where individual cooperation is certainly or almost-certainly inefficacious: cases in which the desired outcome is certain or impossible, and cases in which the outcome has yet to be determined and is possible. Commitment does not require (knowingly) futile cooperation, only cooperation in cases where the desired outcome is a live possibility. In cases where the outcome is certain, and it is therefore certain that each individual act of cooperation will not make a difference to the outcome, the individual is not required to cooperate. In *Rowing*, for example, if the other boat is just about to cross the finish line, but your boat is still 200 meters away from it, then the outcome is settled. Nothing short of an act of God would help you win. In this case, you are no longer required to exert maximum effort. Therefore, commitment requires cooperation only if it is uncertain whether the outcome

can be attained, and even if it is certain or nearly certain that each individual act of cooperation will not make a difference to the outcome.

My account can be thought of as a synthesis and development of recent work in practical rationality that challenges, or at least sidesteps, the prescriptions of orthodox RCT. Sergio Tenenbaum and Diana Raffman (2012) argue that to make sense of vague projects, an extended temporal perspective is needed to overcome the prescriptions of orthodox rational choice theory, which restricts our pursuits to well-defined ends. However, they argue that planning is not necessary for us to carry out vague projects; we simply need to carry out a sufficient number of the relevant actions. This account does not consider the opportunity costs associated with a lack of planning. I show that the concept of commitment requires us to take an extended perspective to both vague and well-defined ends, and its implicit planning helps us act prudently to ensure that the desired outcome will be achieved. Margaret Gilbert (2006) argues that “joint commitment”, wherein two or more parties collectively intend to perform an action as a body, can provide sufficient reason to cooperate in collective action problems. On my account, an individual commitment is enough to provide this reason, and it can do so even if individuals do not intend to act together. Finally, Nefsky (2017) argues that one has a moral reason to cooperate if doing so helps to bring about a good outcome, even if

one's action doesn't make a difference, as long as one's action makes a non-superfluous causal contribution. My account buttresses Nefsky's view by showing that moral reasons for cooperation in collective action problems are structurally similar to, and continuous with, our purely instrumental reasons for action in non-moral situations. Commitment allows us to get the same result without any recourse to moral language. It also gives us the resources to overcome the instrumental, self-regarding calculations that stand in the way of cooperating, even when we know we have moral reasons to do so.

### ***3.4. Why stick to your commitments?***

I have argued that commitment implies action, regardless of the efficacy of each individual action towards the goal. We can use the resources of self-governance to strengthen this instrumental reason for action. A self-governing agent is governed by her commitments and values, not by her occurrent desires (Bratman 2003). Harry Frankfurt's (1971) unwilling addict is the classic example of a non-self-governing agent – one who acts on his first-order desire to smoke, but against his second-order desire (perhaps a commitment) to cease smoking. If self-governance is intrinsically valuable, as Bratman (2018) has argued, it gives each person a reason to act in accordance with their commitment, regardless of how strong a commitment is. However, skeptics such as Niko Kolodny (2008) have pushed back against the intrinsic value of self-governance. I

do not have the space to defend the intrinsic value of self-governance in this paper, nor do I want my account to hang on the resolution to this problem. Nor do I want to further restrict the scope of this reason to act, so that it only applies to those who have an intrinsic desire for self-governance. To leverage self-governance in favor of action, I will therefore appeal to the personal value to the Ordinary Person of being a person who is committed to the bringing about the good.

In the above artificial cases, we appealed to the structure of commitment to derive a reason for cooperation in spite of the inefficacy problem. But the reason for cooperation seems overriding only if the agent is isolated from other conflicting commitments, as happens in games. In real life, however, the Ordinary Person frequently finds conflicts between the commitment to the good and other commitments. Return briefly to the examples of climate change and segregation that we began with. Not all greenhouse gas emissions are caused by frivolous pursuits such as joyguzzling (Sinnott-Armstrong 2005); perhaps one needs to frequently take long-distance flights to visit loved ones across the country because one is committed to maintaining these relationships. Not all segregation is caused by racial animus or indifference to racial justice; much of it stems from the desire to give a good life to one's children by moving to a better neighborhood with more resources. Call these "individual commitments", for

lack of a better term. When commitments conflict, we need to give up one thing in favor of another, even if we do so regrettably. This raises the question: why prioritize the commitment to this particular good – the commitment to alleviating collective harm – rather than the commitment to other important projects and relationships?

To justify the priority of this commitment, we need to delve into the requirements of self-governance. Self-governance requires that an agent adhere to the practical norms of consistency, coherence, and stability of intention with respect to her commitments. I will address each of these in turn. Begin with the norm of consistency, which gives us a weak response to this problem. According to *consistency*, a rational agent should not intend A and B if she believes that A and B are not co-possible (Bratman 2012, 73). So she should simply drop either A or B. However, a rational agent with sufficient information about the collective action problem should, presumably, also have knowledge of the inefficacy problem. The inefficacy problem opens room for the possibility that one can intend to commit to the good, yet also make other seemingly conflicting commitments. After all, if my act of cooperation won't bring about the good, then it seems that I can rationally intend and act on commitments that undermine the good, without undermining my commitment to it. And we can't appeal to the structure of the commitment to the good to override the inefficacy problem, because what is at

issue is why this particular commitment should come out ahead. Because of the inefficacy problem, these commitments are not strictly conflicting. Both the good and ostensibly conflicting commitments are co-possible precisely because of the inefficacy problem. There is a possible world in which I have a high-emissions lifestyle, but due to the sacrifices of others, this world is ecologically sustainable. So it seems that the norm of consistency does not deliver a resolution in favor of action. It only makes the problem more vivid.

The norm of instrumental coherence seems more promising. According to the norm of *coherence*, a rational agent intends what they believe to be the necessary means to their intended end (Bratman 2012, 76). As I've argued, if an intended end is a commitment that can only be collectively realized, then the rational agent must intend to participate in the set of actions that together constitute a necessary means to that end, or the agent can't be said to have that commitment. Coherence then presents the following dilemma. On the first horn, if the agent is committed to bringing about the collective good, then she ought to intend to participate in the set of actions that together constitute a necessary means to that end. This will often require personal sacrifice, sometimes to a significant degree. On the second horn – call this the “soft horn”, because it is less



demanding – if the agent chooses not to intend the necessary means to bring about the collective good, then she has several options.

One is to weaken the strength of her commitment.<sup>24</sup> One can be strongly, moderately, or weakly committed to something. A strongly committed person would indeed be required to undertake the demanding actions required of them, at the expense of self-governance, whereas a less committed person has more slack in their web of reasons. They can do less without compromising their self-governance, simply because they haven't imposed certain rational requirements on themselves by taking on a stronger degree of commitment. But weakening the strength of one's commitment doesn't discharge one from intending at least *some* of actions that constitute part of the necessary means to realizing that commitment. Commitment still requires some degree of action.

Another option is to drop this commitment to the level of a desire, which does not entail adhering to the same rational norms. Nevertheless, a desire gives the agent a *prima facie* reason to satisfy it (Schroeder 2007), so choosing this “soft horn” of the

---

<sup>24</sup> This response may help to address a separate objection that my view is too demanding and would force us to be “moral saints” (Wolf 1982) on pain of irrationality. I thank Leif Wenar for this objection.

dilemma does not entirely get her out of intending some of the actions that help to constitute the necessary means to the relevant end. To get out of intending any of the actions that help to constitute the necessary means, the agent should simply drop any desire or commitment to the collective good.

Both these horns of the dilemma allow a rational agent to preserve coherence. Coherence will not tell the agent to plump for one horn or the other. It does, however, make the cost of inaction vivid. That is, inaction does not come for free, licensed by instrumental rationality. It comes at the cost of *being a person who is committed to the bringing about the good*. It is up to the agent to decide for herself whether that is a price worth paying. And if the agent does not choose either one of these horns, but instead retains the desire for or commitment to the good without intending at least some of the actions that constitute the necessary means to that good, then she is practically irrational on that account.

To sum up, I have developed Bratman's planning framework to excavate a reason for action in collective impact situations. Agents who act in these situations, instead of succumbing to the logic of the inefficacy problem, are acting in the service of their own self-governance – assuming that their commitment to bringing about the good takes priority in their scheme of commitments.

But one might think that plans and commitments change, and that they actually ought to change as we learn more about ourselves and the world (Millgram 2015). The norm of stability of intention thus seems stubborn at best, if not irrational. Why should we stick to our commitments? In particular, why should we stick to our commitment to collective goods, given that we know they are so unlikely to be realized?

I do not have a general answer to the deep theoretical objection about the norm of stability. However, we can again appeal to the specific content of the commitment to the good to buttress the applicability of stability to this particular intention. A better world just *is* something that we shouldn't give up our commitment to. Perhaps this commitment can only be justified intrinsically by the analyticity of the good. But more plausibly, and more in line with the instrumentalist commitments developed in this paper, we should be committed to this good because many of our ordinary desires and commitments cannot be realized if collective harms such as segregation and climate change continue apace. The commitment to the good, then, is unlike other personal commitments (such as the commitment to stick with philosophy no matter what), which should be more amenable to revision due to the fact that less hangs on them. We simply may not be able to realize many of our other desires and commitments if we do not

alleviate collective harms, so we have an instrumental reason to cooperate in order to redress them.

### **3.6. *The content of reasons to cooperate***

What is the *content* of reasons to cooperate in social dilemmas? By itself, my account is silent on the content of such reasons. We need to fill in the theory with substantive views from other areas of philosophy and social science, such as the literature on the right to association (Chapter 2) and the literature on the causes of sorting (Chapter 1).

Taken together with those substantive views, my solution to collective action problems does not imply that those who have access to social and economic privilege ought to seek out, befriend, date, or marry those who lack such privilege simply because they are less privileged. Doing so would be demeaning to the recipients of such social favors. Nobody wants to be sought out for a substantive relationship *because* someone else pities them for being excluded.

