Breaking the trust: the case for regulating anonymous shell companies

Andreas-Johann Sorger

A thesis submitted for the degree of MPhil
at the
University of St Andrews

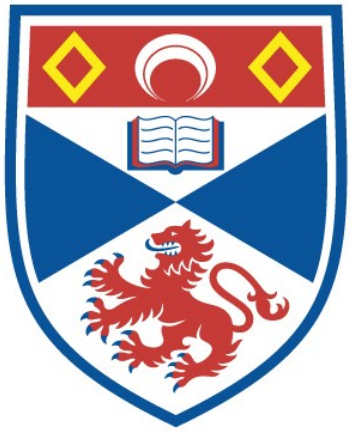

2020

Identifier to use to cite or link to this thesis:
DOI: https://doi.org/10.17630/sta/759

For my parents, Maria and Josef, and Ingrid. Without their support this thesis would not have been written.

Candidate's declaration

I, Andreas-Johann Sorger, do hereby certify that this thesis, submitted for the degree of MPhil, which is approximately 40,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree.

I was admitted as a research student at the University of St Andrews in September 2017.

I confirm that no funding was received for this work.

Date: 04/03/20 Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of MPhil in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date: 9/3/2020 Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Andreas-Johann Sorger, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date: 04/03/20

Date: 9/3/2020

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Andreas-Johann Sorger, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date: 04/03/20 Signature of candidate

General Acknowledgements

I would like to express my heartfelt thanks to Professor Katherine Hawley, Dr Alexander Douglas, Dr Ben Sachs, and Dr Theron Pummer for all their help and support.

Abstract

Anonymous shell companies (ASCs) are corporate entities whose sole purpose is to cloak the identity of their beneficial owner. Due to their strong anonymity provisions, ASCs allow individuals to perform a variety of illicit activities with little chance of being caught. Thus, they have been used in almost every form of economic crime. Though there is universal agreement in the policy sphere that ASCs facilitate a number of negative externalities, policymakers are divided over how they should be regulated. Specifically, policymakers are stuck in an intractable disagreement over the implementation of a public ownership register – a database containing ownership information of every company registered in a particular country. Opponents of this register argue that the public disclosure of ownership information violates a presumptive right individuals have to privacy. Proponents of this register however, deny the existence of this presumptive right. They point instead to the role of transparency in fostering accountability.

The goal of my thesis is to offer a theoretical justification for the creation of a public ownership register. In short, I argue that we can break this impasse by using the value of public trust to justify creating a public ownership register with specific provisions so as to ensure privacy rights are not infringed upon. My argument proceeds in three parts: First, I establish that trust is at least instrumentally valuable. Thus we have a *pro tanto* reason to implement regulations to stop trust being undermined. Second, I offer a novel account of public trust predicated on the assumption of a shared intrinsic commitment to a practice rule. Third, armed with this account of public trust, I identify two distinct mechanisms by which ASCs undermine trust. I conclude by showing how drawing on public trust provides a *pro tanto* reason to implement a public ownership register.

Table of Contents

Introduction

Trusting others is something we do every day – so much so that we often take it for granted. From the mundane, such as asking a stranger for the time, to the profound, like making a promise to a loved one, all of our interactions with others are structured by the presence of trust. It allows us to speak to others without independently verifying every one of their locutions, and – more importantly – gives us the freedom to pursue our own conception of the good with the comfort of knowing that there are people we can turn to if we face difficulties. If anything, a society devoid of trust may resemble a Hobbesian state of nature, where life is characterised by a continuous war of 'all against all' and individuals can rely on no one else but themselves to survive.[1] In short, trust is of paramount importance in fostering cooperation[2], creating a society free from corruption[3], and generating economic prosperity[4] to name just a few. In fact, its significance cannot be understated; we see this sentiment as far back as Confucius, who remarked that of the three things needed for government – food, weapons, and trust – it is trust that is the most important: "Without trust, we cannot stand".[5]

As such, it seems that trust is at least *instrumentally valuable*; that is, its value comes from what trusting others allows us to achieve rather than its presence *simpliciter*. As noted previously, giving testimony would be impossible without the existence of some kind of trust as we would otherwise seek to independently verify every claim our conversation partner utters. This is useful for purposes beyond ordinary small talk – consider the transfer of knowledge. Given that our lives are temporally bound, we only have a limited amount of time in which to pursue projects that we deem worthwhile. Further, as we pursue these projects, we will come across knowledge that we do not

[1] Thomas Hobbes, *Leviathan*, Reprint, Pelican Classics (Harmondsworth: Penguin Books, 1974).
[2] Paul Faulkner, "Finding Trust in Government," *Journal of Social Philosophy* 49, no. 4 (December 2018): 626–44.
[3] Eric M. Uslaner, "Trust and Corruption Revisited: How and Why Trust and Corruption Shape Each Other," *Quality & Quantity* 47, no. 6 (October 2013): 3603–8.
[4] Robert D. Putnam, *Bowling Alone: The Collapse and Revival of American Community* (New York: Simon & Schuster, 2000).
[5] Onora O'Neill, *A Question of Trust*, The BBC Reith Lectures (Cambridge: Cambridge University Press, 2002), 3.

have the time, intellect, or even resources to independently verify. Hence, for the transfer of knowledge to even be possible, we need to have a basis of trust in the testimony of others.[6] Similarly, relationships characterised by mutual trust and trustworthiness are in some sense *meaningful* and therefore important for our own development.[7] To see this, suppose that all our relationships were governed by pure self-interest. Here, we would fail to confide in others, as they would disclose private information only if it were in their self-interest to do so. Therefore, without trust, we would be unable to create and shape meaningful relationships with others that are paramount for our own moral development.[8]

The instrumental value of trust is also observable at the societal level. Societies characterised by high levels of trust – both among their citizens/residents and between a government and their citizens/subjects – tend to be more prosperous and display elements of good governance when compared to societies with lower levels of trust in the public sphere.[9] Similarly, some theorists have argued that trust is necessary for the proliferation and protection of high-value goods that are required for individuals to have a decent life. Feminist philosophers, for example, have argued that individual autonomy flows from the presence of trust, as the necessary conditions for individual autonomy can only emerge in the presence of trustworthy people and institutions.[10] This is predicated on a particular conception of autonomy as a socially-constituted, relational good; if this is true, then it can only emerge if a background of trust is present.[11] It could even be argued that trust is one of the preconditions for theorising about justice, a fact that a few theorists have implicitly acknowledged. John Rawls famously took a definition

[6] Carolyn McLeod, "Trust," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2015 (Metaphysics Research Lab, Stanford University, 2015), https://plato.stanford.edu/archives/fall2015/entriesrust/.
[7] Matthew Harding, "Responding to Trust," *Ratio Juris* 24, no. 1 (March 2011): 75–87.
[8] Ibid., 75.
[9] Eric M. Uslaner, "The Study of Trust," in *The Oxford Handbook of Social and Political Trust*, ed. Eric M. Uslaner, vol. 1 (Oxford University Press, 2017).
[10] Marina Oshana, "Trust and Autonomous Agency," *Res Philosophica* 91, no. 3 (2014): 431–447.
[11] For a defence of a relational conception of autonomy, see Catriona Mackenzie and Natalie Stoljar, eds., *Relational Autonomy: Feminist Perspectives on Automony, Agency, and the Social Self* (Oxford University Press, 2000).

of 'society' as a "cooperative venture for mutual advantage" as his starting point for theorising about justice.[12] Here, the implication is that – before one can start thinking about the principles of justice – one has to first create, or live in, a 'society' in *this* sense. Based on my previous remarks, trust is necessary for cooperation; it seems impossible to consider a model of cooperation without a minimal condition of trust present. As such, it seems that – at least *prima facie* – Rawls' view implies that trust is a precondition for beginning to theorise about justice.

Taken together, these initial remarks show that trust is of paramount importance in achieving a number of goods necessary for a life worth living. Further, trust is ubiquitous in our everyday lives. Yet, despite this – or perhaps because of it – the notion of trust has largely been under-theorised. This is especially true in contemporary political philosophy, where the revival of social contract theory by Rawls influenced many others to describe all relations between individuals in society as if they were two parties in a contract. Put differently, post-Rawlsian political philosophy sought to frame our relations with others in terms of obligations and duties stemming from a contract that we all in some sense agreed to. As such, theorists characterised our relationships in distinctly *impersonal* terms where each party is deemed to be roughly equal.[13] We see this in the very concept of a social contract: All citizens come together to discuss and agree a set of principles to live by. Once this has been done, each party to the contract stands in the same equal relation to everyone else. While undoubtedly a useful concept – especially in specifying the obligations each individual has *qua* citizen – it seems problematic to characterise every kind of relationship in these terms. What happens in cases where a given relationship has a distinct power imbalance, such as between a parent and child or a teacher and student?

In *Trust and Antitrust* – the seminal paper in the contemporary philosophical discussion on trust – Annette Baier posed these exact questions. She argued that the contract

[12] John Rawls, *A Theory of Justice*, Revised Edition (Cambridge, Mass: Belknap Press of Harvard University Press, 1999), 4.
[13] Annette Baier, "Trust and Antitrust," *Ethics* 96, no. 2 (1986): 249–50.

model, where two people of equal standing freely agree on a particular set of actions to perform, cannot account for relationships based on care and cooperation, like those where one party is more vulnerable than the other (e.g. a parent-child relationship).[14] To adequately explain the dynamics of these kinds of relationships, we need to use a different model entirely – namely that of trust. Baier sees trust as a much richer and *personal* concept; she argues that trust is the reliance of one individual on another's goodwill towards them.[15] As such, trust is *normative*: it gives us reason to perform particular actions towards one another. Further, trust is intimately connected to *betrayal*, as well as other reactive attitudes. That is, a given relationship can be characterised as one of trust if – should A breach the trust of B – B feels a sense of betrayal by A's actions towards her.[16] Since Baier's paper, a slew of philosophers have offered different characterisations of the nature of trust. Yet, all of them have in some sense attempted to account for both its normativity and its affective component in their models. The contemporary discussion of trust has, contra the aforementioned contract model, conceived it as being both emotive and personal. No matter what one's preferred theory of trust is, the models deemed most accurate are those that seek to explain the nature of trust in an *interpersonal* setting; that is, a trust relation between two individuals that interact in a non-institutional context.

Recently, theorists have sought to extend their interpersonal model to account for trust in the public sphere. We are frequently being told by the media and other sources that we are in a 'crisis of trust': The 2019 Edelman Trust barometer suggests that the trust inequality between the informed public and the mass population has reached record heights: the population in 15 of the top-26 global markets predominantly distrust their governments.[17] Given these bleak projections, philosophers have sought to offer a more nuanced understanding of the nature of trust in the public sphere – what I call 'public

[14] Ibid., 249–51.
[15] Baier, "Trust and Antitrust"; Paul Faulkner and Thomas Simpson, eds., *The Philosophy of Trust*, First edition (Oxford: Oxford University Press, 2017), 2.
[16] Baier, "Trust and Antitrust," 234–35.
[17] "Edelman Trust Barometer" (Edelman Trust, January 2019), https://www.edelman.com/trust-barometer.

trust' – by viewing it through the lens of their respective interpersonal theory. In this paper, I take issue with this methodological decision. Roughly I argue that, in the same way that the contract model fails to account for the breadth of relations in the moral sphere, extending the interpersonal model to the public sphere obfuscates the diversity of trust relations that are present beyond the interpersonal case. Thus, I call for a reframing of the discussion on trust by suggesting we take an approach that can capture the *impersonal* dynamics of our interactions with others, such as government officials or even the government *simpliciter*. Doing so will shed new light on the subtle ways in which the implementation of new legislation and regulation affects how we trust our governments.

As public trust is instrumentally useful for the promotion and protection of societal goods, any policy or action that seeks to undermine trust ought to be opposed – unless a countervailing reason can be found. Put differently, we have a *pro tanto* reason to implement regulations that prevent the undermining of public trust. I take this to be the starting point of my dissertation. My goal is to show that viewing policy issues through the lens of public trust can aid in finding a compromise position in the event of an intractable policy debate. Though I only focus on one particular policy area, my thesis, in elucidating the relationship between trust and public policy, can perhaps point to avenues of further research in the philosophy of public policy more generally. For instance, we could aim to determine whether we can only have *pro tanto* trust-based reasons to enact a given policy recommendation, or whether public trust can ground a stronger reason to enact the same policy. On a more practical note, we could see how the value of public trust can give us a greater understanding in a number of different policy debates, beyond the one I consider here.

My focus is on the problem of anonymous wealth. A significant amount of global wealth is being stored in bank accounts associated with corporate entities that are designed to keep their owners anonymous. These corporate entities are known as anonymous shell companies (ASCs), which are unlisted companies that have no employees and do not

engage in any substantial business activities beyond the buying, owning, and selling of particular assets.[18] Due to their strong anonymity provisions, ASCs have been used in almost every form of economic crime – ranging from tax evasion to corruption and money laundering. Although it is an area that is receiving significant attention in the policy sphere, ASCs have largely gone unnoticed in the philosophical literature. Most mentions are brief, and only note the potential for anonymous wealth to be taxed for redistributive purposes.[19] Those that do engage with the problem do so as part of a wider project to combat tax evasion. As such, theorists focus on the related problem of tax haven jurisdictions – countries with extremely low tax rates to attract corporations and wealthy elites – and questionable accounting practices like transfer pricing. Put simply, transfer pricing allows corporations to shift their profits to subsidiaries incorporated in tax havens to make it *look* as if their profits are generated in tax haven jurisdictions, even though in reality almost all of their profits are generated elsewhere.[20] Therefore, transfer pricing provides corporations with a legal mechanism through which they evade taxes.

The issue with treating ASCs as part of the fight against tax evasion is that it sidelines the concerns about anonymity in business. That is, this methodological decision assumes that the main reason *why* individuals set up anonymous corporations in tax havens is to take advantage of their favourable tax rates. We see this clearly in the proposed solutions to tax evasion; philosophers like Peter Dietsch have argued for the creation of an international tax body to regulate tax competition in an attempt to prevent countries from setting minimal tax rates.[21] While many wealthy elites use the anonymity provisions of ASCs for tax evasion, there are a number of other purposes ASCs can be used for – such as to facilitate bribery or terrorist financing. Thus, though

[18] Emile van der Does de Willebois et al., *The Puppet Masters: How the Corrupt Use Legal Structures to Hide Stolen Assets and What to Do about It* (Washington, DC: World Bank, 2011).
[19] See, for example, James Henry, "Let's Tax Anonymous Wealth!," in *Global Tax Fairness*, ed. Thomas Pogge and Krishen Mehta, First edition (Oxford: Oxford University Press, 2016).
[20] Gillian Brock, "Taxation and Global Justice: Closing the Gap between Theory and Practice," *Journal of Social Philosophy* 39, no. 2 (June 2008): 165–66.
[21] Peter Dietsch, *Catching Capital: The Ethics of Tax Competition* (New York, NY: Oxford Univ. Press, 2015).

undoubtedly useful in preventing tax evasion, Dietsch's proposal fails to prevent the use of anonymous structures for a variety of other illicit purposes. To this end, I focus on the problem of anonymous wealth, and ASCs in particular, to show that we have an at least *pro tanto* reason to implement more stringent regulations to prevent criminals, terrorist financiers, and even wealthy elites from exploiting our international monetary system for their personal gain. This is not to say that tax evasion and anonymous wealth are not linked; rather, my point is to say that tackling anonymous wealth solely from the perspective of mitigating tax evasion is insufficient.

To understand the extent to which ASCs are used today, and the problems this causes for our global system, we need to examine the political and economic context in which they emerged. In the aftermath of World War One, many European economies were shackled with public debt as a result of financing the war effort. In this environment, and with states promising to compensate both veterans and survivors of the war, many European governments significantly increased the highest marginal tax rate to tap into a hitherto unexplored revenue stream: the assets of wealthy businessmen and the former aristocracy.[22] Further, the industrial revolution had changed the nature of wealth across the globe. While in the past the wealthy elites held fixed assets usually in the form of land, by the 1920s the wealthiest individuals held their assets in the forms of financial securities (i.e. stocks or bonds).[23] These were slips of paper indicating the amount an individual owned of a particular company or of how much they were owed by a corporation or government, and were inscribed with the statement "pay to bearer".[24] As such, anyone who had these slips in their possession was the legal owner of whatever amount they held – which could be millions of dollars. The shift from holding assets in land to financial securities was crucial for two reasons. First, wealth became *liquid*: it was suddenly possible to move assets from one place to another much

[22] Gabriel Zucman, *The Hidden Wealth of Nations: The Scourge of Tax Havens*, trans. Teresa Lavender Fagan (Chicago: The University of Chicago Press, 2015), 9.
[23] Ibid., 9–10.
[24] Ibid., 9.

more easily. Second, it was possible to own vast amounts of wealth *anonymously* – all that was required was a safe place to store all the financial securities one owned.[25]

In this context, the contemporary industry of tax evasion was born – and the country that began this practice is perhaps unsurprising. Though Switzerland had an established banking sector with stringent financial secrecy provisions, initially Swiss banks were primarily inward-looking; that is, as Switzerland lacked the industrial and commercial power of the financial centres of London and Paris, the banks were purely focused on the domestic market.[26] As neighbouring countries increased taxes to considerably higher levels than Switzerland, Swiss bankers saw an opportunity to expand their business and influence. They leveraged their country's banking secrecy laws to attract wealthy clients by promising to handle all their finances with the utmost discretion – regardless of the source. Consequently, foreign capital poured in, causing the banking sector to take-off to such an extent that Switzerland became an international financial centre on the basis of being 'a refuge for foreign wealth'.[27] Until the end of the Second World War, Switzerland's banking industry kept growing as ever increasing volumes of foreign money trickled into the secure vaults of Swiss banks. As they grew in size, the banks broadened the array of financial services they offered: instead of just keeping securities safe, they – amongst many others – became experts at cloaking the identity of account owners.[28] Put differently, they allowed individuals to store their wealth in Swiss bank accounts whilst ensuring that these assets could not be traced back to the individuals that owned them.

In fact, the Swiss banks put this expertise to good use to circumvent the first concerted effort to end banking secrecy. After World War Two, a significant proportion of French wealth stored in Swiss banks was in the form of American financial securities. In an attempt to pressure Switzerland's banking industry, the US Government froze these

[25] Ibid., 11.
[26] Sébastien Guex, "The Origins of the Swiss Banking Secrecy Law and Its Repercussions for Swiss Federal Policy," *The Business History Review* 74, no. 2 (2000): 241.
[27] Ibid., 242.
[28] Zucman, *The Hidden Wealth of Nations*, 26.

assets until the banks released the names of the French owners.[29] Instead of folding, the Swiss banks registered these assets under the names of ASCs incorporated in Panama – a tax haven with lax regulations surrounding company incorporation – so that the true French owners would remain anonymous.[30] Soon after, the US Government unfroze the assets, and Swiss banks began setting up offices in particular tax haven jurisdictions (e.g. Panama, British Virgin Islands etc.) in order to create hubs specialising in specific financial services. In the decades following the Second World War, Swiss financial secrecy laws were exported to small island nations that wanted to attract wealthy elites and multinational corporations in order to find alternative sources of revenue. As a consequence, tax havens became bastions for anonymous wealth; governments in tax havens passed legislation making it cheap and easy to set up shell corporations, while retaining financial secrecy laws that made it very difficult, if not impossible, to determine the owners of these same companies. In some cases, law firms specialising in creating ASCs helped to write the legislation that was then voted on by the governments in tax haven jurisdictions.[31]

This points to the sheer scale of the industry facilitating tax evasion and anonymous wealth. Instead of being confined to one country and their particular banking system, financial secrecy spread across the globe to numerous different jurisdictions in the latter half of the twentieth century. Simultaneously, the volume of wealth held anonymously continued to increase – and this trend shows no signs of abating. Since 2009, our best estimates have shown that the volume of anonymous wealth has exponentially increased – with the vast majority of this growth being attributed to shell companies in non-European tax havens (e.g. Hong Kong).[32] This is a serious issue that deserves our

[29] The reason for this was to ensure that vast amounts of French wealth were subject to US taxes *before* the Government sanctioned the Marshall Plan. For more detail, see Ibid., 19.
[30] Ibid.
[31] Mossack Fonseca, the now-defunct law firm at the centre of the Panama Papers leak, helped to create the offshore financial centre in Niue by drafting financial secrecy legislation that the government approved and signed into law. For more detail, see Michael G. Findley, Daniel L. Nielson, and J.C. Sharman, *Global Shell Games: Experiments in Transnational Relations, Crime, and Terrorism*, Cambridge Studies in International Relations (New York: Cambridge University Press, 2014), 41.
[32] Zucman, *The Hidden Wealth of Nations*, 46–47.

attention; as it stands, in some countries it is possible to incorporate a shell company *anonymously* in a matter of minutes and at very little cost.[33] As such, criminals, money launderers, corrupt dictators, and other reprehensible individuals have utilised ASCs to shield the source of their funds from scrutiny. Thus, tax evasion is only one part of the problem; ASCs facilitate bribery, and allow wealthy individuals to make large anonymous donations to political parties – leading to questions about the impartiality of governments.[34]

Anonymous wealth in general, and ASCs in particular, have been allowed to flourish as a result of loopholes in financial regulation. In the face of every attempt thus far to end financial secrecy, banks, law firms and other corporate service providers – companies specialising in the creation and incorporation of ASCs for clients – have found ways to circumvent the proposed changes. For too long, regulation has been created with the expectation that banks will be honest about the assets of their clients.[35] As we can see from the political and economic context above, this approach is clearly flawed. Banks have assisted their wealthiest clients in ensuring their wealth continued to go undeclared, depriving governments of a significant volume of tax revenue that could otherwise have been put towards redistributive purposes. Recently, however, there has been a shift in the regulatory debate. Backed by public outrage in the aftermath of the Panama Papers leak – a leak of millions of documents from the law firm Mossack Fonseca revealing the scale of the industry behind creating and setting up ASCs – policymakers began calling for more transparency regarding the ownership information of companies incorporated in a particular jurisdiction.

---

[33] A fellow of the Brookings Institute commented that setting up a shell company in the US was "easier than a library card" because, in certain states, an individual did not need to bring any form of ID – see Molli Ferrarello, "One Year after the Panama Papers Leak, Starting a Shell Corporation in the US May Be Easier than Getting a Library Card" (Brookings Institute, April 7, 2017), https://www.brookings.edu/blog/brookings-now/2017/04/07/one-year-after-the-panama-papers-leak-starting-a-shell-corporation-in-the-us-may-be-easier-than-getting-a-library-card/.

[34] On the problems associated with anonymous political donations in the US Context, see Lawrence Lessig, *Republic, Lost: How Money Corrupts Congress - and a Plan to Stop It* (New York, NY: Twelve, 2011); Zephyr Teachout, *Corruption in America: From Benjamin Franklin's Snuff Box to Citizens United* (Cambridge, Massachusetts: Harvard University Press, 2014).

[35] Zucman, *The Hidden Wealth of Nations*, 20.

Specifically, the policy put forward was the construction of a public ownership register. This register would contain information concerning the *beneficial owner* of every company incorporated in a particular country, and would be publicly available for anyone to search. A beneficial owner is the *natural person* who has *de facto* control over the direction and policy of a company.[36] The emphasis on natural person is crucial. Since companies can be owned by other companies, a public register containing merely ownership information would be next to useless as a given ASC would be listed as being owned by *another* ASC in a different jurisdiction. Therefore, there is some consensus amongst policy advisors that creating a public ownership register is the most effective way of putting an end to financial secrecy and its detrimental effects.[37] However, the implementation of such a register in a number of countries has largely stalled as politicians extoll the importance of privacy in the financial sphere. Consequently, the policy debate has largely reached an impasse – thereby maintaining the status quo and allowing ASCs to continue operating with impunity.