My account does, however, imply that individuals have reasons to offer associational opportunities to others, even when doing so is personally costly. In the movie *Forrest Gump*, the other children on the bus look askance at Jenny when she lets Forrest, who is autistic and physically disabled, sit next to her. What I'm suggesting is

that we all have reasons to be like Jenny – we all have reasons to offer associational opportunities to others, costly as this may be. But as I have argued in Chapter 2, others do not thereby gain the right to our association, nor do we have an obligation to continue to associate with them if doing so is burdensome. This reason falls squarely in the realm of personal morality.

But to act socially in ways that promote justice, it is not enough for individuals to offer associational opportunities to others. One reason that it is costly for individuals to offer associational opportunities to others is that they are locked into social and economic structures that discourage such opportunities. Other structures would permit or foster more social opportunities. Return to Andy in the *segregated people dilemma* that I started with. Andy decides not to challenge the segregated status quo in part because the value of his house depends on the fact that he lives in an Andy-only neighborhood. Further suppose that these property values are maintained by restrictive land-use regulations, such as anti-density zoning, which artificially drive up property values by lowering supply. Andy has reasons, derived from his interest in social justice, to act politically as well. For example, he might join a local group that seeks to encourage the development of affordable housing such as apartment complexes within the neighborhood. At the very least, Andy cannot coherently support policies that restrict associational opportunities,

such as by acting as a “neighborhood defender” who blocks affordable housing from being built in his neighborhood (Levine Einstein et al 2020, Dougherty 2020). This is true even if there are other Andy-only neighborhoods around that could absorb this burden, and even if Andy’s successful advocacy would undermine his other prudential interests.

### **3.7. Conclusion**

I have argued that social dilemmas, such as the one presented by the problem of sorting, have the structure of a collective action problem. Collective action problems pose a challenge for theories of morality and instrumental rationality. It seems that any individual can say that they have no moral reason, to act, given that their action will not make a difference to the overall outcome. Instrumental reasons seem even less promising than moral reasons in this context: it is certainly costly for individuals to act, yet it seems highly improbable that their action will make a difference in the desired way.

We should look toward the space of instrumental, diachronic reasons for a reason to cooperate in collective action problems. The Ordinary Person’s psychology, which is thought to give rise to the inefficacy problem, contains the reason for action. What I have tried to do in this paper is to show that the Ordinary Person is irrational if she is committed to the good, but fails to cooperate because she allows the inefficacy

problem to override her commitments. I have built up an instrumental reason for such a person to act in collective action problems. I argued that collective action problems have a similar structure to intrapersonal, intertemporal commitments, so individual commitment to an outcome that requires collective action provides sufficient reason for cooperation. Then, I argued that the requirements of self-governance put a price on noncooperation. Once we examine the requirements of our commitment to a just and good world, we can find a reason to cooperate within our divided selves.

Let me end with some caveats. I have not attempted to generalize my argument to apply to those individuals without the relevant commitment. The account also applies more strongly to collective *harm* situations rather than collective impact situations in general. Despite these qualifications, I hope that I have at least given an instrumental reason to act to those for whom the inefficacy problem is a true problem. The details of this account will be left to future work. The point, for now, is that inaction in collective action problems doesn't come for free.

## **4 The ecology of freedom of association**

### **4.1 Introduction**

Freedom of association is “the right to choose the society most acceptable to us,” in John Stuart Mill’s (1859) canonical formulation. It is both a negative claim-right against the state’s interference on our associations, and a moral permission to associate with whomever one wishes (conditional on the other’s reciprocity). Freedom of association intuitively seems like a desirable cornerstone of liberal society. A liberal state should not interfere in private matters, such as our personal choices of whom to socialize with, befriend, live among, date, and marry. However, freedom of association also implies the right not to associate with others. In other words, freedom of association entails that we are free to exclude others from our company, whether this is through de facto means such as ignoring, shunning, or not choosing certain people, or by enacting club-like rules of membership that prevent certain people from joining associations. It can therefore lead to injustices when we sort ourselves into clearly segregated groups.

Segregation raises a more general tension between liberal society’s ideals of social inclusivity and equality, and the fact that freedom of association licenses exclusion and its attendant inequalities. This tension is not only a theoretical one for liberalism. It is especially salient and urgent for citizens and policymakers when we attempt to



confront the persistence of social sorting. For example, racial residential segregation has never simply been the result of structural factors such as economic inequality, redlining, and unjust policies. It has also been, and continues to be, caused by individual white flight (Shertzer & Walsh 2019, Kruse 2013), support for local zoning laws that restrict density (Rothwell & Massey 2009), and the dynamics of social networks (Hofstra et al 2017, Kossinets & Watts 2009, McPherson et al 2001). This patchwork of causal factors can be seen as exercises of freedom of association, and normatively justified as such. Yet segregation undermines justice. How should we negotiate the apparent tension between freedom of association and the demands of justice?

According to the dominant liberal response, basic liberties such as freedom of association take priority over the demands of justice (Rawls 1993, 313, 341). A regrettable consequence of the priority of basic liberties is that we may have to live with some degree of segregation. While much attention has been paid to the Millian idea that the state should not restrict our liberty, especially in the context of expression and association, another Millian idea – that social forces can restrict our liberty as much as the state can – has been relatively neglected (Mill 1859, Ch.4).

The aim of this paper is to explore the second Millian idea in the context of freedom of association, understood as a basic liberty in a broad sense rather than in

specific interpretations carved out in American First Amendment doctrine. I argue that the apparent tension between freedom of association and justice rests on an overly narrow, negative conception of freedom of association that conceives of its only obstacles as identifiable agents. In place of the negative conception, I develop a positive, ecological conception of freedom of association that foregrounds the social conditions necessary for freedom. Under certain social conditions that interfere with choice, such as socially sorted environments, we lack meaningful freedom of association. A series of free associational choices can, in the aggregate, undermine freedom of association if they result in substantial segregation. I argue that the ecological conception of freedom is truer to the spirit of individual autonomy that liberalism rests on. The upshot is that freedom of association requires interventions on segregation.

The paper proceeds as follows. In Section 4.2, I motivate the tension between freedom of association and justice in greater detail. In Section 4.3, I give an overview of the orthodox view concerning this tension. I show that it assumes a conception of freedom of association on which the only obstacles to freedom are agents such as the state. In Section 4.4, I present the ecological account of freedom of association as an alternative. On this view, we have freedom of association to the extent that the social environment affords us choices of people to associate with. Sorted social environments

therefore constrain our freedom of association. Section 4.5 responds to some objections, and Section 4.6 addresses some implications. Section 4.7 concludes. My upshot novel is that freedom and justice travel together much farther than previously thought: just and diverse social environments are necessary for freedom of association to be meaningful.

## **4.2 Background: homophily and racial segregation**

A cursory observation of almost any social and spatial environment confirms that people predictably cluster into groups along the lines of race, class, ethnicity, religion, sexual orientation, political ideology, lifestyle preferences, and other dimensions, almost as if by some principle of self-organization. Our neighbors, friends, romantic partners, colleagues are more likely than not to look like us, talk and act like us, and have similar socioeconomic backgrounds. These patterns recur across different contexts, from high school cafeterias and neighborhood bars to the reception area at academic conferences. Sociologists call these patterns *homophily*: “the tendency of people with similar characteristics to interact with each other” (McPherson et al 2001) more than a baseline rate of random association would predict.

Racial segregation is one consequence of homophily. Henceforth, I focus on segregation because it is likely the most morally problematic form of sorting, given its history and consequences. It is both cause and consequence of negative stereotypes and

group inequality (Krysan & Crowder 2017). For example, economist Raj Chetty and colleagues find that upward mobility is greatly impacted by the neighborhood in which a child grows up – and neighborhoods in the United States are highly segregated by race (Chetty et al 2014).

What explains racial segregation? Myriad structural factors, such as redlining, neighborhood covenants (Jones-Correa 2000, Massey & Denton 1993), intentionally discriminatory laws (Rothstein 2017), economic inequality (Reardon et al 2015), and “steering” by real-estate agents (Turner et al 2012) are crucial for explaining residential segregation. In contrast, the personal choices of individual actors in this story, such as self-segregation and social closure, have been relatively neglected, even though these individual choices are responsible for a significant portion of segregation as well (Boustan 2013, Logan & Parman 2015, Shertzer & Walsh 2019). Consider the history of integration in Atlanta. Historian Kevin Kruse shows that “white flight” in Atlanta – the movement of white families to the suburbs in response to the civil rights movement – made the integration of schools, neighborhoods, and public spaces extremely difficult to achieve. It is worth noting that whites explicitly cited their freedom of association to justify their resistance to integration. For example, in 1961, Sandra Melkild, a white high school student in Atlanta, sought to transfer from a desegregating high school to an all-

white one, citing her right to “freedom of association”. A year earlier, Thomas J. Wesley Jr., the leader of the Metropolitan Association for Segregated Education, had also invoked the right to freedom of association to resist the desegregation of schools (Kruse 2013).

Racial segregation also persists within integrated institutions, such universities, clubs, and social networks. Crucially, the tension between freedom of association and justice goes right to the heart of our personal lives. Freedom of association includes the freedom of intimate association, “a close and familiar personal relationship with another that is in some significant way comparable to a marriage or family relationship” (Karst 1980). There is growing evidence that assortative patterns in friendship, dating, and marriage maintain social segregation and inequality (Schwartz 2013). Therefore, segregation in the intimate realm undermines justice in a similar way that residential and educational segregation does. In this most private realm, personal choices become even more explanatorily indispensable.

None of this is to say that personal choices are sufficient to explain the full extent of real-world segregation; that assertion would obviously be false. Nor am I claiming that they are the most important factor. My empirical claims are relatively modest: there is some degree of segregation that cannot easily be chalked up to formal structural

factors, and that seems better explained by homophilous social interactions in the informal spheres of life. But this modest empirical claim has weighty normative implications. It brings to light the importance of examining the limits of freedom of association, which gives cover to *all* the above choices regardless of their motivational basis. It does not matter whether we choose our associates based on an innocuous dimension of affinity, such as taste in music, or whether we choose them based on more problematic motivations such as racial preference. Conversely, it does not matter whether we exclude some people from our associations because we find them unpleasant, or because we find their class-based mannerisms repulsive. Freedom of association protects all of these choices.

Therefore, there is an inescapable tension between freedom of association and justice. As philosopher Erin Kelly (2017) puts it,

The challenge is that informal social relations and networks, which are mechanisms of inclusion and exclusion, lie largely in the domain of personal choice and freedom - a domain that is and should be protected by law. An effective approach to countering historically rooted injustice threatens to run up against reasonable personal liberty (p.90).

Yet scholars have not spent much time has been spent dwelling on this challenge. The prevailing view among philosophers and legal theorists is that freedom of association is an inviolable right, regardless of its regrettable consequences for justice. As we shall see next, scholars who disagree on much else agree on this principle.

### **4.3 The priority of freedom of association**

The priority of freedom of association can be traced to Mill and later, to Rawls, whose priority of the basic liberties carves out a special sphere for freedom of association. As Rawls writes in *Political Liberalism* (1993): "freedom of association is required to give effect to liberty of conscience; for unless we are at liberty to associate with other like-minded citizens, the exercise of liberty of conscience is denied. These two basic liberties go in tandem" (p.313).

This special sphere continues to be vigorously defended by classical liberal philosophers and legal theorists, such as Christopher Heath Wellman (2008), Loren Lomasky (2008), and Richard Epstein (2008), who argue that freedom of association trumps concerns of justice. Wellman argues that freedom of association provides a prima facie case for citizens of a country to restrict immigration policy, even against a backdrop of global injustice. Lomasky's "paradox of association" argues that there is a presumptive case in favor of the right to dissociate, which cannot be overridden by the right to associate. Social policies that promote the right to associate over the right to dissociate, such as affirmative action, face a high if not insurmountable burden of justification. Epstein perhaps comes down strongest on the side of freedom of

association. On his view, freedom of association should not be hindered even by antidiscrimination law in the commercial realm.

Even philosophers whose main concern is injustice also share this reluctance to challenge freedom of association. Stuart White (1997) argues that the freedom of association is illegitimate to the extent that it damages the “opportunity interests” of some, e.g. their ability to fairly access income and other goods. However, he carves out a “special right of intimate exclusion,” reasoning that:

our interest in freedom of intimate association, so understood, is arguably so basic to the free development and expression of ethical personality that exclusion rules which protect this interest have a strong presumption of legitimacy even if they also damage others' opportunity interests (p.386).

Elizabeth Anderson (2010) argues that segregation is a problem for justice, for it leads to social closure and the exclusion of marginalized groups from networks of opportunity. However, Anderson concedes that individuals have the right to freedom of association in private life. She sidesteps the tension between freedom of association and justice, arguing that it is a “morally innocent” cause that does not explain the extent of segregation in the contemporary United States – so we need not pay much attention to it (Anderson 2010, 70-71). Like Anderson and White, Andres Moles (2014) heavily criticizes racist and sexist associations on the grounds that “it is more difficult to exercise our moral powers in a society in which racist associations are protected, and this gives



us weighty reasons to interfere with associations so that their existence has positive rather than negative effects" (p.94). On the spectrum of views on freedom of association, his is perhaps the most radically permissive of interference on freedom of association. However, Moles also carves out "brackets of privacy" for unreasonable intimate associations, such as racist friendships, due to their "spontaneous and informal nature" (p.99).

Overall, then, the dilemma between freedom of association and justice is acknowledged to be morally troubling across the spectrum of political views. Nevertheless, the consensus is that at least the more intimate forms of freedom of association should be protected. Call this the *orthodox view*. Only two voices currently stand outside this consensus, and both pay attention to the special harms of exclusion. Kimberley Brownlee (2015) argues that "the impact of isolation is immediate and extreme", and "loneliness is like hunger, fear, or pain" (p.276). Thus, the harms of exclusion should lead us to question the doctrine that freedom of association is fundamental and inviolable, for "our permissions not to associate are hostage to numerous constraints such as necessity, the type of association, burdensomeness, pre-existing commitments and collective responsibility" (p.272). Similarly, political theorist Sonu Bedi (2015) argues that sexual racism, the practice of "prioritizing an individual as

a possible romantic, intimate partner on account of their race in a way that reinforces extant racial hierarchy or stereotypes," is morally objectionable, for it "limits an opportunity that is a social primary good and a capability central to human dignity," (p.998) in Rawls's and Martha Nussbaum's respective terms.

Brownlee and Bedi's arguments have implications for the treatment of groups who tend to be excluded from intimate associations, such as racial minorities, transgender individuals, disabled individuals, and the dependent elderly. On their accounts, exclusion and isolation are moral wrongs. We are not morally entitled to wrong individuals in this way, even if we do so in the course of exercising our freedom of association. Note that Brownlee and Bedi's arguments are pitched at the level of personal morality and are framed in terms of moral permissions. By contrast, the orthodox view is framed in terms of rights against the state.

What accounts for the priority of freedom of association? One explanation is that it protects a valuable sphere of life from interference. As Isaiah Berlin (1969) puts this idea, "some portion of human existence must remain independent of the sphere of social control. To invade that preserve, however small, would be despotism." Even if this "small preserve" of human existence ends up unintentionally causing social injustice, it should remain beyond the purview of justice.

My aim in the rest of this paper is to defuse the tension between freedom of association and justice by turning our attention to the social conditions that can undermine freedom of association. I contend that an ecological conception of freedom of association is needed to do this work. My general approach can be thought of as an extension of Moles's idea that we consider the spillover effects of freedom of association on the social environment, and in support of Brownlee's claim that associative freedom requires certain preconditions to be meaningful (Brownlee 2015, 268).

#### ***4.4 An ecological account of freedom of association***

An ecological account of freedom of association foregrounds the social context in which our associational choices take place. For freedom of association to be meaningful, we need to have multiple choices open to us. These choices are presented to us by our social environment. Different costs are imposed on different choices, and these choices are mediated by our perspectives. Thinking of freedom of association in this ecological sense not only mitigates the tension between freedom of association and justice; it is also better able to account for the importance of negative freedom of association. Therefore, we should first assess whether conditions are in place for our choice of associations to be meaningful. If these conditions are not in place, the importance of negative freedom of association correspondingly diminishes. And, as I argue, these conditions are not in

place in completely segregated environments. Some degree of diversity is necessary for freedom of association.

Before launching into the ecological conception, I should briefly explain why it is not merely an alternative conception of freedom of association, but a *preferable* alternative, even on the orthodox view's own terms. The orthodox view is intuitively appealing because it prioritizes our ability to make choices – an essential part of leading an autonomous life. Freedom of association in this sense is a negative claim-right against the state and other actors, and this claim-right prohibits them from interfering in our associations. It is also a moral permission that allows us to associate with whomever we wish; we have no duty not to associate with whomever we wish, and others have no claim on us that we should associate with them.

Although the orthodox view focuses on intentional actions by the state and other agents, it is external interference by agents per se rather than external interference in any particular form that is troublesome for proponents of negative liberty. Berlin calls this “social control” (1969, 123); Mill refers to this as the “despotism of custom” (1859, Ch.3). If a hurricane interferes with our ability to associate with others, that is not an infringement on freedom, for there is no agency involved. But if social facts, such as conventions, norms, and ideologies that are upheld by individuals interfere with our

ability to freely associate with others, then this should be cause for concern on the orthodox view. As Mill famously points out, social forces and structures can restrict our freedom just as much as direct state or interpersonal interference from identifiable actors can (Mill 1859, Ch.1). Yet this point seems to have been neglected by its proponents.

Consider the story of Hester Prynne, the main character in Nathaniel Hawthorne's *Scarlet Letter*. Prynne has a child out of wedlock. In a Puritan society, this is grounds for social exclusion, and she is forced to wear a scarlet letter that marks her out for ostracism. Prynne has freedom of association in the orthodox sense because no one interferes with her associational choices. They just won't associate *with* her, rendering her negative freedom utterly trivial. This example shows that freedom of association can be limited not just by direct interference from other actors, but by social structures such as norms and ideology. It might seem strange to say that Prynne's exclusion is caused by structures *rather than* by individuals, because it is clearly individuals who are ignoring her. To clarify, this claim isn't meant to be ontological; there is no structure independent of the actions and mental states of individuals. Rather, what I mean is that the social norm "ostracize anyone who wears a scarlet letter" is robust across different microlevel psychologies. It doesn't matter whether people actually think Prynne deserves to be ostracized for having a child out of wedlock, or whether they sympathize with her and

think that this punishment is too harsh. The fact that there is such a norm creates incentives for individuals to shun Prynne simply because that is what everyone else is doing, and norms both rely on and create strong conformity pressures (Bicchieri 2005). Nobody wants to be the odd one out, because they would incur costs by associating with Prynne. In this way, a norm has independent explanatory power, because it generates certain behavioral patterns regardless of the actual beliefs or desires of the individuals who are enacting those patterns (cf. Jackson & Pettit 1992).

A more robust conception of freedom of association should orient our attention toward these social structures. This orientation is compatible with the negative conception insofar as social structures interfere with our ability to form relationships with others, but goes beyond it in identifying non-agential sources of interference on freedom. The ecological conception thus aims to show that freedom of association is a distinctively social achievement – contra the individualistic interpretation offered by the orthodox view – that requires certain social structures to be in place. It does so without relinquishing the importance of individual choice. With this on the table, let's proceed to outlining three commonly accepted conditions for freedom.

#### 4.4.1 Three conditions for freedom

To be free, one must be able to choose; to choose in a meaningful sense, one must first have a menu of options to choose from. How many options? While there needn't be an expansive number of options present, there should be at least more than one. Call this the *Choice Condition*. If there is only one option available, one has not made a genuine choice in selecting that option, even if that option is in fact what one prefers (Berlin 1969). This is the most minimal condition for genuine choice.

Additionally, options must be feasible. For example, feminist philosophers have argued that the formal option to work in a high-powered career is not a substantive option for women if it comes at too high a personal cost, compared to working a less demanding job. Many women are therefore not genuinely free to make the choice to have both family and a demanding career (Cudd 2006, Haslanger 2016). Thus, to make two options count as such, they need to cost roughly the same amount, or at least have costs within a reasonable range. Call this the *Reasonable Costs Condition*.