The goal of my thesis is to offer an argument in favour of implementing a public ownership register. In short, I argue that we have a *pro tanto* reason to implement such a register on the grounds that it stops ASCs from undermining public trust. My argument proceeds in four parts. In Chapter 1, I offer a detailed presentation of the negative purposes for which ASCs are used, in order to clearly indicate the problems they pose for the international community. I begin with a technical explanation showing how ASCs can render their beneficial owners essentially anonymous, before highlighting the illicit activities they have been implicated in. To do so, I draw on data from a variety of sources in both the academic and policy spheres. Once the problem of ASCs has been established, I offer a philosophical reconstruction of the policy debate surrounding the implementation of a public ownership register. While policymakers agree that ASCs

[36] van der Does de Willebois et al., *The Puppet Masters*, 19–21.
[37] A number of academics and policymakers in international institutions and think tanks have argued in favour of a public ownership register – for example see van der Does de Willebois et al., *The Puppet Masters*; Findley, Nielson, and Sharman, *Global Shell Games*; Zucman, *The Hidden Wealth of Nations*, chap. 4.

facilitate a number of negative externalities, they are divided over which value takes priority in this sphere – privacy or transparency. Proponents of financial secrecy have pointed to a presumptive right to privacy over one's financial information, which counts against making a beneficial ownership register public, while proponents of financial transparency explicitly deny this. Given this intractability, I posit a middle-ground compromise position – a public ownership register where individuals can opt-out of having their ownership information made public – and make a theoretical case for this position over the course of my thesis.

In Chapter 2, I outline the avenue I seek to explore in more detail: the role of trust in society. Having established that trust – and by extension public trust – has instrumental value, the goal of this chapter is to assess the merits of our most prominent theories of trust. I start by offering a distinction between the *form* and *psychology* of trust. Trusting, most theorists agree, involves willingly relying on others to perform particular tasks. As such, theorists focus on trying to explain *why* we are so willing to rely on others. Their explanations focus on two aspects: how a trust relation between individuals can be characterised (what I call the form of trust), and the mental attitude of the truster (what I call the psychology of trust), which is crucial in determining the normative content of trust. After showing how three prominent theories of trust characterise the form and psychology of trust in interpersonal settings, I highlight the ways in which these theories have been extended to account for trust in the public sphere. In doing so, I focus on two prominent views of public trust and highlight the deficiencies they face when tackling impersonal relationships in the public sphere. I close with a discussion concerning the need to move to a model of public trust predicated on social norms.

Having established that prominent theories of public trust fail to adequately explain the diversity of trust relations in the public sphere, I put forward my own account of public trust in Chapter 3. First, I present and defend the need for a distinction in the wider trust literature between personal and impersonal trust, where I take public trust to be a type of impersonal trust. The problem with theories of trust in general, and public trust in

particular, is that they all share the assumption that trust is *monist*; that is, we assume we trust everyone in the same way. I think this is mistaken, and is seen most clearly in the public sphere – where we interact with others in a more impersonal manner than in other contexts. As such, I offer an account of impersonal trust largely inspired by Amy Mullin's view, who argued that trust involves the assumption that both parties have a shared *intrinsic commitment* to a social norm deemed authoritative in a given domain.[38] I use her theory as a starting point, before showing how it can be recast as a theory of public trust that is not vulnerable to the same deficiencies plaguing our other conceptions of public trust. As Mullin adopts a largely intuitive definition of social norms, I offer a plausible interpretation of how she views social norms by drawing on what John Rawls elsewhere called practice rules. Practice rules define a set of procedures that an individual must comply with in order to be seen as being a participant in that practice.[39] I argue that seeing public trust through the lens of practice rules provides us with a morally neutral conception of trust that is beneficial for numerous reasons, and close with a discussion concerning the upshots of this view.

Armed with a fully-fledged theory of public trust, I return to the issue of ASCs in Chapter 4 with the aim of highlighting the mechanisms through which they undermine public trust. I distinguish between two different ways that trust can be eroded: Either the intrinsic commitment to a *particular* practice rule is undermined, or an entire practice is subverted such that it no longer fulfils the purpose for which it was created. I draw on case studies to clearly demonstrate how each method operates. I argue that tax evasion, in which individual practice rules governing taxation are undermined, is a clear case of the former, while the practice of giving anonymous political donations is an example of the latter as norms of political behaviour are being completely subverted. Taken together, this establishes two things: First, that ASCs undermine public trust. Second, that ASCs undermine trust in *pernicious ways*. As my conception of public trust is morally neutral, I need to establish that the norms and practices being subverted are valuable

[38] Amy Mullin, "Trust, Social Norms, and Motherhood," *Journal of Social Philosophy* 36, no. 3 (September 2005): 316–30.
[39] John Rawls, "Two Concepts of Rules," *The Philosophical Review* 64, no. 1 (1955): 25.

in order for my argument to succeed. I conclude with a discussion highlighting the ways in which a public ownership register can stop ASCs from undermining trust, before briefly examining the implications of public trust for public policy more generally.

Before starting, it is worth outlining the methodology I employ over the course of this thesis. I see my thesis as belonging to the methodological school of engaged philosophy, which is an alternative way of pursuing applied philosophy. Where applied philosophy starts with a given ethical theory (e.g. utilitarianism), engaged philosophy takes a bottom-up approach that starts with the policy problem in question. It aims to work in reverse by, first, identifying the key issues that plague a given policy debate and understanding why the debate is becoming intractable (which would involve understanding its history, including what form previous legislation took and why that was successful/unsuccessful), before developing a series of options and evaluating these using a combination of theoretical and practical considerations.[40] Put differently, we can distinguish between two clear stages that the method of engaged philosophy puts forward: understanding and recommending. Here, the former stage involves identifying and framing the policy issue, and drawing on the history of past legislation to do so.[41] From this, the latter stage seeks to develop a catalogue of reasonable options that directly solve the problem specified by the first stage, before putting forward reasons to prefer one particular option over the others. The structure of my thesis parallels this: I start by offering a philosophical reconstruction of the policy debate, before showing why we have a *pro tanto* reason to adopt a public ownership register on the grounds that it prevents ASCs from undermining public trust.

[40] Jonathan Wolff, "Method in Philosophy and Public Policy : Applied Philosophy versus Engaged Philosophy," in *The Routledge Handbook of Ethics and Public Policy*, ed. Anabelle Lever and Andrei Poama, 1st Edition (London: Routledge, 2018), 13–24.

[41] The reason the history of past legislation is important is to prevent the development of 'destructive policy cycles' that, for example, plagues foreign aid policy – see Roger Riddell, *Does Foreign Aid Really Work?*, Oxford: Oxford University Press, 2008.

Chapter 1: The Problem of Anonymous Shell Companies

In 2009, Thai authorities seized an arms shipment masquerading as a supply of drilling equipment. The weapons originated in North Korea and were bound for Iran, in clear violation of the arms embargo imposed on both countries by the United Nations Security Council. After months of research, investigators managed to link the arms shipment to a Chinese national named Lu Zhang who had recently emigrated to New Zealand. Lu Zhang was listed as the sole director of SP Trading, the company that had hired the plane and was the apparent mastermind behind the arms deal. As such, Lu Zhang achieved notoriety amongst the international law enforcement community as a brazen and well-connected arms trader. The reality was very different; Lu Zhang worked as a cook at a Burger King in Auckland and was blissfully ignorant of her infamy. She was entirely innocent and played no role in brokering the arms deal. In fact, to this day, the real perpetrators remain unknown and at large.

To supplement her income at Burger King, Lu Zhang was hired by the GT Group – a shady corporate service provider that specialised in setting up ASCs for their clients. For fifteen dollars each, Lu Zhang agreed to sign any documents placed in front of her by employees of the GT Group. What she did not realise was that each of these documents were official incorporation documents listing her as the sole director of one company. At the time, New Zealand incorporation law required that only one 'nominee director' be named per company. As such, corporate service providers exploited this loophole by using vulnerable individuals like Lu Zhang; they paid her so that she would agree to be nominated as the *de jure* director, while the beneficial owners remained completely anonymous. The terrifying fact about this scheme is that – because of this law – the beneficial owners did not need to submit any information about themselves to the authorities in New Zealand. The paper trail ends with Lu Zhang, leaving the masterminds behind this arms deal free to continue carrying out their actions with impunity.[42]

[42] Lu Zhang was convicted of giving false information by a New Zealand court, yet was swiftly discharged from custody after it became clear that she had no knowledge of any of the crimes perpetrated by any of the companies associated with her. For more detail, see Findley, Nielson, and Sharman, *Global Shell Games*, 1–2.

Essentially, ASCs act as a modern-day Ring of Gyges as they allow individuals to perform a variety of illicit activities with little to no chance of being caught. For this reason, they have been used by oligarchs to purchase undeclared luxury apartments[43], by government officials to facilitate bribes[44], and by state actors to bypass international sanctions[45] to name just a few. Though there is near-universal agreement that ASCs facilitate a number of these negative externalities, there is significant disagreement over the implications this has for regulation in general. Critics of ASCs argue that the negative externalities that arise from ASCs provide sufficient justification for greater transparency. That is, the fact that ASCs are used in almost every form of economic crime gives us a strong reason to create and implement a public ownership register, which would be useful in deanonymizing ASCs. Against this view, proponents of financial secrecy have suggested that making beneficial ownership information public is a violation of an individual's right to privacy. While this argument is often fallacious when employed in the policy sphere, there is nevertheless a charitable interpretation – though it does make the argument weaker than its proponents would hope. In what follows I explore the debate between policymakers over regulating ASCs, before examining philosophical considerations that may illuminate certain aspects that have hitherto gone unnoticed.

I begin by outlining the scale of the problem of ASCs. I do so for two reasons: First, to offer a technical explanation as to how anonymous corporate structures can hide an individual's assets from prying eyes. Second, to highlight the extent to which such structures are used, and thus establish that they pose a significant problem for global development. After delineating the problem, I turn to the regulatory debate concerning whether ownership information about *every* company incorporated in a country's

[43] Louise Story and Stephanie Saul, "Stream of Foreign Wealth Flows to Elite New York Real Estate," *The New York Times*, February 7, 2015, https://www.nytimes.com/interactive/2019/admin/100000003460356.embedded.html?

[44] Global Witness, "Congo's Secret Sales," May 2014, https://www.globalwitness.org/en/campaigns/oil-gas-and-mining/congo-secret-sales/.

[45] Findley, Nielson, and Sharman, *Global Shell Games*, 37–38.

jurisdiction ought to be made publicly available. I offer a philosophical reconstruction of this debate to show why it has become intractable. In short, opponents of a public ownership register have pointed to the existence of a presumptive right to privacy that covers the disclosure of beneficial ownership information – something that proponents of a public register explicitly deny. Consequently, each side differs as to which value (privacy or transparency) ought to be promoted, and – in the absence of a decision-making procedure that can show one value to be superior over the other – has caused the debate to reach an impasse. I conclude the chapter by considering how the value of public trust can point to a middle-ground position that may help in resolving the policy debate.

The Scope of the Problem

We have seen that the very purpose of ASCs is to cloak the identity of their beneficial owners. For this reason, attempting to calculate the amount of wealth squirreled away in ASCs across the globe is an arduous task – though it has not stopped some from trying. In 2012, the American economist James Henry made headlines by suggesting the amount of financial wealth in hidden jurisdictions was between 21 and 31 trillion US dollars.[46] It is worth emphasising two points: First, this estimate *excludes* material assets such as yachts or villas that are also frequently owned through ASCs. Second, this wealth has virtually never been taxed, and continues to grow as it is invested – again tax-free – in economic opportunities across the globe. Henry's estimate however has faced methodological criticism; some have argued that his calculations are too coarse-grained as they include legitimate bank accounts held by corporations in multiple different jurisdictions.[47] As such, other economists like Gabriel Zucman have suggested the figure is lower at around $7.6 trillion – which is equivalent to 8% of global GDP. Again, it is worth noting that this is a *minimal estimate*. Like Henry, Zucman's calculations do not

[46] James Henry, "The Price of Offshore Revisited: New Estimates for 'Missing' Global Private Wealth, Income, Inequality, and Lost Taxes" (Tax Justice Network, July 2012), 5.
[47] For a more detailed criticism of Henry's calculation, see Zucman, *The Hidden Wealth of Nations*, 40–42.

take all forms of offshore wealth into account, since we simply lack the data to make a better estimate.

In any case, even our lowest estimates of the wealth held anonymously highlight the sheer scale of ASC use today – and there is evidence to show that this trend is only increasing. The volume of wealth held in tax havens has increased by 25% between 2009 and 2014.[48] Further, if we examine the source of this wealth, we see the issues ASCs pose for global development: 30% of all financial wealth generated across Africa is being held anonymously in tax havens, amounting to a tax loss of $14 billion.[49] Of course, tax havens are not equivalent to ASCs, as many corporations register subsidiaries in tax haven jurisdictions for the purpose of being tax efficient.[50] However, they are nevertheless linked; *anonymous* and *undeclared* wealth in tax havens is often registered as being owned by an ASC, while the incorporation rates of ASCs in certain tax haven jurisdictions have been increasing since 2009.[51] It is worth reemphasising that not all ASCs in tax haven jurisdictions are used for the purpose of tax evasion- they are also used for a number of other illicit activities. For example, the NGO Global Witness uncovered a corruption scandal in the Democratic Republic of Congo, where government officials sold a series of valuable state-owned mining assets to a number of different ASCs registered in the British Virgin Islands – a notable tax haven.[52] These shell companies then resold the mining assets to multinational corporations for a massive profit, which was never seen by the citizens of DR Congo. It is estimated that this sale contributed to a loss of $1.3 billion – almost double the *combined* national budgets for health and education.[53] Once again, due to the anonymity provided by ASCs, it was impossible to determine the beneficial owners of these ASCs.

[48] Zucman, *The Hidden Wealth of Nations*.
[49] Ibid., 53.
[50] For a compelling argument against such practices, see Brock, "Taxation and Global Justice."
[51] Zucman, *The Hidden Wealth of Nations*, 47.
[52] Global Witness, "Congo's Secret Sales."
[53] Ibid.

How can ASCs operate in this manner? Although they are not listed on any stock market and do not engage in any substantial business activity, ASCs can perform nearly all the same activities that people can. This is also being reflected in US law, where a number of recent rulings (e.g. *Citizens United v. Federal Electoral Commission*) have given corporations a number of rights that are normally associated with humans – such as freedom of speech. Furthermore, ASCs ensure anonymity in a multitude of ways. Typically, ASCs are set-up by a corporate service provider (CSP) – a company that specialises in setting up shell companies (anonymous or otherwise) – who offers a variety of services to their clients that can ensure anonymity. One instance may be refusing to accept identification documents from the client, thereby ensuring that no traceable information would appear on the company's corporate record *or* in the files of the CSP. This makes it impossible to trace anything back to the owner in question, as there would be no record of them having opened a company. Other anonymity services include nominee shareholders, where the CSP lists a particular individual (usually an employee or lawyer) as having *de jure* control of a company. Here, the nominee shareholder acts under instruction from the client (who has *de facto* control) without the client being traceable or held liable for any of the company's actions.[54] In more extreme cases, some CSPs offer a full 'reception service' where an employee of the CSP acts as a receptionist on behalf of the client's company.[55]

Though failing to collect proof of identity renders a shell company essentially untraceable, anonymity can to a significant extent be guaranteed even when this information *is* collected. The basic premise here is the following: Capital can cross borders easily, while laws cannot. As such, ASCs can be set up in jurisdictions where it is incredibly difficult to either gain access to corporate records or file a lawsuit in an attempt to reclaim the relevant assets (or both). For instance, in St. Kitts and Nevis, one

[54] Nominee shareholders allowed SP Trading – an anonymous shell company based in New Zealand – to escape all punishment for their role in facilitating an arms trade between North Korea and Iran, while both nations were under an international arms embargo. For more detail, see Findley, Nielson, and Sharman, *Global Shell Games*, 1–2.
[55] Oliver Bullough, *Moneyland: Why Thieves & Crooks Now Rule the World & How to Take It Back* (London: Profile Books, 2018).

has to file a bond of $100,000 with the court just to *begin* legal proceedings, in order to show that the case is not frivolous.[56] So, even if a CSP collects information about a client, there are a number of legislative barriers that make it extraordinarily difficult and costly for any investigator to legally obtain this information. To further compound this problem, ASCs can be owned by other ASCs incorporated in jurisdictions where it is equally difficult to obtain ownership information.[57] To illustrate, suppose that I set up a company, A, and register it in St. Kitts and Nevis. Suppose further that I set up a second company, B, in Panama and list this second company as being a majority owner (i.e. owning more than 50% of all shares) of A. In principle, I could keep doing this *ad infinitum* as there is no limit to the number of companies any given shell company can own.

Suppose that I set up three more companies, with each new company being the owner of the previous company (as above) and with each company being registered in a different jurisdiction that is known to be unfriendly to international law enforcement. We now have the following ownership structure, where -> denotes 'owner of':

Me -> E -> D -> C -> B -> A

Here, I have created a corporate structure that effectively grants me anonymity from any actions I choose to perform. I might use A to offer bribes to government officials in order to circumvent, for instance, environmental legislation that I might see as being bothersome. Suppose I do this. Any investigator thinking that this payment is suspicious would have to access the St. Kitts and Nevis corporate registry to obtain any ownership information, which as indicated above is *not* easy, in order to ascertain its source. Let's imagine they are successful. They access the corporate records only to discover that the majority owner of A is B, another shell company registered in another tax haven country (in this case Panama) that has implemented rigorous legislation to maximise the

[56] Ibid.
[57] van der Does de Willebois et al., *The Puppet Masters*.

difficulty of obtaining information through the corporate registry. Even if the investigator were once again successful, the process would repeat itself at every stage in the aforementioned corporate structure – making it incredibly costly and time-consuming to uncover the beneficial owner of any ASC, even if some ID is collected as soon as a company is set up.[58]

Now, it might be objected that such methods violate international standards; more precisely, corporate service providers that do *not* collect any proof of their client's identity would fall foul of international money laundering standards. These were set up by the Financial Action Task Force (FATF), an intergovernmental organisation tasked with combatting money laundering and terrorist financing. In 2000, they constructed a 'blacklist' of countries that were deemed non-compliant.[59] Although there were no formal sanctions, countries placed on the blacklist found themselves under increased scrutiny and faced significant financial pressure to comply with these standards. As such, by 2007, no country was deemed non-compliant. Thus, it seemed – at least *prima facie* – that the problem was solved. But this did not show the extent of the issue; in a recent study, three academics posed as consultants who were seeking to incorporate shell companies from a number of different corporate service providers (CSPs) based across the world. They found that, overall, 48% of all responses failed to ask for appropriate documents – of which 22.1% didn't ask for *any* form of photo identification.[60] Further, contrary to popular belief, compliance with international standards was higher in tax havens than in OECD countries.[61] In short, we are witnessing a situation in which ASCs continue to facilitate illicit monetary flows at increasing levels with each passing year, thereby posing significant problems for proponents of global development.

[58] For a real-life example of this phenomenon, consider the story of Sarah Pursglove's divorce from the Finnish tech entrepreneur, Robert Oesterlund. In short, Oesterlund created a gigantic web of ASCs and trusts to hide his assets totaling over $400 million in order to minimize both his tax contributions *and* his divorce payouts. For more detail, see Nicholas Confessore, "How to Hide $400 Million," *The New York Times Magazine*, November 2016.
[59] Findley, Nielson, and Sharman, *Global Shell Games*.
[60] Ibid.
[61] Ibid.

Given the myriad ways in which ASCs facilitate illicit and often illegal activity, why have they remained unregulated for so long? Part of the answer may be that most governments assumed the FATF money laundering standards were successful – even if, in practice, this was not the case and multiple CSPs flouted these standards. Further, governments may either have been unaware of the scale of the problem or were complicit in preventing stronger transparency laws from being passed. The 2016 Panama Papers leak showcased the scale of the problem facing policymakers campaigning against ASCs. This leak, which contained 11.5 million documents from the Panamanian law firm and CSP Mossack Fonseca, highlighted the extent to which the global elite abused tax loopholes and corporate structures to perform a number of functions. From tax evasion and money laundering to facilitating corruption, this leak illuminated the seedy underbelly of the international financial system and showcased the myriad ways in which wealthy individuals used their money for illicit purposes. Crucially, the leak also exposed prominent politicians, showcasing how they had stashed their wealth in tax haven jurisdictions in order to avoid the fiscal obligations in their own countries.

Three years on, the information in the Panama Papers spawned investigations in over 82 countries and initiated a flurry of resignations by government officials – including the Prime Ministers of both Iceland and Pakistan.[62] Over $1.2 billion were recovered by tax authorities, and more stringent regulatory and due diligence measures began to be adopted in order to increase the difficulty of money laundering and tax evasion.[63] One measure that has been hailed as a resounding success is the creation and implementation of a public ownership register. In 2016, the British Government launched the register of Persons with Significant Control (PSC), a free and user-friendly database that gives any individual with an internet connection the ability to access

[62] Amy Wilson-Chapman and Douglas Dalby, "Panama Papers Helps Recover More than $1.2 Billion around the World," *ICIJ* (blog), 2019, https://www.icij.org/investigations/panama-papers/panama-papers-helps-recover-more-than-1-2-billion-around-the-world/.
[63] Ibid.

information about the beneficial owners of any company incorporated in the UK.[64] Similarly, the British Government is in the process of implementing legislation that would force its Overseas Territories, such as the British Virgin Islands, to implement a public ownership register. Such registers are crucial for the regulation of ASCs, as they strip ASCs of their absolute anonymity and improve financial transparency in virtue of their public accessibility.

While these are positive signs, they remain insufficient. Though the UK, by being one of the first countries to implement a public ownership register, is one of the global leaders in the fight against ASCs, the PSC register still suffers from a number of significant problems that undermine its effectiveness. Of these, the lack of a verification mechanism to ensure that the information contained within the PSC register is accurate and up-to-date is perhaps the most severe and could sabotage its entire purpose. Without the ability to correct erroneous information, or remove false records, the PSC register risks being flooded with aliases – such as one 'Mr. Xxx Stalin' a purported French resident in East London.[65] Further, Britain's attempt to enforce financial transparency on its Overseas Territories has been met with consistent delays – resulting in significant criticism by members of the Opposition.[66] Ultimately however, the main obstacle to regulating ASCs is that positive steps, like the establishment of a public ownership register, are confined to a few proactive countries rather than being universally adopted. As long as there are countries that consistently fail to improve financial transparency, there will be a jurisdiction in which ASCs and the illicit monetary flows they facilitate can thrive.

---

[64] Global Witness, "Learning the Lessons from the UK's Public Beneficial Ownership Register," October 2017, https://www.globalwitness.org/en-gb/campaigns/corruption-and-money-laundering/learning-lessons-uks-public-beneficial-ownership-register/.

[65] Strictly speaking, it is possible that Mr. Stalin's first name is 'Xxx', though this seems unlikely. For more examples that highlight this issue, see Oliver Bullough, "How Britain Can Help You Get Away with Stealing Millions: A Five-Step Guide," *The Guardian*, July 5, 2019, https://www.theguardian.com/world/2019/jul/05/how-britain-can-help-you-get-away-with-stealing-millions-a-five-step-guide.

[66] It is worth noting that the implementation of the public register has been frustrated by numerous delays – see Patrick Wintour, "MPs Attack Ministers over Delay to Tax Havens' Public Registers," *The Guardian*, January 11, 2019, https://www.theguardian.com/world/2019/jan/11/mps-attack-ministers-over-delay-to-tax-havens-public-registers.

The reason public ownership registers have, to date, only emerged in three countries is largely due to a regulatory impasse. While policymakers agree that ASCs have been used in a variety of economic crimes, they disagree over how they should be regulated. As critics of ASCs extol the virtues of public ownership registers in deanonymizing ASCs, proponents of financial secrecy argue that these regulations infringe an individual's right to privacy and thus oppose their implementation. Advocates of ASCs, it seems, think that the burden lies with the accuser to show that a particular ASC is implicated in wrongdoing before ownership information can be divulged. Opponents of ASCs however explicitly deny this; they argue that the only way in which we can implicate ASCs in wrongdoing is by having access to beneficial ownership information. What emerges is an intractable debate where each side aims to prioritise one value – be it transparency or privacy – at the expense of the other. In the subsequent section, I offer a philosophical reconstruction of this debate in order to pinpoint the reason behind its intractability. From this, I briefly sketch out a potential compromise between both positions and aim to argue in favour of this by considering a value that has hitherto been unnoticed – namely the notion of public trust.