Perhaps more fundamental than either the Choice or Reasonable Costs Conditions is the *Salience Condition*: other things equal, if we see more objectively available options as genuine options for us, then we have a wider range of choices. But before something can count as an option, it needs to be perceived as such. To illustrate,

consider: what counts as food for humans? The set of options that are objectively available for us, in the minimal sense of being edible and reasonably nutritious, is much wider than the set of options that any group of humans will count as food. Insects can be decent sources of nutrition, but people in Western countries generally do not consider them to be food. Some Americans hunt and eat small game such as squirrels and rabbits, while others will only eat livestock such as cattle and pigs. Vegans refuse to consume any animal products at all due to their ethical beliefs. As Ryan Muldoon (2017) puts it:

...the options people perceive and the evidence they use to make choices are mediated by perspectives...Our choices are not over all of the objectively available options, but over the options that we can see. The evidence we marshal to support our choices is not the full set of options, but that which we recognize as salient. (p.262)

Together, these three conditions embody the core idea that freedom requires choice, which in turn requires a menu of reasonably available options. I am not claiming that freedom always increases with the number of options available. At some point, the “paradox of choice” (Schwartz 2004) occurs; more choices induce anxiety and paralysis, compromising our freedom of choice. All I assume is that freedom of choice requires at least more two options, though it is difficult to pinpoint exactly how many options are necessary for maximal freedom. Theorists from the liberal, neorepublican, and feminist traditions generally agree on these conditions for meaningful choice, though they may



not formulate them in the same way. Disagreements about the right theory of freedom concern the sufficiency of choice,<sup>25</sup> but never its necessity.

More succinctly, the ecological conception of freedom of association states that:

An agent has freedom of association if and only if they have access to a menu of objectively available associational options (the Choice Condition), and they perceive some of these associational options as genuine options for them (the Salience Condition), and at least two of these associational options have costs within a reasonable range (the Reasonable Costs Condition).

We lack freedom of association wherever the choice conditions are not satisfied, as they tend not to be in sorted environments. Sorted environments are characterized by rigid boundaries between insiders and outsiders, enforced by social dynamics and economic accessibility factors. This characterizes many actual residential neighborhoods and social networks. These dynamics are not inherent to the concept of segregation itself. It is possible for a segregated environment to have porous boundaries between itself and the broader social world are relatively porous, allowing insiders and outsiders to move freely without penalty. But in our nonideal world, segregation undermines ecological freedom of association.

---

<sup>25</sup> See Philip Pettit (2011) for a nondomination condition and John Lawless (2018) for an argument that freedom requires the social conditions of respect to be in place.

To see this, first consider the Choice Condition as applied to search environments. A search environment is the space in which we form associations. It is often physical; for example, neighborhoods, places of worship, schools, workplaces, bars, and parks are typical physical search environments. With the advent of digital technology, our search environments now increasingly transcend the confines of the physical. Through social media, internet forums, and dating apps, we are able to seek out and connect with people beyond the immediate physical environment. Yet in both physical and digital life, we tend to inhabit “filter bubbles” (Pariser 2012) consisting of people who are roughly similar to us demographically and politically – the result of both personal preference and structures such as algorithms and socioeconomic policy. These sorted search environments limit the diversity of people that we have the potential to encounter. There may be people one wishes to associate with, but who are not available in the search environment. And even if one is perfectly happy to associate with a limited subset of people, the alignment of preference with one’s actual options is not sufficient for genuine choice, as the literature on adaptive preferences makes clear (Elster 2016). Thus, fully sorted environments, in which one encounters individuals only of one type, undermine the Choice Condition.

Even among the associational options that we do have, some carry unreasonable costs. Some of these social costs have to do with search frictions: it takes more effort and social wherewithal to find and connect with someone outside the boundaries of a sorted environment. Other social costs stem from stigma. Some associations are unblinkingly accepted; others are stigmatized. In such a situation, our options are not equal. To illustrate, consider the historical costs of interracial relationships in the United States. Even after the Supreme Court's decision in *Loving v. Virginia*, which struck down state laws that prohibited interracial marriage, there remained much difficulty in forming interracial relationships.<sup>26</sup> This would require transgressing the boundaries of a segregated social network, often incurring significant social stigma or even the prospect of violence. In this social environment, the costs of some associational options are unreasonable. There is some evidence that such relationships are more costly than homogamous (i.e. same-group) ones; sociologists find that couples in interracial marriages are more likely to divorce than their homogamous counterparts (Schwartz 2013, 463-464). One possible explanation for this pattern is that these relationships are

---

<sup>26</sup> A Pew study finds that attitudes opposing interracial marriage have drastically declined, but Herman & Campbell find that these attitudes do not translate into high rates of actual interracial marriage among whites (Livingston & Brown 2017, Herman & Campbell 2012).

less socially accepted and recognized by others. Another possible explanation is that the individuals in these relationships do not share similar enough social schemas – the social scripts and norms that grease the wheels of social interaction by telling us what the appropriate behavior or sequence of events is in well-known situations (Bicchieri 2005, 81-98). Both these external and internal costs are sources of friction that homogamous couples do not encounter. These points can generalize to relationships apart from marriage as well, such as casual friendship, though it is harder to find evidence for this pattern in less socially well-defined relationships. Nevertheless, it is not hard to imagine that in many contexts, homogamous friendships are easier to form because both parties share social schemas. And the more sorted segments of society are from each other, the more difficult it is to converge on shared schemas, because there are fewer opportunities for interaction. The Reasonable Costs Condition is thus difficult to satisfy in sorted environments.

Sorted environments can also restrict the salience of objectively available options, thereby making the Salience Condition difficult to satisfy. Individuals within a sorted environment develop a shared perspective, and one key element of such a perspective is that it delineates insider/outsider membership and criteria of belonging. As Muldoon puts it, “social norms can conspire with perspectives to make it difficult for agents to

realize that there are even other options” (2017, 273). If boundaries are rigid enough, over time, outsiders cease to be considered salient as possible friends, neighbors, partners, etc. For this reason (among others), Anderson (2010) advocates strongly for racial integration. To the extent that segregation renders some people salient and others invisible, or at least only confined to certain roles, it limits our associational options. When the number of salient options dips below a plausible minimal threshold of two or three, our freedom of association is restricted.

In sum: the ecological conception consists of the Choice, Reasonable Costs, and Salience conditions, but it is hard to satisfy these conditions within environments that are completely sorted or sorted to a high degree, such that one either encounters only one or very few associational options. Consequently, freedom of association is undermined in such environments.

#### **4.4.2 The scope and upshots of the ecological account**

The conditions that constitute ecological freedom of association are intended only to apply only to particular social environments, such as neighborhoods, schools, or clubs. In that respect, the argument is modest at its core. However, it can have broader implications. If enough local social environments are segregated, then it becomes harder to move from one environment to another. We can't expand our freedom of association

by simply escaping a segregated environment. For example, if all the neighborhoods affordable to you are segregated, it is not possible to escape segregation by moving. In such a world, the exit strategy for increasing freedom is blocked. Our freedom of association is thus even more compromised in a social world composed of many segregated environments, residential or otherwise social. This characterizes our world at present.

One upshot of the ecological account is that freedom of association is best construed dynamically; it is not a static state that is guaranteed as long as there is absence of interference on our actual choices. Our degree of freedom can vary according to changes in other local social environments. And it is possible that exercises of freedom of association eventually tend towards collective unfreedom, even if they arise from unproblematic initial conditions. These comments are quite speculative, but they are in line with longstanding criticisms of ideal theory (Gaus 2016), as well as with recent cultural evolutionary models on the origins and persistence of inequality that show that vigilant measures are necessary to maintain equality (O'Connor 2019).

In sum, I have argued that meaningful freedom of association requires the joint satisfaction of three conditions: Choice, Reasonable Costs, and Salience. Current discussions of freedom of association ignore the fact that social forces can restrict our

choices just as much as intentional action by agents can. Freedom of association is socially enabled and constrained.<sup>27</sup> Therefore, we lack freedom of association under complete or very substantial conditions of segregation. It follows that freedom of association cannot be invoked as a defense against ameliorating segregation.

#### **4.5 Objections**

Let's return to the problem that motivated this account. Racial segregation, as a subtype of sorting, is unjust in a number of ways, but it seems to be licensed – to some degree – by freedom of association. At least, that is the orthodox view. The ecological conception rejects this opposition and claims that segregation beyond a certain point undermines freedom of association. Thus, a direct implication of the ecological conception is that freedom of association cannot justify a hands-off approach to racial segregation. More strongly, I will argue that it impels the state to support and implement desegregation policy. Before making these arguments, let me first address several objections and implications.

---

<sup>27</sup> See Lawless, "Gruesome Freedom," for a compelling account of freedom in social context. He argues that negative liberty gives us an account of "gruesome freedom" that combines a concern for the interpersonal with a lack of resources to deal with social conditions. My account fits well with his.

First is the “Schelling-Sakoda model” objection<sup>28</sup>. Thomas Schelling and James Sakoda independently showed that a diverse initial environment could settle at an equilibrium of total segregation, as long as individuals preferred to move next to a neighbor of the same color. While it is not clear how much that model explains about the dynamics of segregation in the real world,<sup>29</sup> the Schelling model at least shows that segregated environments can emerge from diverse environments as the *unintentional* result of uncoerced preference. Consider what happens in a typical high school cafeteria: students sort into cliques based on race or socioeconomic status. Nobody aims at creating this situation; it emerges as the result of students congregating with those who are similar to themselves. This is perhaps as close to a Schelling-Sakoda-type situation as we can get in the real world. The objection, then, is that if segregation is the equilibrium outcome of free choice from integrated initial conditions, why is freedom of association compromised? On this objection, we cannot object to the final result of a series of free

---

<sup>28</sup> More commonly known just as the Schelling model (1971). Rainer Hegselmann (2017) argues that the Schelling model is just an instance of Sakoda’s model (1971), which had been published in his dissertation in 1949. But for historical and technical reasons, Schelling’s version became famous, while Sakoda’s was relegated to obscurity.