Privacy v. Transparency: The Regulatory Impasse

Proponents of financial secrecy argue that ASCs are integral to the protection of the privacy rights of (wealthy) individuals. For example, the Panama Papers showed that Emma Watson owned all her property through an ASC. She did so in order to ensure her address remained a secret form her fans, who may have followed or stalked her.[67] Similarly, many wealthy individuals argue that making corporate information opaque and hard to access prevents any individual from determining exactly how much they are worth. This has numerous advantages for their personal security, as it would prevent kidnappers or other criminal gangs from attempting to harm them and extort them for

[67] Bullough, *Moneyland*.

their money.[68] In such cases, this 'privacy defence' seems perfectly reasonable. Further, this view is explicitly endorsed by supporters of ASCs. Take the Wyoming Secretary of State. In response to the Federal Government calling for the release of beneficial ownership information in the aftermath of the Panama Papers leak, his office released the following statement:

"We are not naive as to the importance of the release of these 'Panama Papers,' but we will not compromise the privacy of our customers".[69]

It seems, then, that financial secrecy is justified on the grounds that it infringes an individual's right to privacy – something that cannot be taken lightly. However, although important, defining a right to privacy has proved to be an almost insurmountable challenge.[70] Let us consider what such a privacy right could look like. To sidestep the controversies in this debate, I follow Annabelle Lever in seeing privacy rights as primarily, but not exclusively, rights of solitude, intimacy, and confidentiality.[71] In this case then, proponents of ASCs are backing a stringent right to confidentiality in order to prevent the disclosure of beneficial ownership information. Still, it is unclear what is meant by disclosure. If it simply refers to the disclosure of information to any institutional body, then this becomes a right against even the *collection* of ownership information by government officials. Put differently, if proponents of financial secrecy seek to prevent the disclosure of ownership information to regulatory authorities –

[68] Frederik Obermaier and Bastian Obermayer, "Oligarchs Hide Billions in Shell Companies. Here's How We Stop Them," *The Guardian*, April 3, 2018, https://www.theguardian.com/world/commentisfree/2018/apr/03/public-registries-shell-panama-papers.

[69] Casey Michel, "The U.S. Is a Good Place for Bad People to Stash Their Money," *The Atlantic*, July 13, 2017, https://www.theatlantic.com/business/archive/2017/07/us-anonymous-shell-companies/531996/.

[70] The reason for this is that privacy has been understood differently in different contexts. For instance, in the USA, the Supreme Court enshrined a constitutional right to privacy in the 1965 *Griswold v. Connecticut* ruling, which was understood as a zone covering intimate actions and decisions where the state could not intervene. This constitutional right however is not what is meant here, where privacy covers the control over information that people possess. For a useful overview of this distinction, see Judith DeCew, "Privacy," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, 2018, https://plato.stanford.edu/archives/spr2018/entries/privacy/.

[71] Annabelle Lever, "Feminism, Democracy and the Right to Privacy," *Minerva* 2005, no. 9 (2005): 11.

seeing this as a violation of privacy – then this would prevent the relevant institutional body from collecting ownership information about the companies incorporated in their jurisdiction. Based on public statements from officials in tax havens, this is the kind of privacy right they have in mind.

Such a stringent confidentiality right has a high burden that arguably has not been met; in fact, such a right seems to conflate a descriptive conception of privacy with a normative conception.[72] Here, a descriptive conception refers to instances where privacy can be gained or lost, whereas a normative conception frames privacy in terms of a *right* that can be violated by others.[73] On this reading of disclosure, it would need to show that the *collection* of beneficial ownership information constitutes a privacy violation. Yet, it seems strange to hold that collecting and storing beneficial ownership information in a government database accessible only to authorised personnel would constitute a privacy violation. This bears structural similarities to the NHS database where the public cannot access medical information but doctors (i.e. authorised personnel) can. We can explain this in terms of the descriptive/normative distinction: when a doctor accesses my medical record, I suffer a *loss* of privacy, yet when my neighbour Bill accesses my medical record, I suffer a *violation* of privacy. Therefore, on this interpretation, the burden of proof shifts back to proponents of ASCs to show that collecting beneficial ownership information constitutes a privacy violation.

We can weaken the right to confidentiality by understanding disclosure as referring to the release of beneficial ownership information to the public. As such, the privacy defence suggests that making beneficial ownership information public is a violation of privacy, though the release of this information to government officials would be permissible. What could the basis of such a right be? An intuitive answer is that this information is part of the private sphere, and therefore cannot be demanded. Here, I

[72] This argument has been used to criticize traditional theories of privacy – see Herman T. Tavani, "Philosophical Theories of Privacy: Implications for an Adequate Online Privacy Policy," *Metaphilosophy* 38, no. 1 (January 2007): 4–5.
[73] Tavani, "Philosophical Theories of Privacy."

take the private sphere to refer to an area encompassing one's right to solitude, intimacy, and confidentiality where nobody else can interfere without permission. This answer seems to be in some sense endorsed by the statement from the Wyoming Secretary of State, who – as a government official – recognises an area that he cannot interfere in without good reason. Yet, the concept of a public-private distinction in the case of privacy is a gross oversimplification. If such a distinction *were* to exist, then it would point to the existence of a domain (i.e. the public sphere) in which *any* information is up-for-grabs, as privacy only exists in the private domain.[74] This fails to account for interactions in the public sphere that are best explained in terms of privacy violations – for instance an individual staring at someone else on the London Underground. Therefore, understanding a right to privacy in terms of a public-private distinction fails, since it does not accurately track the ways in which we ordinarily use the term privacy.[75]

What would be required, then, is a right to privacy coupled with a justification that is *context-dependent*. That is, the right to privacy needs to be specified in such a way that it can adapt to a number of different contexts rather than in terms of a sphere covering all the actions and/or decisions one can make without interference by external forces. A plausible route that satisfies this requirement is offered by Annabelle Lever. She argues that a right to privacy can be justified through democratic principles. Roughly, she argues that democracies are committed to the protection of both freedom and equality in both the personal and political domains. Put differently, a government's action will be deemed legitimate if it is used in the process of safeguarding the freedom and equality of its subjects – in both the political (e.g. voting rights) and personal (e.g. sexual equality) domains.[76] Applying these democratic principles to privacy gives us two justifications for privacy rights: a personal and a political one. Therefore, on Lever's view,

[74] Helen Nissenbaum, "Respecting Context to Protect Privacy: Why Meaning Matters," *Science and Engineering Ethics* 24, no. 3 (June 2018): 831–52.
[75] Feminists have also been vocal critics of the conception of a private sphere, arguing that an area of government non-intervention allows domestic abuse and the subjugation of women in the domestic sphere to be covered up. For a compelling argument, see Catharine A. MacKinnon, *Toward a Feminist Theory of the State* (Cambridge, MA.: Harvard University Press, 1989).
[76] Lever, "Feminism, Democracy and the Right to Privacy," 10.

privacy rights are justified as they help protect one's freedom and equality in each of the aforementioned domains. Although it is beyond the scope of this thesis to offer a full defence of each, I will briefly sketch out both justifications to show how they may be useful to proponents of ASCs.

The political justification for privacy has two components. First, privacy rights are instrumentally useful in promoting political participation and allowing individuals to campaign for the protection of their legitimate interests.[77] Second, protecting privacy rights is intrinsically valuable as it safeguards the rights that democracies – at least in principle – are required to defend.[78] To illustrate, consider the Supreme Court ruling in the NAACP v. Alabama case from 1958. Here, the Alabama legislature subpoenaed the NAACP for their membership records in order to reduce their influence in the state.[79] In voiding the subpoena, the Supreme Court ruling made explicit reference to the importance of protecting the privacy rights of individual NAACP members. Their reasoning satisfies both components of Lever's political justification of privacy rights. For one, protecting the privacy rights of individual NAACP members is instrumentally useful for them to continue their fight against segregation and the Jim Crow laws, since they will not be subject to random arrest, harassment, intimidation etc. Further, protecting privacy rights is also intrinsically valuable as it prevents other democratic rights – such as freedom of association – from being undermined.

The personal justification is structured in the same way. Protecting privacy rights is instrumentally useful in promoting one's personal freedoms.[80] By only disclosing information that we would like others to hear, we in a sense control how we are perceived by others which gives us a considerably degree of flexibility in our interactions with others.[81] Protecting privacy rights also has intrinsic value as privacy gives us the ability to give meaning and nuance to our interactions with others; thus willingly

[77] Ibid., 12.
[78] Ibid.
[79] Ibid., 14–15.
[80] Ibid., 20.
[81] James Rachels, "Why Privacy Is Important," *Philosophy and Public Affairs* 4, no. 4 (1975): 323–333.

disclosing personal information to certain individuals becomes a meaningful act and allows us to develop intimate relationships.[82]

On this basis, we can justify why – for example – the privacy rights of women in the workplace should be protected, so that they are not required to disclose information about their marital status or family planning to their employers.[83] Applying this to the case of ASCs, we can see how the privacy defence might work. In cases like Emma Watson, protecting her privacy rights seem necessary in order to ensure that she has the personal freedom to pursue her own projects and goals. Generalising from this case, the privacy defence would have to show that the public disclosure of beneficial ownership information *necessarily* undermines an individual's personal freedom or equality in some way, for instance by increasing the risks wealthy elites have to being kidnapped and extorted for money. If true, then this would constitute a privacy violation and would justify preventing the disclosure of beneficial ownership information. However, the privacy defence would not stretch to every use of ASCs. Consider the case of anonymous political donations; here, anonymous wealth in politics would subvert the political justification for privacy rights and therefore would mandate the creation of a public donor list.

The privacy defence, then, can at most ground a *presumptive right to privacy*. That is, the privacy defence can claim that *ceteris paribus* the public disclosure of beneficial ownership information constitutes a violation of privacy, and therefore public ownership registers should not be implemented. However, if it can be proved that a particular ASC was involved in some wrongdoing, then the privacy defence melts away as the wrongdoing in question would violate either the political or personal justifications for privacy. Thus, proponents of ASCs claim that the onus is on investigators to prove that a given ASC has been implicated in some kind of crime before beneficial ownership information can be disclosed. Now that we have examined the privacy defence in more

[82] Lever, "Feminism, Democracy and the Right to Privacy," 21.
[83] Ibid.

detail, we are in a position to understand why the policy debate has become so intractable. Critics of ASCs have argued that public ownership registers should be implemented purely on the grounds that ASCs facilitate a number of negative outcomes. Further, the transparency provided by a public ownership register is important for the promotion of anti-corruption initiatives as well as for the promotion of democratic accountability. Therefore, a public ownership register would be effective in preventing the negative externalities associated with ASCs.

To see why transparency is invoked in this context, consider the problem policymakers are trying to solve. ASCs protect anonymous wealth; that is, they ensure that undeclared capital remains hidden from prying governments. Therefore, the only way to regulate their use and prevent money laundering and/or outright theft of resources is to make the entire industry of shell companies transparent. This is not controversial; many structurally similar arguments are used to justify more mundane policies, such as requiring public officials to declare any potential conflict of interests before taking office or safeguarding the right of ordinary citizens to make freedom of information requests. In all of these cases, the right individual citizens have to transparency comes from its instrumental value in ensuring that officials can be held accountable for their actions.[84] As Louis Brandeis, former Supreme Court Justice, once remarked "sunlight is said to be the best of disinfectants".[85] If individuals know that their organisation is transparent such that they will be held accountable for any wrongdoing, their behaviour will change accordingly.

However, framing the debate in terms of privacy against transparency creates an unbridgeable divide, since the former is often viewed as a limit to the latter and vice versa. To illustrate, consider a freedom of information request. This allows citizens to request the disclosure of a piece of information deemed in the public interest – unless

[84] Adam Etizioni, "The Limits of Transparency," in *Transparency, Society and Subjectivity: Critical Perspectives*, ed. Emmanuel Alloa and Dieter Thomä (Cham: Palgrave Macmillan, 2018), 179–202.
[85] Louis D. Brandeis, *Other People's Money, and How the Bankers Use It* (New York : Stokes, 1914).

doing so would constitute a *privacy violation*.[86] Conversely, forcing public officials to declare any conflict of interests before they take office involves a loss of privacy on their part, yet is necessary to protect the principle that governments are only accountable to the people. Therefore, applying the privacy-transparency debate to the case of ASCs only serves to reflect this impasse. By highlighting the value of transparency – and by extension public ownership registers – in combatting corruption and increasing accountability, critics of ASCs explicitly deny that individuals even have a presumptive right to privacy. Similarly, proponents of financial secrecy are not moved by appeals to transparency, since they see an individual's presumptive right to privacy as trumping these considerations. To avoid confusion, it is worth reiterating that both sides agree that ASCs do facilitate a number of illicit financial activities. So, any appeal to the harms caused by ASCs will be fruitless in resolving this debate. Thus, the regulatory gridlock we are witnessing is being driven by the inability of policymakers to put forward clear conditions to decide when privacy ought to be preferred over transparency and vice versa.

It seems clear that each side in this debate has a reasonable case for their own positions. After all, it is intuitively plausible to say that Emma Watson should be exempt from publicly disclosing her assets on privacy grounds. Yet, at the same time, as ASCs are being used to illegally avoid taxation and to facilitate other illicit financial flows, this by itself provides a significant reason to implement a public ownership register, which appears to be the most effective way to combat ASCs. If we examine this debate from the perspective of a third party observer, a middle-ground position seems to emerge. Perhaps we could implement a public ownership register where the names of beneficial owners are disclosed, yet where unique identificatory information (e.g. a full date of birth) is excluded to protect an individual's presumptive right to privacy. Alternatively, individuals could apply to have their information withheld from the public register if doing so would infringe their presumptive right to privacy, though they may still have to

[86] Etizioni, "The Limits of Transparency," 188. It is also worth clarifying that there are a number of reasons to deny a freedom of information request that are unrelated to privacy (e.g. national security). I omit these as they are irrelevant for my purposes.

disclose their assets to the relevant government institution. Though this proposal seems plausible, as it stands there are no philosophical reasons to adopt this approach: it would merely be a political compromise. If however, we examine this debate through the lens of a different value – namely public trust – we may in fact have a theoretical basis to adopt this compromise position.

For too long, our best philosophical theories have simply assumed that trust exists between all citizens – an assumption that obfuscates the breadth of interactions and relationships that occur between individuals in any given society.[87] In the context of ASCs, trust exists in a variety of ways; we may trust that our fellow citizens, as equal members of a democratic society, are paying their fair share of taxes. We may trust our governments to shelter us from the excesses of the criminal underworld, or to stop others infringing our rights to pursue our own projects. Such trust is valuable for the development of democratic societies – without it, simple tasks that require minimal cooperation could not be performed (or could only be performed with great difficulty). I aim to argue that the presence of ASCs in society undermines public trust and therefore ought to be regulated on these grounds. For my argument to be successful, two key desiderata must be met. First, it needs to be sensitive to the sheer diversity in use that ASCs have. As ASCs have been used in almost every form of economic crime, any attempt at arguing for their prohibition must apply across all of these domains. Second, it needs to be sensitive to the positive aspects of shell companies as well. In cases like that of Emma Watson, it seems as if ASCs do play a useful role in safeguarding privacy rights. If these desiderata are met, then this gives us a *pro tanto* reason to implement a public ownership register, with the relevant adjustments listed above.[88]

[87] Baier, "Trust and Antitrust," 248.
[88] Henceforth, all references to a public ownership register denote a public ownership register with the amendments outlined in this section, unless otherwise specified.

Chapter 2: The Limitations of Theories of Public Trust

In the film *Titanic*, just before Jack helps Rose stand on the prow of the ship, he asks her the following question: "Do you trust me?". Though this scene is undoubtedly iconic, there is nevertheless a trope in modern cinema where the (usually male) protagonist utters these exact words to their (usually female) companion before performing some action or manoeuvre that will place them at risk. Sometimes these scenes can be romantic, as in *Titanic* or *Aladdin*. Other times, they have the feel of a tired cliché. Nevertheless, common to all of these scenes is the presence of a *normative* concept of trust. That is, in each of these cases, the attitude of trust possessed by the trusting party contains some normative content. Presumably, if Jack dropped Rose halfway through supporting her on the prow, she would feel a sense of betrayal or some other reactive attitude, which by its presence points to some normative content.[89] This is notably different from a thinner non-normative concept of trust based on a prediction (or belief) that the trusted party will perform a given action. So, when we say that we trust our computers to turn on or our cars to start, the trust relation that exists would be of this thinner variety, as neither cars nor computers can display reactive attitudes.

This distinction between a thin and thick conception of trust is common in the literature. While social scientists have almost exclusively focused on the non-normative conception of trust, philosophers have – for the most part – sought to theorise about the normative conception of trust and emphasise its importance in our everyday lives. Thus, the philosophical discussion of trust has focused on the link between trust and reactive attitudes like betrayal, as well as on the significance of trust in giving us meaningful relationships with others.[90] In doing so, philosophers have characterised trust in terms of motives of 'goodwill'[91], the honouring of commitments[92], and the act of directly

[89] Peter F. Strawson, "Freedom and Resentment," *Proceedings of the British Academy* 48 (1962): 1–25.
[90] E.g. see Baier, "Trust and Antitrust"; Karen Jones, "Trust and Terror," in *Moral Psychology: Feminist Ethics and Social Theory*, ed. Peggy DesAutels and Margaret Urban Walker, Feminist Constructions (Rowman & Littlefield Publishers, 2004); Harding, "Responding to Trust."
[91] Baier, "Trust and Antitrust."
[92] Katherine Hawley, "Trust, Distrust and Commitment," *Noûs* 48, no. 1 (March 1, 2014): 1–20.

responding to the reliance of the truster[93] to name just a few. However, the philosophical discussion has solely focused on trust at the individual level; that is, philosophical theories of trust have examined the nature of trust between individuals treated in isolation from the structure of their respective societies. Recently there has been a shift to account for trust at the societal level, with theorists of trust looking to extend their models of individual trust to provide an account of the nature of public trust. As we will see, I take issue with this methodological decision as it obfuscates the diversity of trust relations that exist at the societal level.

My focus is on the value of public trust, which I take to refer to the trust relations that exist between individuals as role-occupiers as well as between individuals and their government.[94] I aim to argue that – because ASCs undermine public trust – we have an at least *pro tanto* reason to implement a public ownership register. We have already seen that public trust is instrumentally valuable in achieving a number of high value goods. The next step in establishing this argument is to provide an account of public trust that adequately captures the nature of trust in the public sphere. This is the goal of the next two chapters. I start by offering a brief literature review of the philosophy of trust. I see the debate about the nature of trust as seeking to account for two aspects of any given trust relation: the form and psychology of trust. After defining each of these aspects, I use this distinction to show how different theorists have characterised the nature of trust. This is useful for two reasons: First, it clarifies the terminology that will be used for the remainder of this thesis. Second, all of the theories of public trust offered so far extend their analysis of trust *simpliciter* to the public sphere. Thus, it would be impossible to understand the nature of public trust without first understanding the nature of trust.

[93] Paul Faulkner, "The Attitude of Trust Is Basic," *Analysis* 75, no. 3 (July 2015): 424–29.
[94] Patti Tamara Lenard, "Trust Your Compatriots, but Count Your Change: The Roles of Trust, Mistrust and Distrust in Democracy," *Political Studies* 56, no. 2 (June 2008): 312–32; Patti Tamara Lenard, "The Political Philosophy of Trust and Distrust in Democracies and Beyond," *The Monist* 98, no. 4 (October 2015): 353–59; Faulkner, "Finding Trust in Government."

In the second half of this chapter, I show how our models of individual trust have been extended to the public sphere. Specifically, I consider two theories of public trust – one offered by Paul Faulkner and the other by Cynthia Townley and Jay Garfield – that are prominent in the literature. In each case, I raise a number of criticisms to cast doubt on the theory's ability to adequately capture the trust relations that exist in the public sphere. Against Faulkner, I argue that he fails to explain the observable evidence showing the presence of reactive attitudes present amongst individuals when their governments breach their trust. Against Townley and Garfield, I argue that their theory of public trust suffers from a lack of clarity that makes their characterisation of trust relations ambiguous at best. I evaluate a number of different interpretations of their theory, before ultimately concluding that each interpretation is inadequate in accounting for the diversity of trust relations in the public sphere. I close the chapter by suggesting a hitherto under-explored route that may be helpful in developing an account of public trust. In short, I argue that the most promising theory of public trust grounds trust in a shared commitment to a social norm deemed authoritative within a given domain.

What is Trust?

Trust is often taken to be a form of reliance on another person, group, or institution. Understood that way, trusting is both an act and attitude; when we trust someone, we rely on them in some manner and, crucially, we do so *willingly*.[95] Theories of trust then seek to explain this willing attitude that we possess in trusting others. Put differently, a key requirement of theories of trust must be to explain why we are so willing to rely on others to act in certain ways. As such, theories of trust usually draw a distinction between the *form* and *psychology* of trust.[96] Here, the form of trust refers to the type of relation that governs a particular trust interaction and corresponds to what or whom

[95] Faulkner, "Finding Trust in Government."
[96] This parallels Faulkner's distinction between the objects and psychological aspects of trust (see Faulkner, "The Attitude of Trust Is Basic."). However, as ultimately the difference lies not in *what* we trust but *how* we trust, I prefer to use the term 'form' to capture this.

is being trusted by an individual. In contrast, the psychology of trust refers to the attitude of the trustee in a particular trust relation, which in turn determines the normative content of trust. In what follows, I offer a detailed explanation of the terminology associated with the form and psychology of trust relations, before briefly showing how each of these components vary across some theories of trust.

Turning first to form, we can trust others in three different ways that correspond to the following three trust relations: three-place, two-place, and one-place trust. In three-place trust, we trust a particular individual to perform a particular action. We can formalise this as 'X trusts Y to ϕ', where ϕ denotes an action. Metaphysically, this is a hybrid notion – it combines an action with an attitude that one has towards the particular action.[97] So, the phrase 'I trust the postman to deliver my post on time' expresses this three-place form that trust takes. Yet we also speak of trusting others more generally. For example, one might say 'I trust my friend' without specifying a particular action for which the other party is trusted. It seems, for this to be true, one would need to trust their friend in *some* capacity (or for some ϕ), but, equally, there is no finite list of actions that one would have to trust their friend for in order to make this statement true.[98] We can thus formulate this two-place relation as 'X trusts Y', where the trust in question is a general attitude one can have towards particular others. Finally, it seems true for us to say of someone that they are trusting, or even perhaps trustworthy. In such cases, trust is merely an undirected attitude – it is simply a way in which certain individuals approach or treat unspecified others.[99] So, one-place trust – formulated as 'X is trusting' – just refers to an attitude where the individual in question has faith in some 'generalisable other'.[100]

[97] Ibid., 424–25.
[98] Ibid., 425.
[99] Ibid.
[100] Faulkner suggests that there is a mono-directional implication from one-place trust to two- and three-place trust respectively, which suggests that the core of our concept of trust is having a particular attitude towards others. See Faulkner, "The Attitude of Trust Is Basic"; Faulkner, "Finding Trust in Government.".

Moving to the psychology of trust, it seems true to say that there are distinct normative attitudes associated with trusting someone. Here, we can distinguish between a 'thin' and 'thick' attitude of trust, where the former relates to a prediction about another's behaviour while the latter refers to some kind of normative expectation that we have of others. Put differently, much like how we trust our computers to turn on or our cars to start, this thin conception refers to a predictive belief about the thing or person that we are trusting.[101] To illustrate, if the postman arrives at 8am every morning, I trust that tomorrow they will arrive at 8am based on their predictable pattern of behaviour. This is the dominant view in the social science literature, where theorists develop trust-based models to examine how individuals trust others given certain conditions. It seems however that we can also identify a thicker, normative conception of trust. After all, if someone close to us were to breach our trust, for instance by disclosing private information, we would feel a sense of *betrayal* – something that cannot be captured on this thin conception.[102] The reason for this is that a 'reactive attitude' like betrayal points to the existence of some normative content: when we hold others morally responsible, we express certain attitudes (e.g. betrayal, gratitude) that stem from their failure or success in performing the required action.[103] As a positive belief contains no normative content, the existence of a reactive attitude by itself points to a thicker conception of trust.