<sup>29</sup> The explanatory relationship between an abstract model and its real-world target phenomena is the subject of a large literature in the philosophy of science. See Elliott-Graves and Weisberg (2014) for an overview that focuses on the Schelling model.



choices simply because it is unjust or otherwise undesirable (as long as each of these choices was made consensually (Nozick 1974, 161-163). Surely the etiology of segregation matters here.

In response, it is quite coherent, even though perhaps counterintuitive, to say that we lack freedom of association as the result of exercising it. As I argued in Chapter 3, we often make a series of individually insignificant, preferred choices that have the perverse result of landing us in an undesirable end state, which is simply the aggregation of all these individual choices. This can occur for one individual over time, as in the case of someone who puts off exercising each day and as a predictable result, never reaches her fitness goal. Or it can occur for many individuals making simultaneous choices that have collective harm, as with the impact of consumer decisions on factory farming or climate change. This relationship between individual actions and their aggregate outcome is a major puzzle in the practical rationality and collective harm literatures (Nefsky 2019), and it is not necessary to resolve that here. All that is needed to respond to the Schelling-Sakoda objection is to note that from the fact that an outcome is the result of individual free choice, it does not follow that the outcome is just. We can become unfree as the result of free choice; freedom of association has the structure of the general paradox. And on the ecological conception,

we do become unfree, because our choices are self-undermining. The burden of proof is on those who would invoke the Schelling-Sakoda model to show why an unintentionally segregated outcome is just.

What about a case in which individuals *intentionally* seek out a segregated environment? This scenario is not confined to cases of obvious bigotry such as mid-century Atlanta, in which white community leaders appealed to freedom of association as a justification for excluding Blacks from suburban neighborhoods (Kruse 2013). Even today, segregation is maintained by similar, though more subtle, mechanisms. Wealthy communities routinely reject proposals for low-income housing with similar appeals to the freedom to construct the character of their neighborhood, including their choice of neighbors (Levine Einstein et al 2020) – appeals that roughly fall under the umbrella of freedom of association. What should we make of segregation that arises from such intentional choices?

In response, consider an analogy to this explanation for women's attrition from the workforce: a well-educated woman has the option of pursuing a demanding career or staying at home with her children. She makes an initial choice to stay home and never regrets her choice. But as a result of that choice, her outside options wither away over time because it is difficult to reenter the workforce. Thus, the only realistic option for her

is to continue to stay at home. It is simply good luck (and perhaps adaptive preferences) that her preferences align with the only option she has, even though this option was gladly chosen. But her preferences may just as well have ceased to align with her choice to stay at home, a situation faced by many unhappy women. In this scenario, the problem is that there is no viable exit option for the woman. The intentional segregationist faces a similar situation: even if they continue to prefer segregation, they have constructed an environment in which they have no other choice, should they eventually come to disprefer segregation. Intentional segregation constrains freedom by foreclosing opportunities.

Finally, the ecological account seems to have an unwanted implication: that minority communities that choose to self-segregate are limiting their own freedom of association. This objection is captured by Tommie Shelby (2014), who argues against Anderson that wholesale integration is not the solution to racial injustice. Black communities might have good reasons to self-segregate in the private realm of life. What is needed, he argues, is egalitarian pluralism: “racial justice requires desegregation and economic fairness but does not require residential integration or proscribe voluntary self-segregation in neighborhoods” (p.274).

In response, my argument does not require *wholesale* integration. It only requires a less than complete degree of segregation, which would be a vast improvement over what we currently have in most neighborhoods, schools and realms of social life. That is, it only requires that people have more than one option in their associational life, so it is compatible with Shelby's argument for the value of voluntary self-segregation.

#### **4.6 Social structure as the site of intervention**

What does the ecological conception imply, practically speaking? Perhaps surprisingly, it does not require that we choose differently as individuals. Consider a privileged activist who is aware of the unjust effects of our associational choices. She takes it to be her duty to befriend and date marginalized people to further the cause of justice and makes a special effort to seek them out. Without her sense of justice driving her decisions, she would reflexively choose to associate with people who are socioeconomically more similar to her. But it seems that something has gone wrong with the activist's way of thinking; there is something rather creepy and inappropriate about it. Her companions would be right to feel indignant and insulted that they are, in some sense, being used as means to the end of justice, instead of being sought out for who

they are as individuals.<sup>30</sup> This feeling would be justified even if the activist and her companions ultimately form deep and genuine bonds. The relationship seems tainted by the role of obligation in its genesis. If this all sounds plausible, it points to the fact that part of what it means to form a relationship is its voluntary, amoral nature; forming an association out of a sense of obligation can undermine the very nature of the association.<sup>31</sup> I do not mean that there are no duties, such as care and loyalty, toward our close associates. Rather, moral obligation should not be a factor in the formation of these relationships.

If freedom of association is undermined by segregation, should the state prohibit homogamous associations from forming? This would be a radical expansion of state power. Fortunately, I do not think my account entails this. Even if our choices are unfree, it does not follow that an external actor has the permission to interfere with those choices. Consider again the longstanding debate over women's choices to stay at home

---

<sup>30</sup> Obviously this example is highly simplistic, but actual cases of this sort do sometimes occur, with complications arising from mixed motives. See Hanna Rosin and Alix Spiegel, "A Very Offensive Rom-Com," Podcast, *Invisibilia* (NPR, April 5, 2019), <https://www.npr.org/programs/invisibilia/710046991/a-very-offensive-rom-com>, for a true story about such a person and the ethical issues she navigates.

<sup>31</sup> Unless, perhaps, in inegalitarian caring relationships such as parent-child relationships where the parent should feel some parental obligation.

or enter into the so-called “feminized” professions. Whether this is due to ideology, adaptive preferences, severely constrained set of options, or all of these factors is unclear. But whatever the provenance of women’s choices, it is widely agreed that the if the state (and other powerful actors) should do anything about these choices, they should *only* expand the range of options available to women, not interfere on their actual choices. And not only should the state expand the range of options available, but it should take efforts to make them feasible choices, e.g. by mandating paid parental leave. The same reasoning applies to interference on our actual relationships. Even if they are formed under conditions of unfreedom, the state may not directly interfere on them. But this prohibition on interference does not stem from respect for freedom of association per se. Rather, it is justified by liberalism’s more general respect for individuals’ choices, regardless of whether those choices are constrained.

It is entirely possible that there is no reason for the state to respect an unfree choice. That is the subject of a large literature on paternalism, which again, we cannot do justice to here.<sup>32</sup> If that is right, then we should accept that the state may interfere

---

<sup>32</sup> See Coons & Weber (2013) for an overview of the paternalism debate.

directly with our personal relationships, though there are probably many good pragmatic reasons for it to refrain from doing so, such as the fact that any particular governmental actors will likely use such power for their own political ends. In the interest of having theory cohere with common sense as much as possible, my argument has tried not to require direct state interference on relationships. But the truth of the ecological conception does not hang on the question of paternalism. Perhaps the state does have a reason to directly interfere on our relationships, but that does not have much relevance for my purposes.

If individuals shouldn't necessarily change their associational choices, and the state shouldn't directly intervene on our relationships, are we simply back where we started – at an intractable standoff between private action and the demands of justice? Fortunately, another site of intervention is available – the social environment itself. This can be accomplished through several policies that have been criticized as infringing on freedom of association.

First and most obviously, desegregation measures need to be much more robust in order to expand the range of associational options that are objectively available. This will go some way toward helping to satisfy the Choice Condition. Without expanding the search environment, the Reasonable Costs and Salience Conditions are not as

relevant. For a case study in strong desegregation initiatives, consider the course chosen by Montgomery County, Maryland, a wealthy suburb of Washington, D.C. In the 1970s, much of the country was resisting the federal government's Fair Housing Act, which required cities to "affirmatively further" fair housing. Montgomery was 92-percent white, and one of the wealthiest in the country. Amidst much resistance by homeowners, the county housing commission spent years advocating for a zoning ordinance that required any development of 50 or more units (now lowered to 20) to set aside at least 15 percent of housing for low-income residents; up to a third of this could be turned into affordable housing by the county (Hannah-Jones 2015). Montgomery County remains one of the wealthiest counties in the country, but now is one of the most diverse.<sup>33</sup> By contrast, Westchester County, a wealthy and mostly white suburb of New York, has been in an ongoing battle with the federal government over its resistance to building affordable housing. On the ecological account, this resistance relies on too thin a conception of freedom of association. There are better, freedom-enhancing reasons to plump for the course taken by Montgomery County.

---

<sup>33</sup> According to Census data, Montgomery County is 43.4% white alone (non-Hispanic or Latino), 60.2% white alone (including Hispanic or Latino), 19.9% black alone, 15.6% Asian alone, 3.4% two or more races, 0.7% American Indian and Alaska Native, and 19.9% Hispanic or Latino.



Beyond housing, we should also take steps to increase ecological freedom of association in the digital realm, which now constitutes a large part of the social environment. Legal scholars and human-computer interaction (HCI) theorists have recently argued that there is no such thing as “platform neutrality” (Hutson et al 2018); the design of digital platforms, including dating websites, social networks, and sharing-economy platforms either promotes or inhibits racial discrimination and online segregation (Edelman et al 2017). In the digital realm, the filtering of types of individuals into different “bubbles” by algorithms, depending on their demographic characteristics or ranking within a system (Pariser 2012), restricts the search environment. Filtering also happens deliberately, as when individuals filter prospective dating matches by selecting ethnicity criteria (Lim et al 2020). On my account, design choices that promote segregation should be prohibited or at least highly regulated for reasons of enhancing freedom of association. Considerations of autonomy (Whitney 2017) are not such an insurmountable obstacle to these design choices.

These policies will go some way towards at least satisfying the conditions for freedom of association. But it is likely that even these won’t be sufficient. What is needed are deeper interventions at the level of social structure – the schemata, norms, and practices that scaffold and give meaning to social life. We can think of social structure as

imposing a *choice architecture* on associational options. To illustrate the idea of a choice architecture, think of a grocery store. The organization of even something as mundane as a grocery store is not random, nor is it the result of a copycat effect in which stores look the same because customers expect them to be that way. Each store is carefully designed to get the customer to spend more money. For example, the dairy case is in the back of the store so that the customer will need to walk through the middle, and in doing so, they will be more likely to make random purchases (Thaler & Sunstein 2008). Similarly, social structure imposes a choice architecture on the social world. It directs our attention to certain features of the world, imposes costs on some choices, and incentivizes other choices.

The norm of homophily, the tendency to associate with people like ourselves, is a core element of social structure that maintains segregation. Homophily is not just a descriptive norm, which people prefer to follow conditional on the belief that everyone else will follow it (Bicchieri 2005, 112). It is also imbued with social meaning, given that it cleaves people into categories based on signals of race, socioeconomic status, and other relevant identities. Homophily is therefore also a social norm: many people have believed and continue to believe that it *ought* to be followed (Bicchieri 2005, 34). Those

who do not follow the norm are liable to be stigmatized, shunned, or otherwise excluded from mainstream society.

But norms do not emerge or operate in a vacuum. They grab our attention and gain a hold over our social life in part because we look for them. On Cristina Bicchieri's influential account of norms, they are activated by cognitive schemata (or scripts), which describe an appropriate sequence of events and embed norms into that sequence (2005, 81-82). Muldoon calls this "shared perspectives" (2017, 268); Haslanger alternatively describes this as "culture" (2017, 149-73). Our schemata, perspectives, or culture categorize people into socially significant groups, and certain norms of association (i.e. different degrees of inclusion or exclusion) are activated for different groups. The homophily norm can generate clusters of people who are similar in race and socioeconomic status because these categories are salient in our eyes. We could, in theory, cluster according to height and hair color. The reason that we don't do so is that those are not socially salient categories. Clustering on those dimensions is secondary and in large part due to correlations with the more salient category, race. What is needed to undermine the homophily norm is a change in social schema, perspectives, and culture. Doing so will bring us closer toward achieving freedom of association.

## **4.7 Conclusion**

Many liberal theorists think that the priority of freedom of association creates an unfortunate roadblock to achieving justice. Intuitively, we should be allowed to befriend, live among, marry, and otherwise associate with whomever we choose – even if this results in injustice. These individual decisions, in the aggregate, create and maintain patterns of segregation. In this paper, I have offered an ecological conception of freedom of association. This account brings to light the social forces that interfere with our freedom of association, and shows that the degree to which we have freedom of association largely depends on the structure of our social environment. Under unjust conditions such as complete segregation, we lack meaningful freedom of association. This reorientation of perspective mitigates tension between freedom of association and justice. It also licenses, and indeed requires, stronger interventions on the social environment in the form of policies and cultural change. One upshot of this account is that it enhances our own freedom of association to support such measures. Freedom and justice can walk together much further than previously thought.

## Conclusion

We have the right to choose the society most acceptable to us. But sometimes, these choices lead to injustice. This dissertation has aimed to set out the nature of this tension and offer a theory of our moral and political duties to address it. In doing so, I draw many seemingly disparate social phenomena under one theoretical umbrella. The clustering of students into racially and socioeconomically homogenous groups in school cafeterias, racial and income residential segregation, private clubs such as fraternities and sororities, and the marginalization of certain racialized and sexualized groups of people in the dating realm all have something in common: they show that there is an irreducibly social dilemma between our associational choices and ideals of justice. This is a deeply uncomfortable tension, for our associational lives are extremely personal, and our associational choices seem to be out of our control. But it is not clear that we should control these choices in the service of justice.

In Chapter 1, I argued that there is an irreducibly social dilemma between our associational choices and ideals of justice: this is the problem of sorting. Because our associational choices are homophilous rather than random, they tend to perpetuate opportunity hoarding: the concentration of social, developmental, and economic opportunities within some advantaged networks. It is inaccurate to characterize the

causes of sorting as purely caused by unjust structures or by stigma against social groups, as there are multiple ways by which opportunity can be hoarded. Some of these ways are apparently innocuous. But all of these choices are protected by the Millian “right to choose the society most acceptable to us.” In the aggregate, these choices undermine justice.

Are those who are privately excluded by this right wronged, and therefore entitled to redress in the form of association? And if they are not, why are those who are publicly excluded wronged, and therefore entitled to such redress? Our moral treatment of public and private exclusion is deeply inconsistent. Chapter 2 further developed this dilemma by offering the concept of associational opportunities as a way to justify this differential treatment. I argued that no one is entitled to any particular person’s association. However, public exclusion wrongs the excluded by denying them associational opportunities. So does private exclusion that is based on discrimination. But private exclusion that is based on personal characteristics, such as egregious behavior, do not wrong the excluded, because the excluded have failed to reciprocate appropriately given an associational opportunity.

Chapter 3 addressed the collective dilemma nature of the problem of sorting. The problem of sorting appears to have the structure of a collective action problem, in which

the harms of sorting emerge only if enough people make relevantly similar choices. But no one has an incentive to make a different, perhaps morally better choices, so sorting is entrenched as an equilibrium. I argued that the concept of diachronic rationality gives us a resource to deal with this feature of the problem. Each person has an instrumental reason to act in ways that transform the search environment to be more conducive to justice. Those who have the desire for justice, but fail to act in such ways, pay the cost of their own self-governance.

Finally, Chapter 4 turned to the concept of freedom of association itself. I offered an ecological conception of this freedom, on which associational freedom requires certain social conditions for genuine choice to be in place. We can constrain our own associational freedom by making many individual choices that construct homogenous social environments. Thus, freedom of association requires dynamic attention to the composition of the social environment and does not imply a laissez-faire approach to social life. The upshot is that completely or substantially segregated environments undermine our associational freedom, even if such segregation came about as the result of initially free choice. A certain degree of diversity is required for associational freedom.

These arguments represent an attempt to grapple with the problem of sorting on liberalism's own terms. From the problem of sorting, we might easily jump to the conclusion that there is an ineradicable tension between negative associational freedom and the imperatives of justice, and thereby conclude that this is one more reason why liberalism is deeply flawed. It might well be. But I hope to have shown that liberalism contains within itself some of the resources we need for reconciling associational freedom and justice. Once we expand our view of individual instrumental reasons, and look outwards toward the social environment and the opportunities or choices it affords, we might see that the freedom of association and the imperatives of justice may yet travel a long way together.



## References

- Aisch, G., Buchanan, L., Cox, A., & Quealy, K. (2017, January 18). Some colleges have more students from the top 1 percent than the bottom 60 percent. Find yours. *The New York Times*. <https://www.nytimes.com/interactive/2017/01/18/upshot/some-colleges-have-more-students-from-the-top-1-percent-than-the-bottom-60.html>
- Allport, G. (1954). *The nature of prejudice*. Addison-Wesley.
- Anderson, E. (1999). What is the point of equality? *Ethics*, 109(2), 287–337.
- Anderson, Elizabeth. (1999). What is the Point of Equality? *Ethics*, 109(2), 287–337.
- Anderson, Elizabeth. (2010). *The Imperative of Integration*. Princeton University Press.
- Andreou, C. (2006). Environmental damage and the puzzle of the self-torturer. *Philosophy & Public Affairs*, 34(1), 95–108.
- Andreou, C., & White, M. D. (2010). *The Thief of Time: Philosophical Essays on Procrastination*. Oxford University Press.
- Arneson, R. (1989). Equality and Equal Opportunity for Welfare. *Philosophical Studies*, 56, 77–93.
- Barlett, R. C., & Collins, S. D. (2012). *Aristotle's Nicomachean Ethics*. University of Chicago Press.
- Bedi, S. (2015). Sexual Racism: Intimacy as a Matter of Justice. *Journal of Politics*, 77(4), 998–1011.
- Bell, D. A. (1998). Civil Society versus Civic Virtue. In A. Gutmann (Ed.), *Freedom of Association* (pp. 239–272). Princeton University Press.
- Berlin, I. (1969a). Two Concepts of Liberty. In *Four Essays on Liberty* (pp. 118–172). Oxford University Press.

- Bicchieri, C. (2005). *The Grammar of Society*. Cambridge University Press.
- Blair, K. L., & Hoskin, R. A. (2018). Transgender exclusion from the world of dating: Patterns of acceptance and rejection of hypothetical trans dating partners as a function of sexual and gender identity. *Journal of Social and Personal Relationships*, 1–22.
- Bokulich, A. (2014). How the tiger bush got its stripes: “how possibly” vs. “how actually” model explanations. *The Monist*, 97(3), 321–338.
- Boustan, L. (2013). *Racial Residential Segregation in American Cities* [Working Paper No.19045].
- Brake, E. (2012). *Minimizing marriage: Marriage, morality, and the law*. Oxford University Press.
- Bratman, M. (2003). A Desire of One’s Own. *Journal of Philosophy*, 100(5), 221–242.
- Bratman, M. (2007). Temptation Revisited. In *Structures of Agency*. OUP.
- Bratman, M. (2018). *Planning, time, and self-governance: Essays in practical rationality*. OUP.
- Brighouse, H., & Swift, A. (2009). Legitimate parental partiality. *Philosophy & Public Affairs*, 37(1), 43–80.
- Broome, J. (2018). Against denialism. *The Monist*, 102(1), 110–129.
- Brownlee, K. (2015). Freedom of Association: It’s Not What You Think. *Oxford Journal of Legal Studies*, 35(2), 267–282.
- Brownlee, K. (2020). *Being sure of each other*. Oxford University Press.
- Buchanan v. Warley, No. 33 (United States Supreme Court November 5, 1917).  
<https://www.oyez.org/cases/1900-1940/245us60>
- Budolfson, M. B. (2018). The inefficacy objection to consequentialism and the problem with the expected consequence response. *Philosophical Studies*.