We are now in a position to see how different theories of trust offer a diverse characterisation of both the form and psychology of trust relations.[104] In her seminal paper on trust, Annette Baier argues that trust is exclusively a three-place relation.[105] The reason for this is due to the psychological component of her theory. Roughly, she argues that trusting someone requires making oneself vulnerable to that person, and

[101] Katherine Hawley, *Trust: A Very Short Introduction* (OUP Oxford, 2012).
[102] Baier, "Trust and Antitrust," 235.
[103] Strawson, "Freedom and Resentment"; Baier, "Trust and Antitrust."
[104] In what follows, I focus on the accounts offered by Annette Baier, Karen Jones, and Paul Faulkner as these three have been the most influential in the debate on public trust.
[105] It is worth noting that theorists who maintain this view often see two-place trust as merely a colloquial shorthand for three-place trust; that is, in saying X trusts Y, we often do just mean that X trusts Y in some particular way (see Faulkner, "The Attitude of Trust Is Basic," 424).

thus involves entrusting them with something that we value and expecting them to perform the action from a motive of goodwill.[106] It is this vulnerability – coupled with the reactive attitudes one would feel if someone were to break their trust (e.g. betrayal) – that distinguishes trust from mere reliance; these features distinguish a normatively thick understanding of trust from a normatively thin understanding of trust as reliance. In a similar vein, Karen Jones has argued for a normatively thick three-place conception of trust. Jones argues that trust involves "accepted vulnerability" over something we take to be valuable; that is, we grant someone else power over something that we deeply care about.[107] For her, 'accepted vulnerability' involves the following two conditions:

1) "The truster foregoes searching for ways to reduce such vulnerability
2) "The truster maintains normative expectations of the one-trusted that they not use that power to harm what is entrusted".[108]

Thus, should one of these conditions fail to hold, the relation in question would not be deemed one of trust. Further, for Jones, all trust relations are grounded by the presence of a (one-place) trusting attitude. For this reason, we often hear victims of trauma claim that their "trust has been shattered"; due to events beyond their control, they suddenly find it very difficult to trust others again.[109]

Both Jones and Baier stress the importance of trust as a *personal* attitude. For them, it is the emotive or motivational element that makes an otherwise simple reliance relation one of trust. This emphasis on motivation is not ubiquitous; for example, Katherine Hawley argues that trust consists in the honouring of commitments, which is less emotive and more *impersonal*.[110] We only trust, she argues, in cases where the trusted

[106] Baier, "Trust and Antitrust," 236–37.
[107] Jones, "Trust and Terror," 6.
[108] Ibid.
[109] Karen Jones, "Trusting Interpretations," in *Trust: Analytic and Applied Perspectives*, ed. Pekka Mäkelä and Cynthia Townley, Value Inquiry Book Series (Amsterdam: Rodopi Press, 2013), 22–23.
[110] Hawley, *Trust*, 6.

party commits to performing a particular action, and we rely on them to fulfil this commitment.[111] Like Baier and Jones, she distinguishes between trust and mere reliance, while also maintaining that trust is mostly (but not exclusively) a three-place relation. Yet, on Hawley's account, the motivation to make a commitment is irrelevant to whether or not a reliance relation is one of trust. One could honour their commitments out of self-interest, a caring disposition and so on – all that matters for trust is the presence and fulfilment of a commitment that one has freely entered into.[112] To illustrate, suppose that a pickpocket is relying on a busker to distract the crowd while he goes about his thievery. If the busker were to suddenly stop playing, the pickpocket could not claim that his trust was betrayed as no commitment was ever made between the two individuals. If however the busker is the pickpocket's accomplice, then any unexpected break in music being played would constitute a breach of trust.[113]

In contrast, Paul Faulkner does not distinguish between trust and mere reliance. Instead, he makes a tripartite distinction between *predictive*, *affective*, and *generalised* trust that captures the diversity present in both the form and psychology of trust. Here, predictive trust is normatively thin and is exclusively three-place, as it relates purely to the positive belief about the outcome of a trust relation (i.e. a positive belief that the trusted will ϕ).[114] Affective trust, on the other hand, is a normatively thick notion that, like Baier and Jones' view, is grounded in an emotive personal attitude. For Faulkner, the reason for the emergence of reactive attitudes in a given trust relation is twofold: first, we expect the trusted party to honour the action they were trusted to perform, and second, we expect that the trusted party is in some sense *moved* by the reliance that the trustee has on them to perform the action in question.[115] Again, affective trust is clearly three-place; we rely on others to perform particular actions and expect them to perform the relevant action *because* we rely on them to. However, we might also have a generalised presumption towards certain other individuals such that, if we were to trust them to ϕ,

[111] Hawley, "Trust, Distrust and Commitment," 10.
[112] Hawley, *Trust*, 6.
[113] I have taken this example from Hawley's book on trust (see Ibid., 4–5).
[114] Faulkner, "Finding Trust in Government," 630.
[115] Ibid., 630–31.

they would be moved by our reliance on them to ϕ. Thus, Faulkner's account can capture cases of two-place trust as follows: X affectively trusts Y iff, whenever X trusts Y to perform some ϕ, X expects that Y would be moved by X's reliance on Y to ϕ. Finally, we can understand *generalised trust* as an attitude of trust towards a general other such that, in any given interaction, there is a trustworthy thing to do and others can generally be trusted to perform this action.[116] This would then capture a one-place trust relation, as it refers to a trusting disposition in the absence of a particular subject.

What this points to, then, is the sheer diversity of conceptions of trust that exist in the literature today. However, as they stand, such views have little to say on the subject of public trust. After all, they have been developed to model the trust relations involved in *interpersonal interactions*. However, the literature on public trust has recently seen a buzz of interest; taking both Jones' and Faulkner's view on trust as starting points, theorists have offered multiple different accounts about the trust relations that characterise the public sphere. Specifically, these theories of public trust have extended the individual model to the public sphere, arguing that *every* trust relation can be captured using this model – regardless whether the trusted party is as a role-occupier, government official, or an institutional body. I take issue with this methodological decision. In short, I think it obfuscates the diversity of trust relations that are present beyond the individual case. In what follows, I show how theories of public trust have built on the individual model, before highlighting their inadequacies. I close by suggesting that developing a theory of public trust based on social norms seems promising.

The Methodological Issues with Theories of Public Trust

Recall Faulkner's tripartite distinction between predictive, affective, and generalised trust. Roughly, he differentiates between the normatively thin conception of predictive

[116] Ibid.

trust (what others call 'mere reliance') from the normatively thick conception of affective trust, where the trusted party is moved to act by the reliance of the truster. Finally, generalised trust just refers to a trusting disposition that others will be trustworthy. Applying his theory to vertical trust, Faulkner makes two key claims. First, he argues that individual citizens ought not to trust *particular* government officials in an affective manner. Thus, they should only predictively trust such officials. Second, he argues that citizens cannot affectively trust the government *in general*. Instead, citizens can have an attitude of generalised trust towards the government; that is, citizens can have the normative expectation that the government will do the right thing and make good decisions on our behalf.[117] In support of the first claim, cases of affective trust involve the expectation that the existence of the truster's reliance will *motivate* the trusted party to act accordingly. However, if we expect government officials to act on the basis of an individual citizen's reliance on them to perform a particular action, then it seems that we would see them as being corrupt.[118] This is because an ordinary citizen could give them a reason to act in such a way as to undermine the very institutions they represent.

Faulkner's point seems to be that, while affective trust in government officials is feasible, it is nevertheless wrong as it would imply that the official in question is corrupt. The reason for this is that government officials ought to remain impartial to the requests made upon them by their citizens; that is, they should be motivated to act on the basis of agent-neutral reasons rather than on agent-relative reasons. Although I take Faulkner's reasoning here to be informative, the conclusion is both bizarre and largely inaccurate. Recall that a key feature of affective trust is the presence of reactive attitudes. On this view, then, it would be *wrong* for ordinary citizens to feel betrayed by, or resentful towards, either their governments or particular government officials, since the mere presence of these attitudes would imply that the relevant authority was corrupt. Yet, we do often employ such language to describe our feelings and attitudes

[117] Ibid., 640.
[118] Faulkner, "The Attitude of Trust Is Basic," 638.

in the event of a government misstep, such as reneging on an electoral pledge. Consider for example the following extract taken from a letter to *The Independent* following a news article about Nick Clegg's knighthood:

"This is the man who *betrayed* millennials over tuition fees. The *oath-breaker* who *reneged* on his promise at the smallest whiff of power and co-presided over the trebling of tertiary education tuition fees".[119]

Had Clegg responded directly to the reliance of this particular citizen (as one of many relying on him to uphold his electoral pledge), Faulkner's account would have labelled him corrupt. This is because the motivating reason for Clegg to oppose the trebling of university fees, and thus uphold his electoral pledge, would have been the reliance of particular members of the electorate. Conversely, it would be morally wrong for the letter-writer to feel betrayed or resentful, as the presence of these reactive attitudes would point to a relation of affective trust between the letter-writer and Clegg. I find this conclusion highly implausible; it seems *prima facie* antithetical to any principle of democracy that, by responding to the reliance of the electorate, an elected government official is deemed corrupt in some manner. In fact, these reactive attitudes seem morally permissible, as they aid in holding goverments and particular officials to account if they breach the trust of the electorate. It could be objected that it is no longer corrupt if an elected official responded to the collective reliance of an electoral subgroup, but this suggestion seems ad hoc. For one, it would be difficult to determine at what exact point an electoral subgroup's reliance on a government official ceases to be impermissible. Given these considerations, at best the burden of proof shifts to Faulkner to provide an argument along these lines. At worst, his conclusion is flawed.

Turning to Faulkner's second claim, he argues that governments *in general* are not appropriate vehicles for affective trust since they cannot engage in trust-based

[119] See https://www.independent.co.uk/voices/letters/nick-clegg-knighthood-lorde-israel-politics-a8128976.html, my emphasis.

reasoning. Therefore, in the same way that it would be inappropriate to affectively trust a computer, it would be inappropriate to affectively trust governments. Instead, as we share the same moral universe as our governments, we can have the normative expectation that they will act in our best interests coupled with the optimistic view that governments tend to accurately reflect this.[120] At present, I just want to focus on the first part of this argument, as it seems unclear why governments *cannot* engage in trust-based reasoning. Presumably, this form of reasoning refers to the awareness that the trusted party is being relied upon to ϕ and should be motivated to ϕ by this reliance. Faulkner gives the example of an ex-convict (Paul) working at the till in a supermarket; here, despite the evidence that Paul may steal money, the store manager (Georgia) expects Paul to be motivated by her reliance on him to not steal from the till.[121] That is, Paul recognises that Georgia is relying on him not to steal and takes her reliance on him as a motivating reason to refrain from stealing. If this is trust-based reasoning, then it seems as if governments do engage in such reasoning – consider the Brexit referendum. The electorate relies on the government to uphold the result of the referendum, and expects the government to take this reliance as a motivating reason to deliver Brexit – even though there is substantial evidence to show that the UK would be worse off.

However, perhaps this is too quick. Strictly speaking, there is an important difference between the ex-convict case and that of Brexit – namely that, in the former, the store manager expects the ex-convict to be moved by *her* reliance. In contrast, the government is not responding to the reliance of individual Brexit voters to deliver Brexit, since the government is simply not aware of each individual person's reliance on them. Put differently, if trust-based reasoning requires a direct response to the trustee's reliance, then governments cannot be the appropriate vehicles for affective trust since they do not engage in trust-based reasoning so defined. While this seems to be the most natural reading of Faulkner's point, it does seem to be unnecessarily individual-focused. We can imagine that the British government in this case is responding to some kind of

[120] Faulkner, "Finding Trust in Government," 640.
[121] Ibid., 633.

collective reliance to deliver Brexit, based on the results of the EU Referendum. Generalising from this case, governments in democracies tend to favour winning the majority of votes in order to possess a democratic mandate to push their agenda. Thus, it seems that we can interpret a democratic mandate as some kind of collective reliance on the government to perform some action – be it to rule appropriately or push through a particular agenda.

If this is accepted, then the aforementioned objection melts away. This is because, although the government is not responding to any particular individual's reliance on them to perform some action ϕ, they are nonetheless responding to some kind of collective reliance purely in virtue of appealing to their mandate to govern in order to pass legislation. As such, governments would both recognise that they are collectively being relied upon to ϕ and would be motivated by this collective reliance to ϕ, where ϕ-ing denotes either governing or implementing a particular set of policy proposals. If this is accepted, then it seems as if governments *do* engage in trust-based reasoning so defined, which again questions the conclusion. Taken together, these issues point to unaddressed inadequacies in Faulkner's account of public trust. It is worth noting however that these issues stem from the methodological decision to ground an account of public trust in a model of individual trust. As noted previously, Faulkner's account of trust is particularly personal; it is based on individuals directly responding to the agent-relative reasons they have to act in certain ways. I think Faulkner is right to point out that his conception of trust struggles to account for affective trust in government officials, as they should in general remain impartial. However, it does not follow that every case of affective trust in government officials is morally wrong – as is evident in the case of Nick Clegg. As we will see, what is instead required is an account of trust that can adequately explain these *impersonal* relations that exist in the public sphere.

Having highlighted the problems associated with Faulkner's view, I now turn to the theory of public trust offered by Cynthia Townley and Jay Garfield. Their account explicitly extends Jones' understanding of trust to the public sphere. As such, they see

trust as a three-place relation grounded in a trusting attitude (i.e. one-place trust) that is important to foster in democratic societies. Put differently, they see trust as a three-place relation between a truster and a trusted party such that the former trusts the latter "within a specific domain or … with a specific good".[122] Further, they hold that trust can *only* be affective; anything that does not spawn a reactive attitude is a case of mere reliance. Another way of framing their view is that, if someone is unreliable, they can remedy the situation by offering some kind of compensation to the individual that was let down.[123] For instance, if I rely on my housemate to keep our kitchen clean before a dinner party and she fails to do so, she could compensate the inconvenience she has caused me by taking over some of my other chores around the house. However, if trust is breached, an entire interpersonal relationship risks being undermined, and therefore "it makes no sense to compensate".[124] Again, imagine I trust my housemate with some private information about one of the guests at my dinner party. Suppose that she divulges this to my guests while I am away and as such breaches my trust. Here, it seems nonsensical to repair the damage done to our friendship purely by promising to take on my chores.

Townley and Garfield's argument is grounded in their intuitions towards two particular cases. The first involves a trust relation that exists between a teacher and their students, while the second involves a trust relation between a university board and their philosophy department (as well as towards the general public). It is crucial to note that they see both as being instances of public trust; that is, they consider *both* interactions between individuals as role-occupiers (i.e. teacher-student) *and* between individuals and an institutional body (i.e. the university board) as instances of public trust. This move is not made frequently in the literature and is what drives my intuition about the inadequacies of applying the interpersonal model of trust to the public sphere. In short, by suggesting that such interactions belong to the domain of public trust, they notice a

[122] Cynthia Townley and Jay L. Garfield, "Public Trust," in *Trust: Analytic and Applied Perspectives*, ed. Pekka Mäkelä and Cynthia Townley, Value Inquiry Book Series (Amsterdam: Rodopi Press, 2013), 97.
[123] Ibid.
[124] Ibid., 96–97.

difference between these kinds of interactions when compared to the archetypal examples of interpersonal trust. Although I think their conception of public trust ultimately fails, they nonetheless make several insightful remarks that are worth elucidating when examining their view. As such, I will present both cases in Townley and Garfield's argument, before highlighting its deficiencies.

Turning to their first case, consider Arthur – a reliable but untrustworthy teacher. In his class, Arthur never misplaces a paper and always ensures his papers are marked and handed back on time, while also ensuring he is never late to any of his classes. However, Arthur is a serial plagiarist and essay broker; that is, Arthur will frequently steal ideas from his students to publish under his own name, while also periodically selling high-quality essays to students to use in other classes.[125] Here, Townley and Garfield argue, Arthur is untrustworthy because he fails to meet the normative expectations his students have of him as a teacher. In a sense, the students are entrusting him with their education – which is extremely valuable – yet he fails to take this seriously and threatens to undermine this by stealing their ideas and facilitating their cheating. Given this, it would be appropriate for the students to feel a sense of betrayal at having their ideas stolen by their teacher, and no amount of compensation by Arthur could ever repair their pre-existing relationship.

Turning to their second case, suppose that the governing body of Vanstone University, a public university, decided to cancel all philosophy degree programmes and instead use the skills of their philosophy department to produce the operating manuals for the equipment used by the Australian military.[126] This decision was undertaken in response to financial pressure the university was being placed under, and, as a result, all faculty members in the department have been relieved of their teaching and research obligations so that they can work on these manuals. Here, they argue that Vanstone University has undermined public trust by failing to act in accordance with the principles

[125] Ibid., 98–99.
[126] Ibid., 100–101.

they have committed to and been entrusted with.[127] Much like the case of Arthur, the University board have failed to meet the normative expectations that both the philosophy department in particular and the public in general have towards them. They were entrusted with the livelihood of a number of philosophers as well as the pursuit of valuable research, and exploited their position of power to damage both of these values. In short, by failing to teach and by failing to respect the academic integrity of its staff, Vanstone fails to act like a proper university and thus undermines public trust. Once again, the employees and public have the right to feel betrayed by the Vanstone board, and the offer of compensation will do little to alleviate the loss of trust.

At first glance, Townley and Garfield seem to offer a plausible account of how trust in the public sphere can be affective. Further, they differentiate between cases of a breach of trust and instances where a role-occupier or institutional body merely acted unreliably. However, the main issue is the lack of clarity that permeates their whole argument. For one, they never use the phrase 'normative expectation' in their paper – this is my interpretation based on their explicit endorsement of Jones' view of trust. Yet, I am not sure this argument is successful as it seems unclear to me whether the parties feeling betrayed in each case are the ones who are entrusting the trusted party with something valuable. Put differently, I am not sure the normative expectations in the relevant sense exist between either Arthur and the students *or* the Vanstone board and the general public. To see this, recall Jones first condition of accepted vulnerability, where the truster "foregoes searching for ways to reduce such vulnerability".[128]

I think it is strange to say that either the students or the public have 'foregone ways to reduce such vulnerability' when this vulnerability is accepted by fiat; that is, if they are placed in a situation where they are forced to accept their vulnerability, it is odd to claim that they have *chosen* not to pursue ways of reducing that vulnerability. To illustrate, compare the following case with the Arthur case from above. Suppose I am a prolific

[127] Ibid., 101.
[128] Jones, "Trust and Terror," 6.

writer who notes down my every thought – from the intimate to the banal – in a journal that I keep under lock and key. Suppose further that I am in a healthy relationship with my partner and that we have been happy together for a number of years. As a show of trust, I give my partner my journal to read – allowing her access to my most intimate thoughts about not only her but other individuals in our social circle. Suppose, for the sake of argument, that at least some of the information contained in the journal is not known to her. In this case, I willingly accept vulnerability over something I see as being extremely valuable and accordingly develop the normative expectation that she will not leak my thoughts to the people I talk about in the journal since I trust her. However, this seems significantly different to the case of Arthur, where the students are in a sense forced to entrust their education to him – since it is rare (if not unheard of) for students to pick and choose their teachers for any given subject.

My point is that the normative value of accepted vulnerability in the journal case is suitably different to the vulnerability in the Arthur case. That is, the act of entrusting my journal to my partner is a meaningful gesture; it is a show of trust towards my partner with the aim of furthering our relationship. As such, it seems clear that both conditions of Jones' theory are present: In entrusting my journal to my partner, I do not seek ways to reduce my vulnerability and I develop the relevant normative expectations over the fact that my journal and its contents are in the hands of my partner. However, it is less clear whether these same conditions hold in the Arthur case. The first thing to note is that the nature of vulnerability is different; rather than freely choosing to place themselves in a vulnerable position, as in the journal case, Arthur's students are forced to accept their vulnerability by fiat. Put differently, the students in his class do not entrust him with the value of their essays and ideas in the same way – they have to submit an essay in order to (at least) pass the relevant class or module. Thus, I take it to be unlikely that the students have foregone 'searching for ways to reduce their vulnerability' if their vulnerable position was one that they were forced to accept, rather than one they freely entered into. Regardless, I think this interpretation is questionable as either the presence of vulnerability by fiat fails the first condition of Jones's theory

or, if it can be shown that it does not, it creates a category error by treating the Arthur case and the journal case as instances of the same phenomenon. To be clear, I am not disputing that there is affective trust in the Arthur case. Rather, I am suggesting that Townley and Garfield's argument, when understood through the lens of Jones' theory of trust, fails to adequately explain why trust exists in the Arthur case.

The reason this interpretation fails is once again due to the methodological decision of applying an interpersonal model of trust to the public sphere. In short, they are supposing that the interactions between Arthur and his students are sufficiently similar to the interactions in personal relationships. A more fruitful approach, I argue, would be to adopt a more impersonal understanding of trust in order to capture the nature of the interactions in both the Arthur and Vanstone cases. Furthermore, an alternative interpretation along these lines is available to Townley and Garfield. When explaining why the Vanstone board undermines public trust, they make the following claim:

"[Vanstone University] has acted contrary to its entrusted purpose by abandoning its *commitment* to the cultural good of the nation in favour of commodifying the resources at its disposal".[129]

Here, it seems that they are drawing on a different understanding of trust. If 'abandoning their commitment' is the reason why Vanstone undermines public trust, then this points to Hawley's understanding of trust – where trust involves honouring the commitments one freely enters into.[130] I think this route is more promising; again, we have a compelling account to explain the presence of affective trust in the public sphere. On this interpretation, the Vanstone board undermined public trust by violating their commitment to academic integrity (towards its philosophy department) as well as their commitment to 'the cultural good of the nation' (towards the general public). Further, by appealing to Vanstone abandoning its commitments, we have an explanation for why

[129] Townley and Garfield, "Public Trust," 101, my emphasis.
[130] Hawley, "Trust, Distrust and Commitment."

both the philosophy department and the general public feel betrayed by the actions of the board. The appeal to commitments has an additional advantage: it does not take a feature inherent in personal relationships and seek to apply it to the public sphere. We make commitments all the time to both close friends and complete strangers; therefore, it seems – at least *prima facie* – that the commitment interpretation of Townley and Garfield's argument can circumvent some of the problems associated with the preceding interpretation.

I think this interpretation of the argument is roughly right. We have a plausible account that explains affective trust in the public sphere while also maintaining a suitably impersonal perspective. However, some pertinent questions remain – specifically related to the target of a given commitment. When we make commitments, we usually make them towards someone – for example, say my plumber commits to me that he will fix the plumbing before noon on Monday. Presumably, then, an archetypal case of trust will involve the trusted party making a commitment to the truster, and the truster relying on them to meet this commitment. Let us call these kinds of commitments a 'direct commitment'. But, is it necessary for trust to involve direct commitments? Hawley disagrees; she argues that there are frequent cases in which we trust individuals to honour their commitments *towards others*.[131] To use her example, suppose Monica, your daughter's friend, promises to Rachel, your daughter, that she will drive Rachel home from a party they are both attending. Placing your faith in Monica, you decide to have a few glasses of wine in the evening such that you are over the legal alcohol limit to drive.[132] Now, imagine that Monica reneges on her promise and does not drive Rachel back. Although Monica has not made a direct commitment to you, it seems clear that you trusted her to honour her commitment to Rachel, and were ultimately let down.

Let us call these kinds of commitments 'indirect commitments'. We can now ask whether you would be entitled to feel betrayed in this situation. That is, are reactive

[131] Ibid., 11.
[132] Ibid.

attitudes appropriate in cases where indirect commitments are not honoured? Intuitively, it seems true to say that it would not be appropriate for you to feel betrayed, though it would be appropriate for Rachel.[133] After all, Monica broke a promise to *her* and not to you. Similarly, it seems true to say that, while it would be suitable for Monica to compensate you for the inconvenience you suffered (perhaps by buying a box of chocolates), such offerings would not placate Rachel who may now no longer trust Monica to keep her promises. What we have, then, is an instance of trust where compensation is appropriate and reactive attitudes are *not* necessarily felt; in short, we have trust without an affective component. Of course, this view is inconsistent with Townley and Garfield's conception of trust as being necessarily affective. As such, they would presumably restrict trust to direct commitments, and view indirect commitments as a special case of reliance. However, doing so has grave implications for their conception of public trust as both cases seem to involve indirect commitments. Teachers by and large do not commit to upholding academic integrity *directly* to their students – their commitment is usually made towards the school as an institution. Even if this is disputed, in the Vanstone case, it seems strange to view the board as having a *direct* commitment to the public to further 'the cultural good of the nation'. At best, this is most likely a direct commitment to the government, who act as a proxy for the people – thus suggesting an *indirect* commitment to the wider public.

Taken together, this interpretation would no longer class Townley and Garfield's two central cases as instances of trust. Further, even if we try to build an account of public trust by including indirect commitments, we only encounter more problems. For one, we can no longer account for the feeling of betrayal that exists in both the Arthur and Vanstone cases. Hawley argues that this points to a dilemma in the trust literature; that is, cases involving indirect commitments imply that the link between trust and betrayal is not as strong as some theorists suggest.[134] Therefore, we are left with a difficult choice: Either we accept that the trust-betrayal connection is not as strong as previously

[133] Ibid., 14.
[134] Hawley, "Trust, Distrust and Commitment."

thought, and attribute the feelings of betrayal in cases of public trust to other factors. Or, we reject one of the intuitions in the daughter case; that is, we either have to reject that the daughter case is an instance of trust, or reject the notion that it is only appropriate for your daughter to feel a sense of betrayal.[135] I am not sure whether this dilemma is escapable. However, by exploring an avenue that has hitherto been largely ignored, we may be able to offer some considerations that makes the first horn easier to accept.