- Chambers, C. (2009). Each outcome is another opportunity: Problems with the Moment of Equal Opportunity. *Politics, Philosophy, and Economics*, 8(4), 374–400.
- Chambers, C. (2017). *Against marriage: An egalitarian defence of the marriage-free state*. Oxford University Press.
- Chetty, R., Hendren, N., Kline, P., & Saez, E. (2014). Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *Quarterly Journal of Economics*, 129(4), 1553–1623.
- Cingano, F. (2014). *Trends in income inequality and its impact on economic growth*. [https://www.oecd-ilibrary.org/social-issues-migration-health/trends-in-income-inequality-and-its-impact-on-economic-growth\\_5jxrjncwxv6j-en;jsessionid=XNBnsqAV0lgrRCdFXC4ErgWh.ip-10-240-5-16](https://www.oecd-ilibrary.org/social-issues-migration-health/trends-in-income-inequality-and-its-impact-on-economic-growth_5jxrjncwxv6j-en;jsessionid=XNBnsqAV0lgrRCdFXC4ErgWh.ip-10-240-5-16)
- Cohen, G. A. (1989). On the Currency of Egalitarian Justice. *Ethics*, 99(4), 906–944.
- Cohen, P. N., & Pepin, J. R. (2018). Unequal marriage markets: Sex ratios and first marriage among Black and White women. *Socius*, 4.
- Collins, S. (2013). Duties to make friends. *Ethical Theory and Moral Practice*, 16(5), 679–695.
- Coons, C., & Weber, M. (Eds.). (2013). *Paternalism: Theory and Practice*. Cambridge University Press.
- Cooper, J. M. (Ed.). (1997). *Plato: Complete Works*. Hackett.
- Cordelli, C. (2015). Justice as fairness and relational resources. *The Journal of Political Philosophy*, 23(1), 86–110.
- Cudd, A. E. (2006). *Analyzing Oppression*. Oxford University Press.
- Damaske, S., Bratter, J. L., & Frech, A. (2017). Single mother families and employment, race, and poverty in changing economic times. *Social Science Research*, 62, 120–133.

- Demographic trends and economic well-being*. (2016). Pew Research Center.  
<http://www.pewsocialtrends.org/2016/06/27/1-demographic-trends-and-economic-well-being/>
- Dworkin, R. (1981a). What is equality? Part 1: Equality of welfare. *Philosophy & Public Affairs*, 10(3), 185–246.
- Dworkin, R. (1981b). What is equality? Part 2: Equality of resources. *Philosophy & Public Affairs*, 10(4), 283–345.
- Edelman, B., Luca, M., & Svirsky, D. (2017). Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*, 9(2), 1–22.
- EEOC v. Consolidated Services Systems, 777 F. Supp. 599 (N. D. Ill. 1991), No. 85 C 8312 (United States District Court, N.D. Illinois, E.D. September 4, 1991).  
<https://law.justia.com/cases/federal/district-courts/FSupp/777/599/2259891/#:~:text=The%20EEOC%20claims%20that%20the,t hroughout%20Cook%20County%20and%20the>
- Elliott-Graves, A., & Weisberg, M. (2014). Idealization. *Philosophy Compass*, 9(3), 176–185.
- Elster, J. (2016). *Sour Grapes: Studies in the Subversion of Rationality*. Cambridge University Press.
- Epstein, R. A. (2008). Should Antidiscrimination Laws Limit Freedom of Association? The Dangerous Allure of Human Rights Legislation. *Social Philosophy & Policy*.
- Fine, S. (2010). *Freedom of Association is Not the Answer*. 120(2), 338–356.
- Fishkin, James. (1983). *Justice, equal opportunity, and the family*. Yale University Press.
- Fishkin, Joseph. (2014). *Bottlenecks: A new theory of equal opportunity*. Oxford University Press.

- Frank, R. (1985). *Choosing the right pond*. Oxford University Press.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy*, 5–20.
- Freeman, K. (1999). HBCUs or PWIs? African American high school students' consideration of higher education institution types. *Review of High Education*, 23(1), 91–106.
- Frey, W. H. (2014). *Multiracial marriage on the rise*. Brookings Institution.  
<https://www.brookings.edu/blog/the-avenue/2014/12/18/multiracial-marriage-on-the-rise/>
- Gaus, G. (2016). *The Tyranny of the Ideal: Justice in a Diverse Society*. Princeton University Press.
- Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360–1380.
- Gutmann, A. (1998). *Freedom of Association*. Princeton University Press.
- Hannah-Jones, N. (2015, June 25). Living Apart: How the Government Betrayed a Landmark Civil Rights Law. *ProPublica*.  
<https://www.propublica.org/article/living-apart-how-the-government-betrayed-a-landmark-civil-rights-law>
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.
- Haslanger, S. (2016). What is a (Social) Structural Explanation? *Philosophical Studies*, 173, 113–130.
- Haslanger, S. (2017). Culture and Critique. *Aristotelian Society Supplementary Volume*, 91(1), 149–173.
- Hegselmann, R. (2017). Thomas C. Schelling and James M. Sakoda: The Intellectual, Technical, and Social History of a Model. *Journal of Artificial Societies and Social*

*Simulation*, 20(3).

Herman, M. R., & Campbell, M. E. (2012). I Wouldn't, But You Can: Attitudes Toward Interracial Relationships. *Social Science Research*, 41(2), 343–358.

Hill, H. (1996). The problem of race in American labor history. *Reviews in American History*, 24(2), 189–208.

Hofstra, B., Corten, R., van Tubergen, F., & Ellison, N. B. (2017). Sources of segregation in social networks: A novel approach using Facebook. *American Sociological Review*, 82(3), 625–656.

Hohfeld, W. N. (1917). Fundamental legal conceptions as applied in judicial reasoning. *Yale Law Journal*, 26(8), 710–770.

Horowitz, J. M., Igielnik, R., & Kochhar, R. (2020). *Trends in income and wealth inequality*. Pew Research Center. <https://www.pewresearch.org/social-trends/2020/01/09/trends-in-income-and-wealth-inequality/>

Hutson, J. A., Taft, J. G., Barocas, S., & Levy, K. (2018). Debiasing Desire: Addressing Bias & Discrimination on Intimate Platforms. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), 1–18.

Jackson, F., & Pettit, P. (1992). Structural Explanation in Social Theory. In D. Charles & K. Lennon (Eds.), *Reduction, Explanation and Realism*. Oxford University Press.

James, W. (1896). *The will to believe*.

Janiak, C. M. (2003). The “links” among golf, networking, and women’s professional advancement. *Stanford Journal of Law, Business and Finance*, 2(8), 317–348.

Jones-Correa, M. (2000). The Origins and Diffusion of Racial Restrictive Covenants. *Political Science Quarterly*, 115(4), 541–568.

Kagan, S. (2011). Do I make a difference? *Philosophy & Public Affairs*, 39(2), 105–141.

- Karst, K. (1980). The Freedom of Intimate Association. *Yale Law Journal*, 89(4), 624–692.
- Kelly, E. I. (2017). The historical injustice problem for political liberalism. *Ethics*, 128, 75–94.
- Kingston, E., & Sinnott-Armstrong, W. (2018). What’s wrong with joyguzzling? *Ethical Theory & Moral Practice*, 21(1), 169–186.
- Kolodny, N. (2008). The Myth of Practical Consistency. *European Journal of Philosophy*, 16(3), 366–402.
- Kossinets, G., & Watts, D. J. (2009). Origins of homophily in an evolving social network. *American Journal of Sociology*, 115(2), 405–450.
- Kruse, K. (2013). *White Flight*. Princeton University Press.
- Krysan, M., & Crowder, K. (2017). *Cycle of Segregation: Social Processes and Residential Stratification*. Russell Sage Foundation.
- Lawless, J. (2018). Gruesome Freedom: The Moral Limits of Non-Constraint. *Philosophers’ Imprint*, 18(3).
- Lens, M. C., & Monkkonen, P. (2016). Do Strict Land Use Regulations Make Metropolitan Areas More Segregated By Income? *Journal of the American Planning Association*, 82(1), 6–21.
- Lim, G., Robards, B., & Carlson, B. (2020, June 7). Grindr is deleting its “ethnicity filter.” But racism is still rife in online dating. *The Conversation*.  
<https://theconversation.com/grindr-is-deleting-its-ethnicity-filter-but-racism-is-still-rife-in-online-dating-140077>
- Lippert-Rasmussen, K. (2014). *Born free and equal?: A philosophical inquiry into the nature of discrimination*. Oxford University Press.
- Livingston, G., & Brown, A. (2017). *Intermarriage in the U.S. 50 Years After Loving v. Virginia*. Pew Research Center.  
<https://www.pewsocialtrends.org/2017/05/18/intermarriage-in-the-u-s-50-years->

after-loving-v-virginia/

- Logan, T., & Parman, J. (2015). *The National Rise in Residential Segregation* [Working Paper No.20934].
- Lomasky, L. (2008). The Paradox of Association. *Social Philosophy & Policy*.
- Madva, A. (2019). Integration, Community, and the Medical Model of Social Injustice. *Journal of Applied Philosophy*.
- Manne, K. (2020). *Entitlement: How male privilege hurts women*. Penguin.
- Markovits, D. (2019). *The Meritocracy Trap*. Penguin.
- Massey, D. S., & Denton, N. A. (1993). *American Apartheid: Segregation and the Making of the Underclass*. Harvard University Press.
- McIntosh, K., Moss, E., Nunn, R., & Shambaugh, J. (2020). *Examining the Black-white wealth gap*. Brookings Institution. <https://www.brookings.edu/blog/up-front/2020/02/27/examining-the-black-white-wealth-gap/>
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415–444.
- Mill, J. S. (1859). *On liberty*. Project Gutenberg. <https://www.gutenberg.org/files/34901/34901-h/34901-h.htm>
- Millgram, E. (2015). *The Great Endarkenment*. OUP.
- Moles, A. (2014). The Public Ecology of Freedom of Association. *Res Publica*, 20, 85–103.
- Moody, J. (2001). Race, school integration, and friendship segregation in America. *American Journal of Sociology*, 107(3).
- Morton, J., & Paul, S. K. (2019). Grit. *Ethics*, 129(2), 175–203.