In Townley and Garfield's discussion about the wrongness of Arthur's actions, they raise the following point:

"Such acts destroy constitutive elements of a teaching/learning relationship, such as respect for intellectual production, academic integrity, and the like".[136]

Now, this point is framed in terms of the effect of betrayal on the students in Arthur's class. Nevertheless, I think it makes an insightful comment concerning the nature of trust and its betrayal. To see this, we must first assume that any given society is replete with *social roles*. This is not controversial; in any society, we occupy particular positions in a social structure – be it an occupational position, such as teacher or doctor, a familial position, such as mother or uncle, a social position, such as friend or partner, and so on.[137] Often, we occupy multiple social roles (i.e. a teacher can also be a mother), but there are instances where we interact from the perspective of just one of these social roles. For example, suppose I am friends with my doctor. If I were to go to a consultation, then it seems for the duration of this meeting she will act from the perspective of being my doctor and adhere to all the 'constitutive elements' of a doctor-patient relationship.

In the public sphere, we often interact with others from the perspective of the social roles we occupy, and these interactions are structured according to a series of social

[135] Ibid.
[136] Townley and Garfield, "Public Trust," 99.
[137] Michael O. Hardimon, "Role Obligations," *Journal of Philosophy* 91, no. 7 (1994): 334–35.

norms. These social norms I take as being the 'constitutive elements' of any given role-based relationship – such as the 'teaching/learning relationship' mentioned previously. Further, I claim that the *reason* we take Arthur to be untrustworthy is precisely due to the destruction of these constitutive elements; that is, the students view Arthur as untrustworthy and feel betrayed by his actions because he violated the social norms that structure a teacher-student relationship. If this is accepted, then this points to a norms-based conception of trust that may be able to solve the dilemma apparent in the trust literature. In the subsequent chapter, I spell out a theory of public trust that can play this role.

Chapter 3: Reconceptualising Public Trust

We have seen that prominent theories of public trust fail to adequately capture the nature of trust in the public sphere. The reason for this, I argued, was the methodological decision to take a model that captures the nature of trust in an individual context and apply this to the societal level. In short, many models of individual trust are highly *personal* in nature and therefore fail to adequately account for the diversity of trust relations in the public sphere. Even impersonal accounts of trust, such as Hawley's commitment-based view, have trouble when used to characterise given instances of public trust – though they do fare better than the personal accounts. In light of these difficulties, I argue that we need to reconceptualise our approach to developing an account of trust in general, and public trust in particular. The intuition that drives my dissatisfaction with the existing theories is that they are all in some sense *monist*. All theorists of trust think that we trust others in exactly the same way, and seek to develop an account that explains the nature of trust in an adequate manner. I think however we need to be more pluralist in our approach. To see this, we just need to compare the case of Paul, the ex-convict working at the till, with the case of Arthur, the plagiarising teacher.

Recall that, despite his past, Georgia trusts Paul not to steal from the till. Crucially, however, she has a normative expectation that the presence of her trust will influence his reasoning in some way. Put differently, Paul takes Georgia's trust on him to play some part in his decision to not steal. I see these kinds of trust relations as being particularly personal, since the truster's act of trusting influences the trusted party in some way. We see this most clearly if we were to try to analyse this *particular* trust relation through the lens of a more impersonal theory of trust. Hawley's view would say that Georgia trusts Paul as he committed to not steal in virtue of taking the job. But this interpretation misrepresents the nature of the trust relation; Georgia puts her faith in Paul by placing him in charge of the till and expects him to take this act of trust seriously by acknowledging its importance when deciding to not steal. As we will see, the social norm view I briefly introduced struggles for the same reason as Hawley's account does:

It is by its nature *impersonal* and therefore cannot account for instances where one party's reliance influences another's decision-making in a meaningful way.

The mistake theorists of trust make is in thinking that *every* trust relation can be couched in these highly personal terms. Thus, when relations of public trust cannot be expressed in this manner, they conclude that trust in the public sphere either *cannot* or *ought not* be normative. Yet this is erroneous, which is epitomised in the case of Arthur. Here, the students display the relevant reactive attitudes of betrayal, resentment, and so on. However, it seems to mischaracterise their reasons for feeling betrayed if we, following Faulkner, suggested that their reactive attitudes were caused by Arthur disregarding their reliance on him when deciding to act wrongly. Instead, it seems that we need to use a more impersonal notion of trust to better account for the nature of this kind of trust. In this chapter, I argue for my preferred impersonal account of trust grounded in social norms. I begin by arguing for a pluralist recharacterization of the nature of trust as having two forms: Personal, as in cases like that of Georgia and Paul, and impersonal, as in cases like Arthur. This distinction is under-theorised in the literature; in fact, it has only previously been offered by Phillip Petit. I suggest Pettit's presentation of this distinction is problematic as it is either exclusively non-normative, and thus cannot account for the normative content of trust. Or, it has the unsavoury implication that, to trust someone non-normatively, they need to be suitably constrained by regulation – which I take to be obviously false.

In light of these difficulties, I offer my own characterisation of the personal/impersonal trust dichotomy. My version of this distinction is purely normative; that is, it is a distinction that only applies within the normative domain of trust. Further, I take most cases of public trust as involving impersonal rather than personal trust.[138] Therefore, as the goal of my thesis is to establish that ASCs undermine *public* trust, I present my preferred account of impersonal trust. I do not aim to offer an account of personal trust

[138] The reason for this is that interactions between individuals as role-occupiers, or between individuals and institutions, tend to lack this dynamic where the trusted party is influenced by the truster's reliance on them. I do however leave it open as to whether such trust relations can be personal.

as it is beyond the scope of this thesis to do so. I argue, following Amy Mullin, that impersonal trust involves the assumption that both parties have a shared intrinsic commitment to a social norm deemed authoritative in a given domain.[139] As Mullin's definition of social norms is largely intuitive, I offer an interpretation of Mullin's argument that views a social norm as a type of 'practice rule' – where this is understood in the Rawlsian sense. Practice rules define a set of procedures that an individual must comply with in order to be seen as being a participant in that practice.[140] After outlining this interpretation, I close the chapter with a discussion of the upshots of this view.

What is Impersonal Trust?

We have seen that two prominent theories of public trust have failed to adequately capture the nature of trust relations in the public sphere – even when interpreted through the lens of different theories of trust. I argued that the reason for this was due – in part – to the ubiquitous application of theories designed to capture *personal* trust relations. That is, the application of accounts of trust that work best in cases where individuals have some kind of personal relationship only serve to obfuscate the nature of trust between individuals as *role-occupiers* and between individuals and institutions. Having established this negative thesis, I hope to call for a reframing of the literature on affective trust. In short, I argue that we need to distinguish between *impersonal* and *personal* trust within the affective domain. This distinction has been largely under-theorised within the trust literature – with the exception of Phillip Pettit, who sought to distinguish between two normatively thin conceptions of trust.[141] I briefly highlight the inadequacy of Pettit's account, before arguing for a reconceptualization of this distinction based on a shared commitment to social norms.[142] That is, I argue that impersonal trust occurs in cases where individuals are committed to a shared social

[139] Mullin, "Trust, Social Norms, and Motherhood."
[140] Rawls, "Two Concepts of Rules," 24.
[141] Philip Pettit, "Republican Theory and Political Trust," in *Trust and Governance*, ed. Valerie Braithwaite and Margaret Levi (Russell Sage Foundation, 1998).
[142] Mullin, "Trust, Social Norms, and Motherhood."

norm, while personal trust refers to cases where individuals have some kind of personal relationship. I take a key feature of personal relationships and interactions to be the ability to give others agent-relative reasons to act in some way, which is something that existing personal accounts of trust can capture well. Therefore, the aim of my account is to supplement the personal model in order to offer a more accurate characterisation of public trust. I close with a brief discussion concerning the upshot of my view.

At this juncture, it is worth exploring the following question: Is this distinction I draw between impersonal and personal trust ad hoc? That is, we can ask whether the distinction I am drawing is done purely to avoid the problems associated with the aforementioned theories of public trust or whether there are substantive reasons to consider building an impersonal theory of trust. In response, I begin with the observation that our interactions with our colleagues, our doctors, our teachers etc. are fundamentally different from the interactions we have with our close friends, partners, and relatives. We saw in the discussion concerning privacy that the amount and nature of the information we disclose to others is used to delineate our meaningful relationships from others. Furthermore, our interactions in the public sphere are structured in particular ways and are predicated on the roles we occupy in society. We can understand the social roles I have in mind as being *institutionally defined*. In a sense, the positions we occupy – be they political, occupational, familial, and so on – grant us particular duties, rights, and obligations that are explicitly tied to the social function these positions serve.[143] So, we have expectations that police officers risk their lives to combat crime in virtue of their occupational role while we do not have the same expectations towards individuals outside this social role.

The point is, then, that our social roles structure our interactions. Put differently, not all interactions are of a personal nature but can be either role-based (i.e. an interaction between people *qua* role-occupiers) or involve groups. Given that theories of trust seek to explain the psychology behind trust relations, it seems clear that the attitudes

[143] Hardimon, "Role Obligations," 334.

involved when we interact on the basis of our social roles will be fundamentally different from the attitudes we have towards our close relations. Therefore, the distinction between impersonal and personal trust is not ad hoc; rather, we do have substantive reasons to draw this distinction – namely that the underlying motivational attitudes we have when trusting others are not always the same. Phillip Pettit's conception of trust echoes this view. For him, impersonal trust refers to instances where we trust others on the basis that they are in constrained by the regulations of their job.[144] In such cases, we do not trust on the basis of some belief about the trusted party or their intentions; rather, we trust them because they are constrained such that they are bound to act in the way that we expect. To illustrate this kind of trust, suppose that I am about to take the Central Line on the London Underground. Suppose further that I hear an announcement stating that, due to a signal failure, all Central Line services are severely delayed. Here, I trust that this information is correct *without* forming an opinion regarding the announcer or their intentions, since I believe them to be constrained by their job to not give out false information. As such, it seems as if impersonal trust operates within the realm of predictive trust (i.e. is normatively thin).

In contrast, we might trust another individual due to a more positive disposition. Pettit suggests that, if the following two conditions hold, then we have grounds to trust another individual on a *personal* basis:

1. The trusted party is aware of my reliance on them
2. My reliance on the trusted party "triggers a cooperative disposition", which gives them reason to act in the way that I trust them to.[145]

Now, this characterisation of personal trust is ambiguous. On one interpretation, this 'cooperative disposition' might just be that the trusted party becomes aware of the value of cooperating in a given case.[146] If this is true, then this *also* points to a

[144] Pettit, "Republican Theory and Political Trust," 297.
[145] Ibid., 637.
[146] This interpretation is taken from Faulkner, "Finding Trust in Government," 638.

normatively thin understanding of trust, as it suggests that we engage in trust relations so long as it is in our interests to do so. On a second interpretation, a 'cooperative disposition' might be understood as a motivating reason to act in the relevant way. As such, this disposition is nearly identical to Faulkner's characterisation of affective trust, since it involves taking an individual's reliance as a *motivating reason* to ϕ.

However, under both interpretations, Pettit's distinction is problematic. On the first interpretation, *both* types of trust are predictive, which makes his account unable to tackle clear cases of affective trust. For instance, recall the case of Paul, the ex-convict working at the till; here Paul is trusted neither due to being sufficiently constrained so as to ensure that they refrain from stealing, nor are they trusted because it is in the interests of *both* parties to cooperate in this manner. Instead, the trust involved here points to a thicker conception whereby Paul is *directly responding* to Georgia's reliance on them to refrain from stealing from the till, rather than being the result of recognising the potential value in cooperating with the store manager. On the second interpretation, although Pettit can account for the case of the ex-convict, equating impersonal trust with predictive trust leads to the unsavoury implication that trusting someone based on predictable patterns of behaviour *necessarily* requires them being sufficiently constrained by regulations.

As both interpretations of Pettit's distinction appear to be problematic, it may be fruitful to reconceptualise our understanding of personal and impersonal trust. To reiterate, this distinction is important as there are important dynamics to trust relations that cannot be captured on a monist account. We saw this in our comparison between the cases of Paul and Arthur; in the former case, the act of reliance is deemed to be normative in some way. Put differently, when Georgia places Paul in charge of the till, she expects her reliance on him to play some role in his reasoning. In contrast, no such dynamic is in place in the Arthur case. Here, it seems more accurate to describe the students' feelings of betrayal as deriving from the wrongful actions that Arthur performed rather than the disregard he showed for their reliance on him. Therefore, it

seems more accurate to characterise the trust involved in the Arthur case in a more impersonal manner, where Arthur breached the trust of his students by breaking one of the constitutive elements of a teacher-student relationship. In light of this comparison, I hesitantly propose the following characterisation of the personal-impersonal distinction within the realm of (normative) trust. In personal trust, the source of normativity is internal to the individual; the attitudes and/or mental states of the truster are supposed to influence the trusted party in some way. In impersonal trust however, the source of normativity is external to the individual: it is the constitutive elements of the relationship between the truster and trusted that influences one's reason for acting.

I find this distinction at least *prima facie* plausible. Further, it seems to adequately capture the essence of the difference between the cases of Arthur and Paul. My hesitation in proposing this distinction lies in the fact that I have only established it negatively; my argument for this characterisation entirely rests on the inadequacies of monist accounts of trust, which my preceding analysis has hopefully shown. It is beyond the scope of this thesis to offer a positive defence of this particular characterisation – though I think this would be a fruitful project for further research.[147] Instead, in what follows, I provide my preferred account of impersonal trust. This is because impersonal trust can better capture the nature of trust in the public sphere. To gain a greater insight into the cases I have in mind when discussing impersonal cases, consider the following example. Suppose that Magnus is an avid chess player such that he frequently enters chess tournaments. In any given match, it seems plausible to say that Magnus trusts his opponent to play at his best (i.e. to not 'throw the game'), to play fair (i.e. playing by the rules) and so on.[148] Further, there is a sense in which Magnus would feel betrayed if it turned out that his opponent cheated or 'let him win'. How do we account for this sense of betrayal? On Faulkner's view, Magnus would have to rely on his opponent such that

[147] If the reader feels this move is unjustified, then it is worth noting that the remainder of my argument in no way rests upon this *particular* characterisation of the distinction between personal and impersonal trust. Instead, my argument relies on there being *a* distinction between personal and impersonal trust, which I have established through my negative thesis that purely monist accounts are inadequate to capture trust in the public sphere.

[148] I have taken this example from Mullin, "Trust, Social Norms, and Motherhood," 318.

his opponent took Magnus' reliance as a motivating reason to play fair. Intuitively however, it seems implausible for Magnus' opponent to respond to his reliance in this manner. After all, they are opponents and it is therefore unlikely that they would give any consideration to each other's expectations, let alone be moved by them. Instead, it is much more likely that both players are committed to a constitutive element of competitive chess.[149]

Amy Mullin offers an account of trust on precisely these grounds. She argues that, in every case of trust, the truster trusts the trusted party insofar as the former assumes that the latter *shares* the former's commitment to a social norm *for its own sake* rather than for instrumental reasons.[150] Mullin's point of departure from the trust literature is that insufficient attention has been given to the role of social convention.[151] The traditional view is that social convention plays some role in ensuring a 'climate of trust' by regulating general behaviour through either sanctions or occupying clearly defined social roles (e.g. having a job).[152] Thus, the traditional view sees social convention as having an indirect role in sustaining or developing trust relations in particular contexts. Contra this, Mullin argues that social conventions are central in determining the content and limitations of trust. The reason for this is that, in examining and dismissing a number of motivations to enter a trust relation, she found that – common to all – was the expectation that the trusted party would adhere to a particular social norm.[153] If this is true, then social norms directly affect the content and limits of a trust relation. This is because they determine both whether we can trust someone and whether it would be appropriate to trust someone in a given context.

[149] This characterization lends itself more favourably to Townley and Garfield's conception of public trust. However, in what follows, I hope to offer a more precise account concerning the nature of impersonal trust.
[150] Mullin, "Trust, Social Norms, and Motherhood," 316.
[151] Ibid., 317.
[152] Russell Hardin, *Trust and Trustworthiness*, The Russell Sage Foundation Series on Trust, v. 4 (New York: Russell Sage Foundation, 2002); Baier, "Trust and Antitrust."
[153] Mullin, "Trust, Social Norms, and Motherhood," 320.

From this, Mullin suggests that trust is the *assumption* of a shared *intrinsic*, rather than instrumental, commitment to a social norm that is deemed to be authoritative in some domain.[154] In drawing this distinction between an intrinsic and instrumental commitment, she can account for the difference between affective and predictive trust respectively. That is, in cases where a commitment to a social norm is assumed to be intrinsically shared by both the truster and trusted party – for example a norm of good friendship – the truster has some normative expectation that the trusted will act in the appropriate way. However, if a commitment is assumed to be instrumentally shared, then this points to a model of predictive trust as the trusted will only act in the appropriate manner so long as it remains in their interests to do so. It is worth noting that Mullin does not make this distinction between predictive and affective trust; rather, she suggests that trust requires a normatively thick conception associated with an intrinsic commitment to a social norm while mere reliance can be understood in thinner normative terms. I think her terminology is somewhat misguided as it would preclude the ability to *rely* on inanimate objects like cars and computers, since these cannot in any sense be committed to a social norm. As such, I reframe her distinction as one between predictive and affective trust, and leave it open as to whether trust in inanimate objects is a third and more instrumental type of trust.

Similarly, it seems as if Mullin's account is sensitive to the distinction between three-, two-, and one-place trust. Her account of a social norm is broad and encompasses rules that regulate behaviour in particular ways. As such, social norms on Mullin's view can be role-dependent (e.g. rules governing the behaviour of a friend, husband, teacher etc.), based on following the 'rules of a game' (e.g. refraining from making an illegal move in chess), idealisations of certain roles (e.g. norms governing what it means to be a 'good' employee, boss, friend etc.), as well as principles that govern particular interactions or contexts (e.g. standing on the right on the London Tube).[155] Thus, on

[154] Ibid., 322.
[155] Ibid., 320. This list is not intended to be exhaustive; rather, in the absence of a precise definition of social norms, the point of these examples is to give an indication as to the types of rules or principles that may qualify as social norms on Mullin's account.

Mullin's view, the form a trust relation takes is entirely dependent on the nature of the shared norm. To illustrate, suppose that A promises B that she will ϕ. Here, it would only be true to say that B trusts A to uphold her promise if both parties assume that the other shares their intrinsic commitment to the social norm of keeping promises. Notice that the norm in this case lends itself to a three-place trust relation. In contrast, suppose that C and D are best friends. As such, C and D behave in accordance with a number of different norms that they both assume the other is intrinsically committed to as friends; for instance, C and D regularly confide in one another whenever they need advice, or may organise spontaneous trips to the pub at short notice. Now supposing that they are both intrinsically committed to a number of these social norms of friendship, it seems that we can reformulate this as a two-place trust relation such that C trusts D across a variety of domains. Finally, given the ubiquity of social norms in society, it seems that an individual can be described as trusting so long as they have a predisposition to assume that others will act share their own intrinsic commitment to the social norm that is authoritative in the relevant context. Such cases would thus correspond to one-place trust.

There are two further advantages of Mullin's account. Firstly, it allows room for trust relations to be *interpretive*. In certain cases, an individual may feel as if they trust someone to perform a particular action – even if the trusted party in such cases does not have a commitment to the same social norm. To illustrate, suppose that Todd is a patient of Dr. Young.[156] Now, it seems conceivable that Todd trusts Dr. Young to tell him the truth about his diagnosis – even if it would cause him pain or discomfort. Yet, suppose that Dr. Young was also Todd's sister-in-law. In such cases – as Dr. Young occupies multiple social roles – Todd may perceive their relationship as being closer than merely that of a doctor and patient. Perhaps he thinks that, as family, it would be permissible for him to call her outside working hours and expect her to give him medical advice. Similarly, we can imagine that Todd may expect preferential treatment (e.g. jumping a queue for a non-urgent operation) from Dr. Young given their familial

[156] I have slightly amended this example from Mullin's paper – see Ibid., 319.

connection. Let us suppose Todd interprets his relationship with Dr. Young in this way. However, being a good doctor, Dr. Young does *not* have an intrinsic commitment to these same norms; for her, she is intrinsically committed to the social norms that govern being a good doctor – namely helping those in need and treating everyone fairly. As such, if she were to not live up to these expectations, Todd would feel betrayed and interpret Dr. Young as betraying his trust – even if this trust is not present from an objective perspective.

Second, Mullin's norms-based theory of trust is *morally neutral*. That is, she does not see trust (and by extension trustworthiness) as being something morally virtuous that ought to be praised. Instead, she suggests that whether trust is morally right or wrong depends on the social norm that is being adhered to.[157] That is, trusting another individual would be immoral only if the shared social norm deemed authoritative is itself immoral. This point is crucial as it allows Mullin to condemn certain high-trust environments (e.g. the Mafia) as being unequivocally wrong – something that other accounts of trust may struggle to do. Although I elaborate on this point later on, viewing trustworthiness as dependent on the social norms undergirding a given trust relation allows Mullin to point to instances when distrust is permissible.[158] Put differently, if the social norms that structure a given trust relation are themselves immoral, then trust in such cases would be immoral whereas distrust would be moral. To make her point, Mullin uses a case of a nanny caring for the child of self-interested parents. She argues that, if the parents instruct the nanny to take their child on playdates with a known bully – whose parents just so happen to be in a position that can benefit the careers of the first child's parents – it would be moral for the nanny to refuse to subject the child in her care to the machinations of the bully even if this would betray the trust of the child's parents (by violating the social norm that, say, babysitters adhere to the wishes of their employers).[159]

---

[157] Ibid., 322–23.
[158] For an excellent contribution on the importance of distrust in democracy, see Meena Krishnamurthy, "(White) Tyranny and the Democratic Value of Distrust," *The Monist* 98, no. 4 (October 2015): 391–406.
[159] Mullin, "Trust, Social Norms, and Motherhood," 322.

My account of impersonal trust is largely inspired by Mullin's. In cases where we interact with one another as role-occupiers, our interactions are governed by a number of social norms that we, as occupiers of those social roles, are necessarily committed to. Furthermore, as these social norms play a large role in determining behaviour, it seems plausible to claim that we trust others *qua* role-occupiers insofar as they share our intrinsic commitment to the relevant social norms that are deemed authoritative as part of their social role. To illustrate, a student might trust their teacher to be doing their job well insofar as they assume the other is also intrinsically committed to norms that govern quality in teaching, such as displaying a commitment to knowledge/learning etc. The same is true when it comes to interactions between citizens and their governments; there are a number of social norms that structure these interactions, and governments are only trustworthy insofar as they consistently display a commitment to such norms. These norms might be related to democratic procedures, such as voting, or might be something greater and refer to the principles a government is committed to upholding. In any case, governments and/or particular government officials would undermine the trust of the public if they were to suddenly deviate from the norms that structure these interactions. This is because the public would no longer be able to assume that these government officials would share their intrinsic commitment to the relevant social norms.

The above characterisation, while I hope persuasive, is nevertheless somewhat vague. Part of this problem stems from Mullin's use of a broad conception of social norms, where anything from specific role-based norms to general evaluative norms that determine whether or not one is performing their given role well is included in her definition. What complicates this further is that, at certain points in her paper, it seems as if Mullin equates 'social norms' with 'social practice', which is questionable and risks being a category error. At other points, she distinguishes between the two, yet it is unclear what she means by either. In what follows, I offer an amenable interpretation of her conception of social norms. The reason I do so is for reasons of clarity; Mullin's

definition alone does not tell us how these social norms emerged or how individuals in a given society become aware of them. She merely stipulates that social norms exist, and lists a number of rules that can be seen as social norms without explaining *why* these rules are social norms. To make this step, I offer an interpretation of Mullin's understanding of social norms based on what John Rawls has elsewhere called a 'practice rule'. I first offer a definition of practice rules, before illustrating why social norms can plausibly be construed as a type of practice rule. I defend the plausibility of this interpretation by, first, showing that they share a number of intrinsic features and, second, by showing that this interpretation is beneficial for the clarity of Mullin's theory.