- Muldoon, R. (n.d.). *Diversity Isn't What Divides Us. Division Is What Divides Us*. Knight Foundation.
- Muldoon, R. (2017). Perspectives, Norms, and Agency. *Social Philosophy & Policy*.
- Muldoon, R., Smith, T., & Weisberg, M. (2012). Segregation that no one seeks. *Philosophy of Science*, 79(1), 38–62.
- Nagel, T. (1995). *Equality and Partiality*. Oxford University Press.
- Nefsky, J. (2011). Consequentialism and the problem of collective harm. *Philosophy & Public Affairs*, 39(4), 364–395.
- Nefsky, J. (2017). How you can help, without making a difference. *Philosophical Studies*, 174, 2743–2767.
- Nefsky, J. (2019). Collective harm and the inefficacy problem. *Philosophy Compass*, 14(4).
- Nguyen, C. T. (2019). Games and the art of agency. *Philosophical Review*, 128(4), 423–462.
- Nguyen, T. (2020, September 29). Why it's so difficult to abolish fraternities and sororities. *Vox*. <https://www.vox.com/the-goods/21492167/abolish-greek-life-campus-covid>
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Basic Books.
- O'Connor, C. (2019). *The Origins of Unfairness: Social Categories and Cultural Evolution*. Oxford University Press.
- Orfield, G., & Frankenberg, E. (2014). *Brown at 60: Great Progress, a Long Retreat, and an Uncertain Future*. The Civil Rights Project, UCLA.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Pariser, E. (2012). *The Filter Bubble: What the Internet is Hiding From You*. Penguin.
- Paul, L. A. (2014). *Transformative experience*. OUP.

- Pendall, R., Puentes, R., & Martin, J. (2006). *From Traditional to Reformed: A Review of the Land Use Regulations in the Nation's 50 largest Metropolitan Areas*. The Brookings Institution.
- Pettit, P. (2011). The Instability of Freedom as Noninterference: The Case of Isaiah Berlin. *Ethics*, 121, 693–716.
- Pettit, P. (2015). *The robust demands of the good: Ethics with attachment, virtue, and respect*. Oxford University Press.
- Phillips, A. (2004). Defending Equality of Outcome. *Journal of Political Philosophy*, 12(1), 1–19.
- Piketty, T., Saez, E., & Zucman, G. (2016). *Economic growth in the United States: A tale of two countries*. <https://equitablegrowth.org/economic-growth-in-the-united-states-a-tale-of-two-countries/>
- Putnam, R. D. (2000). *Bowling alone*. Simon and Schuster.
- Quinn, W. (1990). The Puzzle of the Self-Torturer. *Philosophical Studies*, 59, 79–90.
- Rawls, J. (1971). *A theory of justice*. Harvard University Press.
- Rawls, J. (1993). *Political Liberalism*. Columbia University Press.
- Rawls, J. (2001). *Justice as fairness: A restatement*. Harvard University Press.
- Reardon, S. F., & Bischoff, K. (2011). Income inequality and income segregation. *American Journal of Sociology*, 116(4), 1092–1153.
- Reardon, S. F., Fox, L., & Townsend, J. (2015). Neighborhood Income Composition by Household Race and Income, 1990-2009. *Annals of the American Academy of Political and Social Science*, 660(1), 78–97.
- Reeves, R. V., & Rodrigue, E. (2015). *Single black female BA seeks educated husband: Race,*

*assortative mating, and inequality*. Brookings Institution.  
<https://www.brookings.edu/research/single-black-female-ba-seeks-educated-husband-race-assortative-mating-and-inequality/#:~:text=There%20is%20a%20growing%20trend,spouses%20with%20similar%20educational%20attainment.&text=Households%20with%20two%20college%20graduates,households%20with%20less%20Deducated%20couples>.

- Richman, B. (2006). How community institutions create economic advantage: Jewish diamond merchants in New York. *Law & Social Inquiry*, 31(2), 383–420.
- Rosenberg, N., & Stucki, B. W. (2019, June 26). How USDA distorted data to conceal decades of discrimination against Black farmers. *The Counter*.  
<https://thecounter.org/usda-black-farmers-discrimination-tom-vilsack-reparations-civil-rights/>
- Rosin, H., & Spiegel, A. (2019, April 5). A Very Offensive Rom-Com [Podcast]. In *Invisibilia*. NPR. <https://www.npr.org/programs/invisibilia/710046991/a-very-offensive-rom-com>
- Rothstein, R. (2015). The racial achievement gap, segregated schools, and segregated neighborhoods: A constitutional insult. *Race and Social Problems*, 7(21), 21–30.
- Rothstein, R. (2017). *The Color of Law: A Forgotten History of How Our Government Segregated America*. Liveright.
- Rothwell, J., & Massey, D. S. (2009). The Effect of Density Zoning on Racial Segregation in U.S. Urban Areas. *Urban Affairs Review*, 44(6), 779–806.
- Rothwell, J. T., & Massey, D. S. (2010). Density zoning and class segregation in U.S. Metropolitan Areas. *Social Science Quarterly*, 91(5), 1122–1142.
- Ryan, A. (1998). The City as a Site for Free Association. In A. Gutmann (Ed.), *Freedom of Association* (pp. 314–329). Princeton University Press.
- Sakoda, J. M. (1971). The Checkerboard Model of Social Interaction. *Journal of Mathematical Sociology*, 1(1), 119–132.

- Schelling, T. C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1, 143–186.
- Schelling, T. C. (1978). *Micromotives and Macrobehavior*. W. W. Norton.
- Schroeder, M. (2017). *Slaves of the passions*. Oxford University Press.
- Schwartz, B. (2004). *The Paradox of Choice*. Harper Perennial.
- Schwartz, C. (2013). Trends and Variation in Assortative Mating: Causes and Consequences. *Annual Review of Sociology*, 39, 451–470.
- Schwartz, C. R., & Mare, R. D. (2005). Trends in educational assortative marriage from 1940 to 2003. *Demography*, 42(4), 621–646.
- Shelby, T. (2014). Integration, Inequality, and Imperatives of Justice: A Review Essay. *Philosophy & Public Affairs*, 42(3), 253–285.
- Shertzer, A., & Walsh, R. P. (2019). Racial sorting and the emergence of segregation in American cities. *The Review of Economics and Statistics*, 101(3), 415–427.
- Silver, C. (1997). The racial origins of zoning in American cities. In *Urban planning and the African American community: In the shadows*. Sage Publications.
- Sinnott-Armstrong, W. (2005). It's not my fault: Global warming and individual moral obligations. In *Perspectives on climate change: Science, Economics, and Ethics* (Vol. 5). Elsevier.
- Thaler, R., & Sunstein, C. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.
- Thompson, D. (2016, August 18). Donald Trump and “economic anxiety.” *The Atlantic*. <https://www.theatlantic.com/business/archive/2016/08/donald-trump-and-economic-anxiety/496385/>

- Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5).
- Tilly, C. (1998). *Durable Inequality*. University of California.
- Trounstein, J. (2018). *Segregation by Design: Local Politics and Inequality in American Cities*. Cambridge University Press.
- Turner, M. A., Santos, R., Levy, D. K., Wissoker, D., Aranda, C., & Pitingolo, R. (2012). *Housing Discrimination Against Racial and Ethnic Minorities 2012*. The Urban Institute.
- U.S. Census Bureau releases new educational attainment data. (2020). U.S. Census Bureau. <https://www.census.gov/newsroom/press-releases/2020/educational-attainment.html#:~:text=In%202019%2C%2040.1%25%20of%20non,from%2013.9%25%20to%2018.8%25>.
- Valentini, L. (2016). What's wrong with being lonely? Justice, beneficence, and meaningful relationships. *Aristotelian Society Supplementary Volume, XC*, 49–69.
- Wellman, C. H. (2008). Immigration and Freedom of Association. *Ethics*, 119(1), 109–141.
- Wenar, L. (2021). Rights. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/spr2021/entries/rights/>
- White, S. (1997). Freedom of Association and the Right to Exclude. *Journal of Political Philosophy*, 5(4), 373–391.
- Whitney, H. M. (2017). *The Autonomy Defense of Private Discrimination*.
- Williams, B. (1962). The idea of equality. In P. Laslett & W. G. Runciman (Eds.), *Philosophy, Politics, and Society, Series II* (pp. 110–131). Basil Blackwell.
- Williamson, V., Gode, J., & Sun, H. (2021). "We all want what's best for our kids": *Discussions of D.C. public school options in an online forum*. Brookings Institution. <https://www.brookings.edu/wp->

content/uploads/2021/03/Discussions\_DC\_public\_school\_options\_online\_forum\_Brookings-Report.pdf

- Wilson, V. (2020, September 16). Racial disparities in income and poverty remain largely unchanged amid strong income growth in 2019 [Economic Policy Institute]. *Working Economics Blog*. <https://www.epi.org/blog/racial-disparities-in-income-and-poverty-remain-largely-unchanged-amid-strong-income-growth-in-2019/>
- Yi, S.-H., & Hoston, W. T. (2020). Demystifying Americanness: The model minority myth and the Black-Korean relationship. *Journal of Ethnic and Cultural Studies*, 7(2), 68–89.
- Zipf, G. (1949). *Human behavior and the principle of least effort: An introduction to human ecology*. Addison-Wesley.
- Zwolinski, M. (2006). Why not regulate private discrimination? *San Diego Law Review*, 43, 1043–1062.

## Biography

Valerie-Jean Soon studied Philosophy and Economics at Wellesley College, where she graduated in 2013. She obtained her M.A. in Philosophy at the University of Houston in 2015, and began graduate study in Philosophy at Duke University in 2015. She has published “Implicit bias and social schema: a transactive memory framework” (2020) in *Philosophical Studies*, “An intrapersonal, intertemporal solution to an interpersonal dilemma” (2021) also in *Philosophical Studies*, and “Social structural explanation” in *Philosophy Compass* (forthcoming). She is a 2020-21 recipient of the Ottis Green Graduate Fellowship at Duke.