Rawls distinguishes between two different types of rules: Ones that justify a social practice, and ones that justify particular actions *within* a given social practice.[160] In short, we need to distinguish between a *summary rule* and a *practice rule*. Here, a summary rule is the product of all the individual decisions people would take when applying the same general principle (e.g. the principle of utility) to a particular kind of case.[161] As such, the actions of individuals are prior to the formulation of the rule. In contrast, a practice rule outlines a set of procedures that one must comply with in order to be seen as performing a new kind of action.[162] To illustrate, consider the game of basketball. Although the actions involved in basketball – such as throwing a ball, running etc. – can be performed in any context, one could only be described as 'scoring a basket' or 'performing a layup' if they could also be described as playing basketball. Put simply, practice rules outline a set of actions that, when individuals perform such actions, they can be described as participants in the relevant practice. Therefore, contra summary rules, practice rules are logically prior to the *particular* actions of individuals. My goal is to show that social norms can be conceived as a set of practice rules.

Social norms and practice rules share a number of intrinsic features; as such, it is at least plausible that social norms can be interpreted as a kind of practice rule. First, both

[160] Rawls, "Two Concepts of Rules."
[161] Ibid., 19.
[162] Ibid., 25.

practice rules and social norms are *internally justifiable*; that is, both are justified by appealing to the norm or practice itself. Some theorists argue that it is precisely their internal justifiability that distinguish social norms from moral norms.[163] In his paper, Rawls makes a similar point concerning practice rules. He argues that practice rules are justified solely by appealing to the standards of the practice itself (i.e. by appealing to practice-dependent standards). The reason for this is that any practice is constituted by the rules that govern it and therefore can only be justified by invoking these rules. For example, consider the act of promising. Suppose that – in response to being asked why he broke a promise – Boris referred to the principle of utility and argued that he had to break his promise 'for the greater good' (i.e. to maximise utility). It seems that, in this case, Boris fails to understand the *very practice of promising*; if promises could be broken on the basis of any general principle, the practice would simply cease to exist (or would be eroded beyond recognition).[164] For Rawls, then, promising as a practice only exists because of a set of rules that *prevent* the application of general principles to its particular case. Put differently, the purpose of promising is to strip oneself of the right to act in accordance with general principles when performing the action of promising. As such, "the only justification of promise-keeping that he can offer, Rawls maintains, is one that appeals to the constitutive rules of promising".[165]

Second, both social norms and practice rules are *socially constructed*. We use social norms to in some sense ensure compliance with a particular pattern of behaviour. As such, they are necessarily constructed by our behaviour. The same is again true of practice rules – recall Rawls' claim that practice rules are logically prior to the particular actions of individuals. At this juncture, it is worth clarifying the ambiguity of this statement. For Rawls, this just means that the actions that define a particular practice can only emerge once that practice has been codified by practice rules. Take the example of basketball once again. Presumably, in the development of the game,

[163] Geoffrey Brennan et al., *Explaining Norms* (Oxford University Press, 2013), 69–70.
[164] Rawls, "Two Concepts of Rules," 16–18.
[165] Tamar Schapiro, "Compliance, Complicity, and the Nature of Nonideal Conditions," *Journal of Philosophy* 100, no. 7 (2003): 335.

individuals performed actions that *looked like* they were playing what we would now call basketball (e.g. throwing a ball into a basket, dribbling with the ball etc.). However, these actions only became known as basketball once the practice rules had been established. It is in this sense that practice rules are logically prior to the particular actions of individuals – until a practice has been codified by a set of rules, it cannot be said that individuals are participating in such a practice (even if the actions they perform appear similar). Third, social norms are (or purport to be) *normative* in some sense. We often think that social norms give us reasons to act in specific ways; that is, the existence of norms give us reasons to behave in ways that we may not otherwise, such as by wearing black at funerals.[166] Again, it seems the same is true of practice rules; they are normative in the sense that they define procedures compliance with which determines the type of action one performs. As such, if one takes themselves to be participating in a practice – say attending a funeral (or being respectful at a funeral) – they would only do so if they are in compliance with the rules governing that practice.

It seems, then, that we have reason to treat social norms as practice rules. However, before considering the upshot of this view, two important qualifications are in order. First – at least *prima facie* – it seems as if practice rules struggle to account for norms of convention (e.g. driving on the left in the UK). Put differently, a number of collective action problems are solved by either constructing – or appealing to – social norms. But these do not seem to be the same as practices. I argue this is misguided as we often set up institutions and laws to protect and reinforce these conventions. Consider the driving example; in the UK, all the cars are produced with the steering wheel on the right, driving lessons are regulated to ensure all new drivers are told to drive on the left, and so on. Therefore, it is conceivable to view all of these regulations as the development of a practice of driving in the UK, of which a constitutive rule is driving on the left.[167] Second,

[166] Brennan et al., *Explaining Norms*, 6.

[167] Although it is not a necessary condition of practices that they be unique, it seems more plausible to view driving in the UK as a practice when we consider that particular rules about driving in the UK (e.g. particular speed limits) are not necessarily the same in other countries. Motorcycle drivers, for example, have different gestures to mean 'thank you' in different places – as shown on this forum: https://advrider.com/f/threads/sticking-a-foot-out-to-say-thanks.630059/

understanding social norms as practice rules would treat legal norms and social norms as the same kind of phenomenon. In their analysis, Geoffrey Brennan et al. view this as problematic as legal norms are often followed out of fear of being sanctioned by the relevant authority. This argument is overly simplistic; as we will see in the subsequent chapter, a number of legal norms are followed out of a commitment to the norm *itself* rather than out of fear of being sanctioned, and thus can be conceived as a type of practice rule.[168] Yet, there may nonetheless be a reason for this distinction on the grounds that legal norms and social norms may be normative in different ways.

Now, if one is committed to legal positivism – where the law is determined exclusively by social facts (i.e. social conventions)[169] – then this is not an issue. This is because, under a positivist picture, the normativity of the law is *only* internally valid, much like practice rules. However, under a Dworkinian picture, the law is in some sense determined not only by social facts, but *moral facts* as well.[170] As such, the law derives its normativity from the moral principles that inform their content, which points to a different source of normativity than those found in practice rules.[171] Although I am sympathetic to legal positivism, this is nevertheless troubling as a theory of public trust should not be committed to particular conceptions of what the law is. Notice though that, for Dworkin, the law is *non-exclusively* determined by social facts (rather than exclusively determined by moral facts). As such, there may be a number of domains within the law that can fit a conception of practice rules, as they will be internally justified. Further, we may be able to apply Rawls' distinction between summary and practice rules to accommodate individuals sympathetic to Dworkin's conception of the

[168] Norms of taxation are one such example, although I will elaborate on this point in the subsequent chapter. For more detail, see James Alm, "Measuring, Explaining, and Controlling Tax Evasion: Lessons from Theory, Experiments, and Field Studies," *International Tax and Public Finance* 19, no. 1 (February 2012): 54–77.

[169] Here I use 'social conventions' instead of 'social practice' to avoid confusion with my use of the term 'practice' and 'practice rule'. For a compelling account of legal positivism, see Scott J. Shapiro, *Legality* (Cambridge, Massachusetts: Harvard University Press, 2013).

[170] Scott J. Shapiro, "The 'Hart-Dworkin' Debate : A Short Guide for the Perplexed," in *Ronald Dworkin*, ed. Arthur Ripstein (Cambridge University Press, 2007), 22–49.

[171] For the initial presentation of his position, see Ronald M. Dworkin, "The Model of Rules," *The University of Chicago Law Review* 35, no. 1 (1967): 14–46.

law. That is, we might say that Dworkin is committed to the view that the legal system *as a whole* should be designed according to some general moral principle(s), while particular laws and legal practices may be justified by appealing to standards internal to that system. It is beyond the scope of this thesis to outline what such a summary rule could look like, but what has been said is sufficient to accommodate the Dworkinian view.

Understanding social norms as practice rules provides us with numerous benefits when considering the notion of public trust. For one, we can non-arbitrarily explain the existence of social roles. Recall that under my conception, social roles are institutionally defined. This I think is naturally read as being a particular instantiation of a practice rule; just as practices define a set of procedures, such as scoring a layup in basketball, they would also define a set of roles that must be occupied – such as the positions on a basketball team (e.g. point guard). Further, each social role is then associated with constitutive rules that define what an occupant of that social role ought to do or how they ought to behave. In the same way a point guard is tasked with dribbling the ball up the court in basketball, firefighters are tasked with saving individuals from burning buildings or police offers are tasked with arresting lawbreakers. The constitutive rules governing each practice allow us to clearly identify the social norms that particular individuals are committed to in virtue of their social position. In this way, our expectations that others are aware of the norms they are in some sense committed to becomes less mysterious (i.e. the ways in which norms can be shared becomes easily explained).

It seems clear that Mullin's theory can be recast in these terms without losing any explanatory value. When two individuals assume that the other is also intrinsically committed to a given practice rule, then we can say they trust one another. In addition, framing social norms in this way gives us a number of insights into the ways trust can collapse. We have already seen a case where one individual interprets the practice rules governing the other's role in a way that is incompatible with the way these rules are

actually instantiated (i.e. the Todd case). On this conception, we can outline further cases where the fault lies with the *trusted* party rather than the truster. Suppose for example that Reem is about to participate in the practice of a business negotiation with Graham, a third party. It seems plausible that, amongst others, a business negotiation would be governed by practice rules including a principle of 'good faith', however defined.[172] Now, imagine that Graham enters these negotiations purely with the intention of stalling Reem (perhaps because this may give him or someone else a better deal, or even because he just wants to), in direct violation of this rule of good faith. It seems we can say two things: First, by not displaying (and adhering to) an intrinsic commitment to the practice rule of good faith, Graham violates Reem's trust – which would explain why she may feel betrayed and resentful as a result of these negotiations. Second, it seems we could add that, in addition to breaching trust, Graham fails to negotiate with Reem – instead he is merely "babbling", since he fails to comply with the procedures that constitute negotiation.[173]

This insight is valuable as it can explain *why* we feel betrayed. We trust that others will act in accordance with the constitutive rules of a given practice and, when this fails to occur, we feel betrayed for having thought we were participants in such a practice. Further, I think this also points to the importance of having intrinsic commitments in trust relations. To see this, consider the following situation:

> Suppose a police force is filled with corrupt cops such that any individual attempting to adhere to the practice rules governing policing merely serves to distort the ideal of justice police officers (otherwise) try (or are expected) to uphold. Suppose further that this behaviour is encouraged by the Chief of Police, Captain Dudley, who has been using his influence to gain control of a drug trade formerly controlled by the Mafia. Dudley has abused his position to secure his

[172] I have taken this example from Schapiro, "Compliance, Complicity, and the Nature of Nonideal Conditions," 337.
[173] Ibid.

> power by framing innocent civilians and torturing or even murdering his fellow police officers.[174]

Here, the police are completely violating most of the practice rules that govern their practice – only procedurally adhering to the rules that would see them be identified as police officers (e.g. wearing a uniform). In such cases, where the police become the crime syndicate, it seems that the general population suffers a complete loss of trust in the police force and their ability to fulfil their purpose. The reason for this is that they can no longer be perceived as having an intrinsic commitment to the social norms that determine what it means to be a police officer. Put differently, it is no longer the case that individuals assume that police officers are intrinsically committed to the practice rules of policing, which in turn points to the complete lack of trust. In the absence of such intrinsic commitments, societies of distrust emerge, bringing with it significant negative values such as crime, corruption and so on. The strength of my account is such that it can explain these phenomena and can suggest reasons why the innocent individuals who were framed (as well as the general population) feel betrayed by Captain Dudley simply by pointing to the violation of the practice rules that govern policing.

In summary, I have defended a norms-based account of public trust, where social norms are understood as rules that constitute a particular kind of social practice. I did so by first showing how mainstream theories of public trust fail to adequately capture the diversity of trust (and distrust) in the public sphere, before outlining how my preferred account can succeed in the areas where the other theories fail. However, my norms-based view is strictly a theory of impersonal trust, which includes broad swathes of interactions in the public sphere. As such, it needs to be supplemented by one's preferred conception of personal trust – be it a motive-based account as advocated by Baier and Jones, a second-personal account like Faulkner's, or even a suitably-amended

[174] This is (slightly) adapted from the plot of the film *L.A. Confidential* and is the main case for Tamar Schapiro's argument that – in instances where rule violation is widespread – pursuing actions that are otherwise morally wrong may be permissible. For more detail, see Ibid., 347–50.

commitment based account similar to the one put forward by Hawley. Although I have only hinted at this point throughout this chapter, it seems as if my theory is compatible with each of these views on trust, which I take to be a strength of the account. In what follows, I hope to apply this conception of trust to the case of ASCs, which motivated this project in the first place. If it can be shown that ASCs undermine the shared intrinsic commitment to practice rules that govern useful social practices, then this would be a compelling argument for their regulation.

Before investigating this avenue further, I will close this chapter with a brief discussion of the case that motivated our exploration of social norms. Recall the case of Rachel, your daughter, and Monica, your daughter's friend who promised to drive her home and who ultimately reneges on this promise. Here, you trusted Monica to fulfil her promise to take Rachel home, and so decided to drink too much wine such that you were no longer able to drive. Again, it seems that we have an instance of a trust relation between you and Monica, yet one where you are not entitled to feel betrayed by her act of reneging on her promise. Applying the norms-based account just defended to this case gives us an interesting implication. The practice of promising defines two distinct roles – a promiser and a promisee – such that the promiser has an obligation to the promisee to fulfil the promised action, unless they have a valid excuse.[175] Presumably Monica in this case fails to have such an excuse. Now, this means that Monica and Rachel – as promiser and promisee – have a shared commitment to this practice rule of promising.

But what of you as a parent? As Monica did not promise *you* to take Rachel home, there is no shared commitment between you and Monica. Yet, I would contend that the trust relation can be expressed in the form of one-place trust (i.e. a trusting attitude). That is, you may have a predisposition to think that others will do the trustworthy thing by displaying an intrinsic commitment to the social norm in a given context, and thus trust Monica in *this* sense – namely you trust her to fulfil her promise to Rachel. When Monica reneges on her promise, while you may not feel betrayed, you may nevertheless see

[175] Rawls, "Two Concepts of Rules," 14.

that the practice of promising has been diminished or infringed upon in some sense. Perhaps you might think that ‘promising no longer means what it used to’, which may affect the ways in which you trust others to keep their promises. Put differently, Therefore, although the relation between trust and betrayal may be weakened, there is a sense in which this weakening may have a substantive (i.e. non-ad-hoc) explanation predicated on the different forms trust relations may take.

Chapter 4: Impersonal Trust and Anonymous Shell Companies

We are now in a position to see how anonymous shell companies (ASCs) undermine public trust. As noted previously, ASCs are a blight on the international system. They have been used in almost every form of economic crime, as they effectively guarantee anonymity to their ultimate beneficial owners. Consequently, they are a favoured device for criminals and tax evaders – amongst many others – to hide their wealth from the reach of the state. I argued in Chapter 1 that, in spite of universal agreement concerning these negative externalities, ASCs have remained unregulated as policymakers are caught in an intractable debate between promoting either transparency or privacy. In light of this impasse, I suggested we turn towards the value of public trust in order to provide a *pro tanto* argument for the implementation of a suitably-amended public ownership register. In what follows, I show that ASCs undermine public trust through two different mechanisms.

My argument rests on impersonal trust, because it best captures public trust. Recall that I understand impersonal trust as the assumption that there is a shared intrinsic commitment between individuals to a social norm deemed authoritative within a given domain, where social norms are understood as types of practice rules. Thus, when A impersonally trusts B, A assumes that B shares A's intrinsic commitment to a particular practice rule that governs the action A is trusting B to perform. When we say that public trust has been (or is being) eroded, we are referring to this assumption no longer holding between individuals in a given society. So, in cases of public trust, this assumption operates at the *societal level*; that is, in cases where public trust is high, any truster can assume that the trusted party shares the same intrinsic commitment that they do purely in virtue of being in the same society.[176] This clarification is important as it avoids the otherwise problematic implication that, in the event that one individual's intrinsic commitment to some norm is undermined, public trust is also undermined. To see this, recall the case of Monica – your daughter's promise-breaking friend. I argued that the

[176] Note, my use of 'truster' and 'trusted' here refer to any agent capable of entering into a trust relation and therefore includes individuals, groups, institutions etc.

best way to interpret this case was that your intrinsic commitment to the practice rule of promising was perhaps shaken, but not eroded. While you may now participate in the practice of promising with greater apprehension, it does not follow from this one instance that public trust in the practice of promising has been undermined.

However, if more and more individuals begin to break promises with reckless abandon, then there will be a point when individuals in this particular society will cease to trust others when they promise to do something. At that point, we can say that public trust in the practice of promising has been eroded because it will no longer be the case that any two individuals can assume that the other shares the same intrinsic commitment to the practice rule of promising. I argue that the widespread use of ASCs within a given society erodes public trust in much the same way. Take the case of taxation. As we will see, there is significant empirical evidence to suggest that social norms play a crucial role in fostering tax compliance. If, however, individuals within one's close circle begin to evade tax, then overall tax compliance rates begin to decrease. The best way to interpret this data, I hope to show, is by applying my conception of public trust. Thus, the first half of this chapter is devoted to exploring this argument. After examining the empirical evidence showing that tax compliance and social norms are interlinked, I recast this literature in terms of practice rules without losing its explanatory power. I then demonstrate how viewing tax compliance through the lens of practice rules gives us reason to think that we have an intrinsic commitment to the rules governing taxation, which suggests trust plays a role in ensuring tax compliance. Therefore, by undermining the intrinsic commitments individuals have to norms of taxation, the increased use of ASCs by individuals will – over time – undermine the society-wide assumption that others share an intrinsic commitment to the practice rules of taxation.

This is however not the only way in which ASCs undermine public trust. The use of ASCs can also facilitate the subversion of practices such that they no longer fit the purpose for which they were created. In such cases, I argue, public trust is eroded as the practice rules themselves have, in a sense, been made redundant. The key example here is the

case of Captain Dudley we considered in the previous chapter. If the police, through Captain Dudley's influence, become a crime syndicate, then public trust in the police force will vanish since individuals can no longer assume that any police officer will share their intrinsic commitment to practice rules of policing. In the second half of this chapter, I argue that the continued use of ASCs risks subverting the practice of democracy in a similar manner. Specifically, by acting as a vessel to funnel anonymous wealth into the American political system, ASCs undermine one of the central practice rules that govern any democracy: that the legislature is beholden *only* to the people. In making this case, I consider the context of American politics in the aftermath of the landmark Supreme Court ruling in *Citizens United v. FEC* in 2010, which declared all limits on campaign spending to be unconstitutional.[177] I begin by examining the debate about corruption in the USA from the Constitutional Convention in 1787 to *Citizens United* to contextualise my argument, before showing how ASCs undermine the legislative process in an unaccountable manner – thus eroding public trust.

Before starting, two clarifications need to be made. First, though these are not the only areas in which ASCs operate, the mechanisms by which ASCs undermine trust operate either at the individual level, where the aggregation of individual trust-undermining actions ultimately undermines public trust, or at the societal level, where a given practice is subverted entirely. Thus, while I focus only on two areas that ASCs operate in, I identify two key mechanisms that can be generalised and thus I hope to show that – even in areas not considered here – ASCs undermine trust in one of these two ways. Second, it is worth reiterating that my conception of public trust is *morally neutral*. As such, for my trust-based argument to work, it has to be shown that the practices being undermined by ASCs are in some sense *socially valuable*. That is, if trust is central to certain practices *and* these practices are socially valuable, then it follows that we have a *pro tanto* reason to stop ASCs undermining these practices (unless a counter-balancing reason can be found). I assume that both taxation and democratic principles – the two areas I consider – are socially valuable; taxation is often lauded for its role in creating a

[177] Citizens United v. Federal Election Comm'n, 558 U.S. 310 (January 21, 2010).

fair distribution of wealth.[178] Similarly, there are a number of reasons for thinking that democracy is at least instrumentally valuable, if not intrinsically valuable.[179] It seems plausible that all other practices that ASCs undermine would also be socially valuable (at least in democratic states), and thus this argument can generalise.

Claim 1: ASCs undermine the intrinsic commitment of ordinary citizens to particular social norms

Tax evasion involves the use of unlawful methods, such as under-reporting or failing to declare one's income, to ensure that one does *not* fulfil one's tax obligations under the law. Typically, this is distinguished from tax *avoidance*, which involves wholly legal mechanisms to reduce one's tax burden.[180] On this distinction, tax evasion involves breaking the law whereas tax avoidance in principle obeys the letter of the law, yet may use a legal loophole that contravenes the spirit of the initial law.[181] However, this distinction fails to adequately capture a number of actions, largely performed by multinational corporations, that exist on a spectrum somewhere between both extremes and are likely to be deemed illegal if brought up in court.[182] The point is, tax avoidance has become a term that is used far too liberally, and often by corporations in an attempt to legitimise their otherwise-illegal tax minimising actions.[183] As such many

[178] For example, see Brock, "Taxation and Global Justice"; Gillian Brock, "Global Justice," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, 2017, https://plato.stanford.edu/archives/spr2017/entries/justice-global/. Note, I think even libertarians would accept this point; they would either concede that taxation is valuable for redistribution yet object to it on other grounds. Or, they would see (minimal) taxation as being valuable to sustain a minimal state that protects the rights and freedoms of its citizens.

[179] For an overview of this literature, see Thomas Christiano, "Democracy," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Metaphysics Research Lab, Stanford University, 2018), https://plato.stanford.edu/archives/fall2018/entries/democracy/.

[180] Tax Justice Network, "Tax Avoidance," accessed May 22, 2019, https://www.taxjustice.net/faq/tax-avoidance/, https://www.taxjustice.net/faq/tax-avoidance/.

[181] Although intuitively plausible, this 'letter-spirit' distinction when applied to the law runs into problems in many cases where loopholes are exploited. For a full discussion, including an argument that exploiting a legal loophole is somewhat unproblematic, see Leo Katz, *Why the Law Is So Perverse* (Chicago ; London: University of Chicago Press, 2011).

[182] Tax Justice Network, "Tax Avoidance."

[183] For an alternative theory of tax avoidance that can circumvent some of these issues, see David Quentin, "Risk-Mining the Public Exchequer," *Journal of Tax Administration* 3, no. 2 (December 6, 2017), http://jota.website/index.php/JoTA/article/view/142.

cases labelled as tax avoidance are, in fact, cases of tax evasion. Now, ASCs are usually, but not exclusively, used in tax evasion schemes as they provide a mechanism for individuals to have undeclared assets *without* risk of being discovered. My goal is then to show how using ASCs in this manner – be it for tax evasion or tax avoidance – can undermine the intrinsic commitments individuals have to the rules or norms governing the practice of taxation. To do so, I need to first establish that practice rules are involved in preventing tax evasion.

In the 1970s, rational choice theory – due to its prominence – was ubiquitously applied by economists to a wide range of policy areas that *prima facie* were not usually associated with traditional economics.[184] Initially, the main assumption of rational choice theory was that humans are self-interested such that, when faced with a set of options, they would only choose the option that maximised their utility. As such, the first theoretical examination of tax evasion concluded that, given the audit and fine rates, it is in one's self-interest to evade taxes.[185] Put differently, this 'portfolio model' of tax evasion plausibly suggested that rates of tax compliance are dependent on the probability of being caught, which on their view was measured by examining the rates at which random audits and fines are imposed on an individual. Therefore, governments should increase the frequency of their audits and the severity of the punishments for tax evaders in order to solve the problem.[186] However, the portfolio model implies that – based on the central assumption of (early) rational choice theory – the *only* reason people pay their taxes is out of fear of being detected and/or punished, since *ex hypothesi* the would-be tax evader is weighing the economic benefit of tax evasion against the risk of punishment.[187]

[184] This is sometimes also called the 'economics of crime' model or approach. For a useful overview, see Alm, "Measuring, Explaining, and Controlling Tax Evasion."
[185] Michael G. Allingham and Agnar Sandmo, "Income Tax Evasion: A Theoretical Analysis," *Journal of Public Economics* 1, no. 3/4 (November 1972).
[186] Ibid.
[187] Alm, "Measuring, Explaining, and Controlling Tax Evasion."

There are a number of problems with a purely rational choice approach – for one, it is far from obvious that humans are utility-maximising agents that *always* choose the option that best promotes their interests.[188] However, examining the myriad criticisms of rational choice theory is unfortunately beyond the scope of the thesis. All that is required here is to show the limits of the portfolio model for tax evasion, which can be done by appealing to the available empirical evidence. Put simply, the portfolio model fails to account for the high rates of tax compliance *despite* low enforcement rates. The chance of one's tax returns being subject to a rigorous audit is less than 1% in almost every country, while any penalties for tax evasion are not always enforced and, when they are, do not exceed the amount owed to the state.[189] Yet, notwithstanding these enforcement rates, countries with the lowest tax compliance have significantly fewer cases of tax evasion than predicted by the portfolio model. As such, it seems clear that the rational choice model of tax evasion is inaccurate, and there are other factors involved in an individual's unwillingness to evade tax.

One prominent explanation is that social norms play a significant role in ensuring tax compliance.[190] The thought is that the commitment individuals have to norms of taxation give them a reason to fulfil their fiscal obligations that extends beyond the fear some individuals may have of being caught. Now, I argue that we have reason to think this commitment is intrinsic – in fact, it seems our rejection of the portfolio offers some support for this claim. The reason for this, I contend, is that instrumental commitments tend to be grounded in self-interest, since such a commitment is maintained insofar as it is aligned with the agent's wider interests. As we have seen that pure self-interest

[188] There is a wealth of literature, based on developments in psychology, challenging the assumption that humans are inherently self-interested. For an interesting (albeit brief) overview, see Ibid., 62–64. For an accessible overview of behavioural economics see George A. Akerlof and Robert J. Shiller, *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism* (Princeton: Princeton University Press, 2009).

[189] Alm, "Measuring, Explaining, and Controlling Tax Evasion," 119. To illustrate, consider the fine Apple had to pay to the EU for tax evasion. Here, the fine was calculated was based on just 10 years of unpaid taxes even though the favourable tax arrangement had existed for much longer – see https://www.theguardian.com/business/2016/aug/30/apple-pay-back-taxes-eu-ruling-ireland-state-aid

[190] Donna D. Bobek, Amy M. Hageman, and Charles F. Kelliher, "Analyzing the Role of Social Norms in Tax Compliance Behavior," *Journal of Business Ethics* 115, no. 3 (July 2013): 451–68.

accounts fail to explain the high rates of tax compliance, it seems to follow that the empirical evidence points to non-self-interest-based accounts – which, if we accept the social norm view, would point to an intrinsic commitment to a social norm of tax compliance. Perhaps this is too quick; strictly speaking, the portfolio model suggested that individuals pay taxes out of fear of being caught *by the state* but did not factor the fear of censure by one's peers. This would then imply that people pay taxes ought of fear of censure *by their peers*.

Although plausible and not captured by Allingham and Sandmo's account, I do not find this convincing. For one, there are huge informational asymmetries between tax evaders and their fellow citizens such that it would be extremely difficult for the average individual to uncover that their neighbour failed to pay taxes. Further, ASCs and other offshore vehicles only serve to increase these informational asymmetries and make it harder to uncover perpetrators of tax evasion. If the effect is indirect (e.g. if the state uncovered tax evasion and disseminated information about the evaders to the public), then it seems that this alternative framing of the self-interest argument just collapses back into that offered by the portfolio model, since one of the reasons people pay taxes would be a fear of getting caught. There is also a third option; perhaps the social censure tax evaders face is so high that individuals pay taxes in order to avoid the *risk* of being made into a social pariah. Again however, this seems unlikely – at least at the individual level. There have been a number of instances where prominent individuals have been arrested, fined, and even imprisoned on charges of tax evasion to very little public outrage.[191] Therefore, it seems that the empirical evidence supports the view that tax compliance is *not* driven by pure self-interest. In what follows, I draw on the social science literature showing how social norms play some role in preventing tax evasion. As I aim to show that ASCs undermine impersonal trust, I show that this literature can be recast in terms of practice rules without changing the results. Therefore, if ASCs

[191] Consider for instance the case of Uli Hoeness, who was imprisoned for tax evasion and reinstated as President of Bayern Munich upon his release – see https://www.dw.com/en/uli-hoeness-steps-down-from-bayern-munich-presidency-oliver-kahn-joins-board/a-50203363. Interestingly, the same is not true at a group level; as we will see, the Panama papers

undermine the intrinsic commitments we have to the rules governing the practice of taxation, impersonal trust is also undermined.

There is significant empirical evidence showing that social norms play some role in reducing levels of tax evasion. In particular, there are a number of experimental studies run by economists showing that the presence of norms of tax compliance drastically reduces levels of tax evasion.[192] Although the papers in question examined a multitude of different factors associated with social norms, the general conclusion seems to be that social norms play a significant role in ensuring tax compliance. It is beyond the scope of this thesis to outline the intricacies of each paper, yet two points are worth highlighting. First, in instances where enforcement rates are low, information sharing about tax avoidance schemes *reduces* levels of tax compliance.[193] That is, the more agents discover that others are flouting norms of tax compliance, the more they themselves begin to think that they should not abide by these norms of taxation. Thus, as more agents weigh the payoffs received by their neighbours through tax evasion schemes in their own decisions, the more overall tax compliance rates decrease. It seems then that information about the actions of others can undermine one's commitment to a social norm. Put simply, if one is faced with the prospect that everyone in their social group is failing to abide by a given social norm, then it is likely that one's own intrinsic commitment to that norm will fall away.[194]

Second, social norms have both a direct and indirect effect on tax compliance. Specifically, an economic study found that one's *subjective* and *personal* norms directly influence one's compliance with tax law, while *injunctive* and *descriptive* norms have an indirect effect on one's compliance.[195] The four-way distinction of social norms is drawn from the social psychology literature and, although interesting, the two categories

[192] For a helpful overview, see Alm, "Measuring, Explaining, and Controlling Tax Evasion"; Bobek, Hageman, and Kelliher, "Analyzing the Role of Social Norms in Tax Compliance Behavior," sec. 1.
[193] Adam Korobow, Chris Johnson, and Robert Axtell, "An Agent–Based Model of Tax Compliance with Social Networks," *National Tax Journal* 60, no. 3 (September 2007): 589–610.
[194] I will draw on this discussion later in examining the particular mechanisms by which ASCs undermine one's intrinsic commitment to norms of taxation.
[195] Bobek, Hageman, and Kelliher, "Analyzing the Role of Social Norms in Tax Compliance Behavior."

relevant for my purposes are injunctive and subjective norms, since these have the necessary features for impersonal trust.[196] Roughly, subjective norms correspond to the perception one has about what their close relations ought to do in a given situation, whereas injunctive norms refer to the perception of what most people think others ought to do in a given situation.[197] The difference between the two is essentially the subject of perception. That is, an injunctive norm is more akin to the "moral rules of the group" one is a part of, whereas a subjective norm is the perception one has about the injunctive norms one's close relations hold.[198] The point is that – in line with the findings in the preceding paragraph – the subjective norms one holds *directly* influence tax compliance while injunctive norms have an indirect influence, where this is understood as shaping the *content* of one's subjective norms.

Taken together, the empirical evidence points to a strong link between social norms and tax compliance. Further, as noted in our rejection of the portfolio model, there is moderate support for our *intrinsic* commitment to these norms of taxation, as an instrumental commitment to these norms can usually be couched in terms of self-interest and thus would be vulnerable to the issues outlined above. Note that this support is only moderate; in my rejection of the portfolio model, I only considered arguments stemming from a fear of censure. That is, I only showed that arguments from self-interest, where this is understood as *not* being censured, fail to adequately explain high rates of tax compliance. However, it could be argued that individuals adhere to norms of taxation not due to an intrinsic commitment to these norms, but rather due to a belief that taxes are the most effective ways of redistributing wealth in society. This would remain an instrumental commitment to norms of taxation; individuals view taxes as being useful tools in the pursuit of a good (e.g. equality) that they see as having intrinsic value. I suggest that, by reframing this literature through the lens of practice

[196] For a full outline of this theory of social norms, see Robert B. Cialdini and Melanie R. Trost, "Social Influence: Social Norms, Conformity, and Compliance," in *The Handbook of Social Psychology*, ed. Daniel Todd Gilbert, Susan T. Fiske, and Gardner Lindzey, 4th ed (Boston: McGraw-Hill, 1998).
[197] Ibid.
[198] Bobek, Hageman, and Kelliher, "Analyzing the Role of Social Norms in Tax Compliance Behavior."

rules, we may be able to see why a commitment to norms of taxation may in fact be intrinsic.

The distinction between injunctive and subjective norms outlined previously can be captured in terms of practice rules. Recall that injunctive norms refer to the 'moral rules of the group', while subjective norms refer to the ways in which individuals think their close relations (i.e. friends, colleagues, family etc.) interpret these injunctive norms. I think that injunctive norms just refer to practice rules themselves, whereas subjective norms refer to how these practice rules are interpreted by the individuals we take to be important in our lives. To see this, consider the game of football. It is a practice rule in football that the role of a goalkeeper is to stop their team from conceding a goal. This has the features of an injunctive norm. All footballers and football fans think that this is what goalkeepers ought to do and, if a goalkeeper started playing far away from their team's goal, most people would think they were doing something wrong. However, certain goalkeepers have interpreted this rule in different ways; Manuel Neuer for example became famous for venturing outside his team's box to clear away any dangerous passes. Thus, the players on Neuer's team would alter the way they play football, based purely on their perception of Neuer's interpretation of one of the practice rules of goalkeeping.

The point is, given that injunctive and subjective norms can be adequately captured through the lens of practice rules, we can draw on a wide body of empirical evidence to show that practice rules play some role in ensuring individuals do not evade taxes. Further, by viewing taxation as a practice with its own set of practice rules, we can understand how impersonal trust plays a role in ensuring tax compliance. First, can taxation be viewed as a practice? This I take to be uncontroversial; taxation is an institutionally defined activity with a set of rules (e.g. only individuals earning above a certain threshold are expected to pay) that are logically prior to the actions of individuals. Put differently, what distinguishes taxation from charitable giving or paying a fee is the creation of rules that denote 'paying taxes' as a new form of activity, which

is the very definition of a practice rule. Second, if taxation is a practice, we have reason to think that an individual's commitment to the rules governing taxation are intrinsic rather than instrumental. Taxation *as a practice* is usually justified on instrumental grounds, for example by appealing to the effectiveness of taxes in redistributing wealth. Crucially however, this instrumental justification for taxation is predicated on the assumption that *every* eligible taxpayer (or at least almost every eligible taxpayer) will pay their taxes – otherwise the whole practice would collapse.

As such, it seems plausible to say that, while the practice itself may be instrumentally justified, individual taxpayers have an intrinsic commitment to the practice rules of taxation. We have already seen how self-interest explanations fail to adequately explain the high levels of tax compliance. In addition, if taxpayers were instrumentally committed to the practice rules of taxation, they would opt-out of paying if they perceived alternative arrangements as being a better tool to achieve their respective ends. Hence, we can plausibly say that intrinsic commitments to the practice rules of taxation – and therefore impersonal trust – play a role in ensuring tax compliance. Further empirical evidence can help to establish this point: In the aftermath of the Panama Papers leak, when it was discovered that criminals, politicians, and wealthy elites were using ASCs to avoid paying taxes, people across the globe displayed a wide range of reactive attitudes. In Iceland for example, after the Prime Minister, Interior Minister, and Finance Minister were all implicated in the leak, 6% of the entire population took to the streets in protest, which was unprecedented in Icelandic politics.[199] The protestors frequently characterised their reactions as feelings of "anger, shame, and betrayal" – all of which are reactive attitudes that offer support for the presence of impersonal trust.[200]

How do ASCs undermine the assumption that individuals in the same society share an intrinsic commitment to norms of taxation? I argue that the aggregation of individual

[199] "Voices from Iceland as PM Resigns," *BBC News*, April 5, 2016, sec. Europe, https://www.bbc.com/news/world-europe-35967320.
[200] Ibid.

trust-undermining actions will, over time, make it no longer plausible for this assumption to hold. The general point I am making is simply that, as more and more individuals evade taxes, the less plausible it will be for any truster to maintain the assumption that the trusted party shares their intrinsic commitment to practice rules of taxation. At that point, public trust in certain practice rules of taxation (e.g. that all taxpayers have to pay) has eroded, though the practice of taxation would still exist. It would however be a practice where individuals are either instrumentally committed, or not committed, to certain practice rules governing it. We can characterise this trust-undermining mechanism in terms of a snowball effect; some individuals begin to use ASCs to evade taxes and may – directly or indirectly – disclose this to their close relations.[201] Some of these close relations will find that their intrinsic commitment would be shaken, but remain intact while others may feel betrayed by the actions of their close relations. Yet, a few may also see their intrinsic commitment to norms of taxation erode, and thus begin to find ways to start evading taxes as well. Again, this seems to be in line with the empirical evidence as, when individuals share information about tax evasion schemes, overall tax compliance rates decrease.[202]

Even if my specific characterisation of this mechanism is disputed, the general point I think holds. As information about trust-undermining activity is shared, the intrinsic commitments individuals have to a given practice rule will erode. Over time, if this process keeps happening, public trust in that specific practice rule will disappear as the society-wide assumption that there is a shared intrinsic commitment to that rule will no longer hold. Given the ease at which ASCs can be set-up – and the scale of the industry involved in setting them up – it is plausible to conclude that they have played an integral part in undermining public trust in taxation. We see this for example in the case of the Panama Papers; wealthy elites across the globe used firms like Mossack Fonseca to set up ASCs and evade taxes in their respective countries. As noted in the introduction, what

[201] There is empirical evidence to support this claim; for instance Bradley Birkenfeld, a former banker for UBS, testified that he was flown by his then-company to events attended by wealthy elites in order to convince them to put their money in UBS accounts, which would remain secret and thus allow individuals to evade taxes.

[202] Korobow, Johnson, and Axtell, "An Agent–Based Model of Tax Compliance with Social Networks."

started as a service by Swiss banks to allow affluent Europeans to avoid high taxes following World War I snowballed into a gigantic industry that offered the same service for *anyone* on Earth. The only reason public trust remained intact for so long was the secretive nature of the industry itself. However, once the sheer scale of the industry was exposed, the society-wide assumption that individuals are intrinsically committed to norms of taxation became untenable – thereby undermining public trust.

Claim 2: ASCs facilitate the subversion of democratic practices

We have seen how the aggregation of individual actions can contribute to the erosion of public trust. Turning to my second claim, I aim to show that ASCs facilitate the subversion of democratic practices and thus undermine trust. To do so, I focus on the context of American democracy, as I think it has undergone a number of institutional changes that place it at a greater risk of being subverted in this way. The link between practice rules and the preservation of American democracy is not new. Two political scientists, Steven Levitsky and Daniel Ziblatt, argue that the commitment to two unwritten practice rules that have dictated political behaviour in the US since its inception are being eroded at an alarming rate, thus threatening democracy.[203] Specifically, they suggest that the erosion of norms of mutual toleration, where each party views the other as legitimate, and forbearance, where congressional powers are used with restraint by government officials, is causing American democracy to backslide towards a dictatorship.[204] Whether or not the erosion of democracy can be linked to the erosion of these two particular norms is beyond the scope of this thesis, though I am sceptical as they fail to consider one prominent area that has dangerous implications for democratic values – namely the role of money in politics. In what follows, I first highlight how the debate about money in American politics has changed from the Constitutional Convention in 1787 to the Supreme Court ruling in *Citizens United v. FEC* in 2010. These changes, I argue, have created a situation where ASCs have the ability to influence

[203] Steven Levitsky and Daniel Ziblatt, *How Democracies Die*, Kindle Book, First edition (New York: Crown, 2018).
[204] Ibid., 23.

politics to a significant degree. This, in turn, has the potential to destabilise the commitment we have to democratic principles – thereby undermining trust.

At the Constitutional Convention, delegates from across the fledgling United States came together to discuss the central laws that would govern their new country. Americans sought to distance themselves from the 'Old World' decadence they felt was embodied by European states in general, and England in particular – where prized offices were handed out by the nobility to their close relations.[205] As such, the Americans were particularly concerned about *corruption*. In fact, they saw corruption as being the greatest threat to their young republic, since it would recreate the system of dependency between the ruling classes and their subjects that they had rebelled against during the Revolutionary War. For these reasons, delegates at the Convention spent more time discussing the prevention of corruption than instability, factionalism, and violence.[206] Thus, it could be said that one of the goals of the Convention was to create strong legal bulwarks to prevent the decadence of Europe from affecting the American polity. The historical record provides significant evidence for this interpretation; for instance, James Madison wrote the word 'corruption' 54 times in his notebooks from the Convention.[207] Similarly, Alexander Hamilton, when defending the US Constitution, remarked "Nothing was more to be desired than that every practicable obstacle should be opposed to cabal, intrigue, and corruption".[208]

They had good reason to be concerned. Just two years earlier, in 1785, Benjamin Franklin was given a majestic snuff box, adorned with 408 of the finest diamonds, from King Louis XVI upon leaving his position as American Ambassador to Paris.[209] The giving of gifts to departing diplomats was a common activity in Europe at the time, yet from the American perspective was seen as emblematic of the corruption that plagued the

[205] James D. Savage, "Corruption and Virtue at the Constitutional Convention," *The Journal of Politics* 56, no. 1 (February 1, 1994): 175; Teachout, *Corruption in America*, 63.
[206] Teachout, *Corruption in America*, 57.
[207] Savage, "Corruption and Virtue at the Constitutional Convention," 177.
[208] Alexander Hamilton, James Madison, and John Jay, "Federalist No. 68," in *The Federalist Papers* (New York: Palgrave Macmillan, 2009).
[209] Teachout, *Corruption in America*, 25–26.

'Old World'. As such, the snuff box caused great consternation amongst Franklin's American contemporaries once he returned – with many of them labelling Franklin a corrupt influence. Their worry was that, by bringing the snuff box back with him, Franklin was importing the corrupt traditions that the Americans had fought so hard to free themselves from. This concern was in some sense encapsulated in the Articles of Confederation – the precursor to the Constitution – which prohibited any American office-holder from accepting "present, emolument, office, or title of any kind whatever, from any king, prince, or foreign State".[210] After much discussion, this same clause was incorporated into the US Constitution (as Article I Section 9) for reasons of anti-corruption.[211]

The root cause of this worry was one of influence. American delegates at the Constitutional Convention realised that the giving of lavish gifts could in some sense buy *influence*: receiving lavish gifts would make elected officials more amenable to the foreign powers in question, and thus undermine their objectivity in representing American interests.[212] Note that this is different from a narrower conception of corruption as *quid pro quo*, which would involve the giving of a gift in exchange for a specified favour. Rather, the Americans were worried that foreign countries with significant capital could – through cash or gift-giving – influence the opinions of the American legislature for their own ends. This definition came to be the dominant understanding of corruption in American society and acted as an unwritten norm regulating American political practice over money until the 1970s, following the Supreme Court ruling in *Buckley v. Valeo*.[213] Put differently, the dominant reading of corruption in America disapproved of the giving of gifts or cash on the grounds that it risked giving powerful foreign countries or wealthy individuals undue influence in the American political system. My goal in this section is to show that the legacy of *Buckley*,

[210] Ibid., 20.
[211] Savage, "Corruption and Virtue at the Constitutional Convention," 182.
[212] Teachout, *Corruption in America*, 30–31.
[213] Teachout, *Corruption in America*.

coupled with the growing use of ASCs, means that the American polity may face these same worries in the 21st Century – over 200 years after the Constitutional Convention.

In 1976, the Supreme Court ruled in *Buckley v. Valeo* that limits on election spending constituted a violation of the First Amendment. The implication of this ruling was the formal recognition that the spending of money constituted a *speech act* rather than a potential way to influence public or legislative opinion in particular ways.[214] In making their ruling, the Court recognised the civic interest in preventing corruption – yet did not see limits on election spending as having a role to play in facilitating corrupt practices. Consequently, the framing of the debate around corruption suddenly changed; where previously the giving of cash or gifts was seen as a potentially corrupt action, if done towards elected officials, now it was seen as a form of speech that needed protection as per the rights grounded in the First Amendment. That is, the legacy of the *Buckley* ruling was a bizarre judicial framework where the value of anti-corruption was pitted against the value of protecting the rights individuals have to freedom of speech.[215] This framework became the basis for *every* legal challenge to the presence of expenditure limits in politics, culminating in the landmark ruling in *Citizens United* where this hitherto dominant interpretation of corruption was supplanted.

In the *Citizens United* ruling, the Supreme Court reached two consequential verdicts. First, by viewing political spending as a type of speech act, they ruled that *any* limitations on political expenditure constituted a violation of the right individuals (and corporations) have to freedom of speech.[216] Here, political expenditure refers to the use of money or gifts for *any political purpose* – including the giving of donations to the political candidates and parties of one's choice. In short, then, this ruling gave wealthy elites a carte blanche to spend however much they wanted on furthering their own political causes and ambitions, which undoubtedly gives them an undue influence on

[214] Buckley v. Valeo, 424 U.S. 1 (January 29, 1976).
[215] Teachout, *Corruption in America*, 6–7.
[216] Citizens United v. Federal Election Comm'n, 558 U.S. 310.

the political process in the US.[217] Second, Justice Kennedy made a clear break with the influence-based understanding of corruption; writing in the majority opinion, he argued that corruption was a particular brand of *quid pro quo*, where money or gifts are given in exchange for an explicitly defined, specific legislative act.[218] By explicitly interpreting corruption in this way, Kennedy equated corruption with vote-buying and therefore legitimised actions that for over 200 years were labelled corrupt.[219] Now, for example, it is completely legitimate for wealthy elites to spend money on political campaigns *with the intent* of influencing the respective candidate – so long as no explicit conditions or agreements are attached to the donation.

Of course, the giving of gifts or donations does not *necessarily* point to the presence of corruption: gifts, after all, can just be gifts. Yet, it is striking that the fears and worries raised by the Framers of the Constitution once again threaten to raise their head and undermine the independence of American democratic institutions. In this context, I argue that the mere presence of ASCs risks subverting the goals of democracy. In much the same way that Captain Dudley's machinations turned the police department into a crime syndicate, thus subverting the purpose of the police department, the use of ASCs in the sphere of American politics post-Citizens United risks transforming a democracy into a plutocracy with a thin democratic veneer. This will do more than just undermine public trust in a particular domain. Instead, it will jeopardise public trust in our democratic values and institutions and make it nigh impossible to develop a non-corrupt and prosperous society. Part of the threat ASCs pose to American democracy stems from the wider worry about money corrupting politics; after all, ASCs are effective tools for cross-border financial flows and thus would be used to funnel money into one's preferred candidates. Yet, there is nevertheless a distinctive issue posed by ASCs in this context: the giving of political donations *anonymously*. In such cases, while the politician is often made aware of the identity of the donor, the public has no way of knowing who

[217] Lessig, *Republic, Lost*; Teachout, *Corruption in America*, 7.
[218] Citizens United v. Federal Election Comm'n, 558 U.S. 310. This interpretation of corruption was echoed by Justice Roberts in 2014 – see his concurring opinion in McCutcheon v. Federal Electoral Commission 572 U.S. 185 (April 2, 2014).
[219] Teachout, *Corruption in America*, 232–33.

is funding particular campaigns or political ends. As we will see, this threatens to undermine democratic accountability and – with it – the very practice of democracy itself.

I argue that the significant presence of money in American politics risks undermining one of the central practice rules of democracies: that the legislature is beholden to the people. Put differently, money is being used by wealthy elites to influence American politics such that it threatens the independence of the legislature.[220] Further, the *Citizens United* ruling will only serve to exacerbate this effect. My argument proceeds in two parts. First, I highlight the problems associated with money in politics in general, and show why ASCs are tarred with this same brush. Second, I examine how the presence of anonymous wealth in politics compounds this problem even further, by emphasising the implications this has for democratic accountability. In both steps, I show how trust is undermined and thus gives us a *pro tanto* reason to set up a public ownership register. Note that my argument is not *viewpoint neutral*; central to my argument is the assumption that the wealthy have a distinctive set of political views and interests that differ from other citizens with less material wealth. After all, if political views and interests were randomly distributed amongst the different economic classes, we would not take issue with the presence of money in politics.[221] However, given the current social, political and economic context, I take this assumption to be uncontroversial.

In the US context, money in politics is governed by two different mechanisms. First, money is used as a gatekeeper to prevent either individuals or ideas from entering political discourse. This has the effect of distorting the agenda-setting and deliberative processes of an institutional body, with elected officials promoting values and policies that would benefit the elites at the expense of the general population.[222] We see this if

[220] Lessig, *Republic, Lost*.
[221] Thomas Christiano, "Money in Politics," in *The Oxford Handbook of Political Philosophy*, ed. David Estlund (Oxford University Press, 2012), 254.
[222] Ibid., 245.

we closely examine American election campaigns; as the contest between both parties to secure a majority in Congress intensifies, the volume of wealth needed to launch and sustain a campaign increases. Thus, Democrats and Republicans regularly rely on their legislators to raise funds from both corporations and the affluent, in order to increase the amount of money each party can use to fight and win an election – to the point where members of Congress spend between 30 and 70 percent of their time fundraising.[223] This in turn decreases the amount of time elected officials spend legislating; between 1983 and 1997, the number of committee meetings for the purpose of regulating government expenditure fell from 782 to just 287 in the House of Representatives.[224] Simultaneously, as legislators seek to source as much funding as possible for the party, they may make policy concessions that will directly benefit the wealthy elites at the expense of the rest of society. For instance, after losing their majority in the Senate and the House in 1994, the Democratic Party adopted a much more pro-business approach to woo Wall Street financiers into donating vast sums of money to the party so that they could more effectively campaign against the Republicans.[225]

Taken together, the use of money as a gatekeeping mechanism results in the legislative process being subverted for the aims of serving the interests of the elite rather than the population at large. Empirical evidence provides further evidence for this conclusion: the political scientist Larry Bartels has argued that the votes of US Senators are *not at all* responsive to the interests of those in the bottom third of the income bracket while only being minimally responsive to those in the middle third.[226] Though shocking, this may not be exclusively due to the influence of money on the agenda-setting powers of governments; there is also a long history of money being used to influence both legislative and public opinion. For instance, multinational corporations and wealthy

[223] Lessig, *Republic, Lost*, 138.
[224] Norman J. Ornstein, Thomas E. Mann, and Michael J. Malbin, *Vital Statistics on Congress 2008* (Washington, DC: Brookings Institution Press, 2008).
[225] Lessig, *Republic, Lost*, 96.
[226] Larry M. Bartels, *Unequal Democracy: The Political Economy of the New Gilded Age*, Second edition (New York: Russell Sage Foundation, 2016).

elites often fund lobbying groups with the aim of influencing elected officials to adopt particular policy platforms that directly benefit them. In American politics, lobbyists exert such influence that the promotion of corporate interests serves as a 'constraint on widespread reform'.[227] In the public sphere this is no different; privately-funded think tanks and political groups tend to have much better funding than grassroots organisations, and thus have the ability to run more frequent advertising to influence the opinion of the general public so that they agree with the political interests of the best-off in society.[228] John Roemer, for instance, has argued that the erosion of a social ethos in the contemporary United States is partly due to the role of privately-funded neo-liberal think tanks advocating an agenda of fervent individualism and deregulation.[229]

As corporate vehicles that can facilitate the transfer of money from one area to the next, ASCs are symptoms of the same problem. The affluent use ASCs to fund political campaigns, lobbying groups, and other political organisations to buy political influence and thus subvert one of the central practice rules of democracy: that an elected government is beholden *only* to its people.[230] In doing so, ASCs bring to life the worries that plagued the original Framers of the Constitution – namely that cash and gift-giving can undermine the objectivity of elected officials, and thus make them mere instruments for the ends of others. This is especially true in the aftermath of *Citizens United*, as the affluent are no longer beholden to constraints on their political expenditure; rather, they now have the freedom to spend as much as they desire to pursue their own political ends. I take this to be a compelling argument against the *unconstrained* use of money in politics, of which ASCs are a part. Yet, there is a sense in which transparency would mitigate some of these effects; if eligible voters had access to information about political donations and a given official's voting record, they could make a decision for themselves as to whether their representative was unbiased or

[227] Lessig, *Republic, Lost*, 184–85.
[228] Christiano, "Money in Politics," 249–50.
[229] John E. Roemer, "Ideology, Social Ethos, and the Financial Crisis," *The Journal of Ethics* 16, no. 3 (September 2012): 273–303.
[230] Lessig, *Republic, Lost*.

compromised in some manner. This brings us to the crux of the issue: ASCs compound the problems of money in politics further by ensuring that political donors remain anonymous to the public. Put differently, the affluent can use their finances to fund lobbying efforts and electoral campaigns of sympathetic politicians and thus buy political influence *without revealing their identity to the public*.

This is how ASCs risk subverting the entire practice of democracy. While the unconstrained use of money in politics undermines the political independence of government institutions, channelling these funds through ASCs risks making an elected government unaccountable to its own people. This is because eligible voters would have no way of knowing who is shaping the agenda of their governments – thereby compromising the entire purpose of a democratic system. This may seem conspiratorial and far-fetched, as if some hidden puppet-masters were pulling the strings of our elected officials. Yet there is evidence to suggest it is already happening. In 2016, the Guardian obtained the 'John Doe files', a series of leaked documents exposing how Scott Walker, the ex-Republican Governor in Wisconsin, raised millions of dollars in donations from wealthy conservatives and used this to swing elections in his favour since 2011.[231] In turn, Walker adopted a highly pro-business approach to attract as much funding as possible, perhaps epitomised by the passing of a bill granting "effective immunity to lead manufacturers from any compensation claims for lead paint poisoning" just after receiving a donation of $750,000 from NL Industries – one of America's leading lead producers.[232] Further, in the 2012 US election, ASCs funnelled $17 million to Super PACs, which are committees that can raise unlimited funds to campaign for or against particular candidates.[233] In both cases, ASCs facilitated the use of money to influence

[231] Ed Pilkington and the Guardian US interactive team, "Scott Walker, the John Doe Files and How Corporate Cash Influences American Politics," *The Guardian*, accessed September 25, 2019, http://www.theguardian.com/us-news/ng-interactive/2016/sep/14/john-doe-files-scott-walker-corporate-cash-american-politics.
[232] Ibid.
[233] Brendan Fischer and Blair Bowie, "Elections Confidential: How Shady Operators Used Sham Non-Profits and Fake Corporations to Funnel Mystery Money into the 2012 Elections" (US PIRG Education Fund, January 2013).

public and legislative opinion in favour of certain interest groups that do not reflect the will of the general public.

The implications this has for the preservation of democratic principles are significant. One of the central practice rules of democracy is that the legislature is chosen by 'the people' and is therefore beholden to them and *no one else*.[234] The assumption that all citizens and their governments share an intrinsic commitment to this practice rule is partly what drives individuals to vote in elections. On the one hand, individuals feel politically effective; that is, individuals see themselves as making their opinions known to the government purely on the basis of their voting a certain way.[235] On the other, the individual-level commitment is bolstered by governments responding to the demands of the electorate. If governments are effective at acting on the basis of the wishes of the electorate, then this strongly correlates with high voter participation.[236] However, in cases where governments cease to function in this manner and instead prioritise corporate interests, public trust vanishes entirely. This is because the very practice of democracy has been subverted; governments cease to be beholden to individual voters, who in turn become disaffected with the democratic procedure.[237] In such cases, governments procedurally adhere to certain practice rules of democracy (e.g. holding regular elections) while disregarding those that are crucial for its preservation (e.g. having an accountable legislature). Where this occurs, public trust – much like the case of Captain Dudley – ceases to be present as there is no longer the assumption that voters and governments share an intrinsic commitment to practice rules of democracy. As ASCs facilitate the flow of anonymous wealth to politics, they risk subverting democracy in this manner and therefore undermine public trust in democratic institutions.

[234] Lessig, *Republic, Lost*, 160.
[235] Steven J Rosenstone, Keith Reeves, and John Mark Hansen, *Mobilization, Participation, and Democracy in America* (New York: Longman, 2003), 144; Kevin Chen, *Political Alienation and Voting Turnout in the United States, 1960-1988* (San Francisco: Mellon Research University Press, 1992), 214, 217.
[236] Rosenstone, Reeves, and Hansen, *Mobilization, Participation, and Democracy in America*, 144.
[237] Lessig, *Republic, Lost*, 168.

In summary, ASCs undermine public trust in two distinct ways. When ASCs are used by some to evade taxes, they undermine the intrinsic commitments *individuals* have to norms of taxation. Over time, the aggregation of these individual trust-undermining actions lead to the erosion of public trust, as the general assumption that there is a shared intrinsic commitment to practice rules of taxation is no longer tenable. Under this mechanism, public trust is undermined by the erosion of an intrinsic commitment to particular practice rules. Yet, ASCs also risk subverting entire practices. Here, there is no longer an intrinsic commitment to *any* of the rules governing a given practice. By facilitating the flow of anonymous wealth into politics, ASCs risk creating a system where the legislature is no longer *exclusively* beholden to the people, which causes public trust in democratic institutions to vanish. This was a serious worry for the Framers of the US Constitution and today, over 200 years after its writing, there are signs to suggest that this worry is being realised. If my analysis is accepted, we can conclude that ASCs undermine public trust in socially valuable practices, which gives us reason to regulate them. The consensus amongst policymakers suggests that the creation of a public ownership register provides the best way to regulate ASCs. As we will see however, a simple public ownership register also poses serious questions of trust. Therefore, I conclude with a discussion concerning how the instrumental value of public trust gives us a *pro tanto* reason to adopt a suitably amended public ownership register.

Conclusion: Regulation and Public Trust

We began with the observation that trust is ubiquitous in our everyday lives. Without some kind of trust, cooperation would be near impossible, our societies risk being venally corrupt, and even our lives may lack meaning. Thus, it seems true to say that "a complete absence of trust would prevent [one] even getting up in the morning".[238] On this basis, we can conclude that trust is *at least* instrumentally valuable; its value comes from what trusting others allows us to achieve rather than its presence *simpliciter*. Again, empirical evidence has pointed to *inter alia* a strong correlation between high levels of trust and economic prosperity – as well as between trust and low levels of corruption. This was the foundational premise of my argument; taking the instrumental value of trust as my starting point, I sought to show that we can use trust to raise an insightful argument against the problems caused by anonymous wealth in general, and anonymous shell companies in particular. In short, I argued that ASCs undermine public trust in pernicious ways, which hinders our ability to achieve the instrumental benefits of high-trust societies. This in turn, gives us a *pro tanto* trust-based reason to implement policies that regulate ASCs. Now, I conclude by showing how this same *pro tanto* trust-based reason gives us a theoretical justification for the implementation of a suitably-amended public ownership register.

I began by providing significant evidence to supplement my claim that anonymous wealth is one of the most significant policy issues we face as a global society. Our campaigns to promote equal societies, foster global development, and combat transnational crime are all stymied by the presence and use of ASCs in our international financial system. They do so by allowing wealthy elites and multinational corporations to evade taxes, thus potentially creating shortfalls in government revenue that can only be met by increasing the fiscal burden on the middle and lower classes. Further, ASCs facilitate the sale and re-sale of natural resource assets by corrupt dictators who pocket massive profits instead of reinvesting this wealth in poverty alleviation projects. This is

[238] Niklas Luhmann, *Trust and Power* (Chichester: Wiley, 1979), 4.

a particular problem amongst a number of developing countries and regions: between 20% and 30% of all financial wealth across Latin American and African countries are being held in offshore tax havens, amounting to a tax loss of up to $21 billion.[239] Again, though tax havens and anonymous wealth are not equivalent, they are inextricably linked. The history of ASCs is bound with the growth and development of tax havens, initially in the form of Swiss banks before being exported to a multitude of different Caribbean and Pacific islands seeking an alternative revenue stream, and to this day the vast majority of companies incorporated in tax haven jurisdictions are ASCs.

In light of these negative externalities, some policymakers have argued that the creation of a public ownership register is the most effective method to tackle ASCs. Public ownership registers are essentially a database of information: they contain information on the beneficial owner – the natural person with *de facto* control of a company – of *every* company incorporated in a given country. As such, public ownership registers have the potential to deanonymize ASCs, thereby allowing governments to track monetary and financial flows within their jurisdictions much more effectively. Implementing such registers however has proved to be controversial. I established that the major obstacle to the creation of a public ownership register is the existence of an impasse within the policy sphere. Though policymakers agree that ASCs have negative externalities, they disagree over the best way that these should be combatted. Specifically, proponents of financial secrecy argue that – by disclosing beneficial ownership information to the public – public ownership registers violate the presumptive right individuals have to privacy. They suggest that, as public ownership registers can be used to track down the addresses of famous celebrities, it impacts their ability to exercise their privacy rights. Instead, it needs to first be shown that an ASC is implicated in wrongdoing *before* any ownership information is exchanged. Opponents of financial secrecy explicitly deny this. They argue that the only way to ensure no illicit activity is being undertaken is for *all* beneficial ownership information to be made public.

[239] Zucman, *The Hidden Wealth of Nations*, 53.

The longer this debate remains unresolved, the longer the status quo remains unchallenged. Currently, ASCs are facilitating a wide range of illegal activities that is causing harm and suffering to individuals across the globe. This point is uncontested between both sides, and yet the intractability of the policy debate is preventing any action from being undertaken to resolve this issue. I pointed to the existence of a compromise position: we could construct a public ownership register with certain provisions to ensure that the privacy rights of high-profile individuals were not infringed. For instance, individuals could apply for special provisions to allow them to disclose their ownership information to the government, but not to the public. Alternatively, the data stored on the public information register could be presented in such a way that no individual could be uniquely identified. This position however is merely a political compromise; we have no theoretical grounds to adopt this position beyond the realisation that doing so would be practically useful. I argued that, by considering the value of public trust, we could provide a theoretical justification for a public ownership register. Reaching this justification however requires, first, an adequate account of public trust and, second, an argument showing *how* ASCs undermine public trust.

In establishing the first claim, I argued that existing theories of public trust were inadequate on the grounds that they were *monist*. Theorists of trust often assume that every trust relation can be couched in the same terms; that is, they think that we *always* trust others in the same way. In film for example, the affirmation of trust (i.e. a positive response to the question 'Do you trust me?') is often seen as a meaningful and romantic gesture. In affirming trust, the companion takes her reliance on the protagonist to influence his decision-making in some way. Further, the companion would feel betrayed if, after affirming her trust, the protagonist ignored her or let her down. Now, suppose the accountant at my start-up posed me the same question. In affirming my trust, I am neither making meaningful gesture nor seeking to influence my accountant's decision-making in any way. Yet, if my accountant were to defraud my company, I would nevertheless feel betrayed as I expected him to act in accordance with the constitutive elements governing the relationship between an employer and employee. The

differences between these cases, I argued, points to the existence of a pluralist conception of (normative) trust. By distinguishing between personal and impersonal trust, we can account for the diverse range of trust relations that we see in any given society. I characterise this distinction based on the source of normativity; in personal trust, the source of normativity stems from the attitudes and/or mental states of the truster (i.e. the normativity is internal). In impersonal trust however, the source of normativity maps onto an external feature – namely the constitutive elements of a given relationship.

With this distinction in mind, I presented my preferred account of impersonal trust. I focused solely on impersonal trust because this better captures the trust relations in the public sphere. In short, I argued that impersonal trust involves the assumption that there is a shared intrinsic commitment to a particular social norm, where I interpret social norms as types of practice rules. Here, practice rules define a set of procedures that an individual must comply with in order to be seen as being a participant in that practice.[240] I defended this interpretation on two distinct grounds: first, practice rules share many of the same intrinsic features as social norms (e.g. internal consistency, normativity etc.). Second, adopting this interpretation sheds new light on the nature of trust in the public sphere. For one, it makes our understanding of social roles seem less mysterious; we can simply understand them as a particular instantiation of a practice rule. Taken together, I showed that this interpretation overcomes the deficiencies that plague other theories of public trust and argued that, when applying it to the case of ASCs, we can clearly see how they undermine public trust.

In establishing the second claim – to show how ASCs undermine public trust – we need to first clarify what it means for public trust to be undermined. I showed that, when we talk about public trust, we are referring to an assumption at the societal level that there is a shared intrinsic commitment to a particular practice rule. If it is no longer the case that this assumption is tenable, then public trust has eroded. Using this conception, I

[240] Rawls, "Two Concepts of Rules," 25.

argued that ASCs undermine public trust in two distinct ways. First, the aggregation of individual trust-undermining actions will – over time – lead to a situation where the general assumption that there is a shared commitment to a *particular* practice rule no longer holds. Second, by threatening to subvert entire practices, ASCs risk eroding the intrinsic commitments individuals have to an entire set of practice rules – thereby undermining public trust as it will no longer be possible to assume that there is a shared commitment to a given practice. Taken together, we can identify general mechanisms by which ASCs erode public trust and thus establish the presence of a *pro tanto* reason to promote regulatory policies that stop ASCs from undermining public trust.

We are now in a position to see how the lens of public trust can provide us with a theoretical justification for the creation of a suitably amended public ownership register. Central to this argument is the assumption that public ownership registers are the most effective policy tools to stop ASCs. As this position is endorsed by a number of policymakers and NGOs, I take this assumption to be uncontroversial. Now, we have seen that public trust is instrumentally valuable. Thus, we automatically have a *pro tanto* reason to regulate actions that undermine it. If my above analysis is accepted, and the reader agrees that ASCs undermine public trust, then we can establish the following negative thesis: We have a *pro tanto* reason to implement a public ownership register as they stop the undermining of public trust by ASCs. This however is too quick. If a government created a public ownership register and forcibly made beneficial ownership information public, then – though this would stop ASCs from undermining public trust – it would blatantly disregard the individual privacy rights of its citizens. Therefore, implementing such a policy would undermine trust to a countervailing degree.

To see this, we can think of a government or a state as a particular practice of authority. We can distinguish (at least democratic) governments from general authorities on the grounds that governments safeguard the rights of its citizens. Put simply, it is a practice rule of (at least democratic) governments that they protect and maintain the rights of its citizens. Thus, if governments sacrifice the protection of rights in favour of achieving

some overarching good, then *inter alia* public trust would be eroded. This is because the assumption that governments share an intrinsic commitment to the protection of the rights of its citizens would no longer be tenable. Yet, we can amend a public ownership register in certain ways so that privacy rights are *not* undermined. Recall that, in characterising the position advocated by proponents of financial secrecy, I used Anabelle Lever's democratic theory of privacy to show that this presumptive right to privacy is only justifiable if it serves to protect one's freedom and equality in either the political or personal domains. So, on this conception, for the public disclosure of beneficial ownership information to be deemed a privacy violation, this disclosure would have to undermine an individual's personal freedom and equality, or political freedom and equality.

As I briefly discussed at the end of Chapter 1, this presumptive right to privacy is much weaker than proponents of financial secrecy would like. It would, for instance, justify privacy provisions for celebrities like Emma Watson, whose personal freedom would be undermined if her beneficial ownership information were to be disclosed. Recall that Watson owns her property through an ASC to avoid having her address uncovered; as such, protecting Watson's privacy in this case would give her the freedom to pursue her own ends and goals without being accosted by crazed fans at every opportunity. It would not, however, justify privacy provisions for political donors in the context of American politics, since anonymity in these cases does not safeguard the donor's freedom or equality in either the personal or political domains. I contend that, in most cases, publicly releasing beneficial ownership information will undermine neither freedom nor equality in either domain. After all, most companies are transparent about who the beneficial owner is. In fact, the only time beneficial ownership information is hidden is if an ASC (or other similar corporate vehicle) is incorporated in a given jurisdiction. Of course, there may be some legitimate uses for ASCs, as in the Emma Watson case. For the most part however, ASCs are founded with the explicit purpose of keeping the beneficial owner hidden from prying eyes.

In light of these comments, we can offer the following reformulation of the negative thesis: We have a *pro tanto* reason to implement a *suitably amended* public ownership register as this stops public trust being undermined by ASCs. Here, I take 'suitably amended' to refer to the addition of certain procedures that respect the presumptive right individuals have to privacy. Again, the negative thesis is simple: public ownership registers deanonymize ASCs and thus stop them from facilitating illicit monetary flows. For example, ASCs allow individuals to hide their assets from their own governments, which facilitates tax evasion. However, if a public ownership register were in place, this becomes near impossible as tax investigators simply have to look up the name of the company in the register of the country it is incorporated in to find who the beneficial owner is and whether they have declared their assets. Thus, by being sensitive to privacy considerations, we have shown how public trust can establish a theoretical justification for a suitably amended public ownership register.

Before closing, we can ask whether this argument meets the two desiderata I set out in Chapter 1. I suggested that, to be successful, an argument for the regulation of ASCs had to be sensitive to two different considerations. First, the argument has to be sensitive to the sheer number of contexts ASCs operate in. Due to their versatility, ASCs have been used in cases ranging from money laundering to brokering illegal arms deals. As such, any argument that *only* supports piecemeal regulation (i.e. regulation across some but not all contexts) would fail to succeed. Second, the argument has to be sensitive to the positive aspects of ASCs. Put differently, if an alternative argument for the regulation of ASCs were not sensitive to the privacy considerations of wealthy elites, then we would simply reach another impasse where privacy rights are traded off against a different desirable value that we seek to promote. My argument is at least partially successful in meeting these desiderata. As we have seen from the preceding discussion, the second desideratum has been met: a *pro tanto* trust-based justification specifically argues against a public ownership register *simpliciter* on the grounds that it would undermine the public trust between the government and those whose privacy rights would be violated.

With respect to the first desideratum, my argument is only partially successful. Though my theory of public trust can account for the myriad ways in which ASCs are trust-undermining, the *pro tanto* trust-based justification I offer for regulating ASCs only functions in a democratic context. Put differently, my theory of public trust can determine the mechanisms by which ASCs undermine trust in both democratic and non-democratic contexts, simply by pointing to the erosion of a general assumption that any truster shares an intrinsic commitment to a practice rule with any trusted party. However, to ground an argument for regulation, either public trust must have some instrumental value or the social practices sustained by public trust must be socially valuable. In either case, these benefits correlate much more readily with democratic systems than non-democratic systems. This is not to say that ASCs in non-democratic contexts should remain unregulated; after all, they facilitate the theft of natural resource wealth by corrupt dictators, which is detrimental to the aims of global development. Yet, a justification for the regulation of ASCs in non-democratic contexts cannot be grounded in a *pro tanto* trust-based reason – at least on the account of trust I defend here.

Ultimately though, I fulfilled the goals set out at the start of this project. I established that, at least in democracies, we have a *pro tanto* trust-based reason to implement a public ownership register that is sensitive to privacy considerations. Further, I offered a novel pluralist reconceptualization of the nature of trust *simpliciter* that I see as being able to better capture the diversity of trust relations in the public sphere. Taken together, these two points have significant implications for the philosophy of trust in particular, and the philosophy of public policy more generally. If my pluralist conception of trust is accepted, then this creates a new conceptual space for thinking about trust where we try to merge existing theories in novel and interesting ways in order to better understand the basis of the distinction between personal and impersonal trust. For instance, perhaps we could pair a morally neutral account of impersonal trust with a morally loaded account of personal trust in order to preserve the notion that trusting and being trustworthy is – at least in some sense – a virtue. Moreover, though I am

sceptical of this, there is scope for further research to examine whether it is possible for relations of personal trust to exist in the public sphere. While I explicitly stated that most trust relations in the public sphere would be impersonal, I left it open as to whether public trust could ever take on a personal flair. If, though it remains to be established by future research, public trust can be personal, we may be able to identify stronger, intrinsic arguments in favour of particular policy positions, as opposed to merely pro tanto arguments based on impersonal accounts of public trust.

To sum up, I argued that ASCs pose a significant threat to our global system: they facilitate corruption, money laundering, and all manner of illicit financial flows. Further, despite universal agreement concerning these negative externalities, ASCs remain largely unregulated as a result of an impasse in the policy sphere. Specifically, proponents of financial secrecy suggest we have a presumptive right to privacy, which would prevent the disclosure of public ownership information *unless* it can be shown that a given ASC is implicated in some wrongdoing, while opponents explicitly deny this premise. I argued that shifting the argument to the value of public trust can provide us with a way out of this impasse. In short, by showing that ASCs undermine public trust in pernicious ways, I argued that we have a *pro tanto* reason to regulate ASCs. Further, viewing the policy debate through the lens of public trust allows us to rigorously justify the comprise position of a suitably amended public ownership register on the basis of a theoretical reason grounded in trust, rather than merely on the basis of practicality.

Bibliography

Akerlof, George A., and Robert J. Shiller. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton: Princeton University Press, 2009.

Allingham, Michael G., and Agnar Sandmo. "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics* 1, no. 3/4 (November 1972).

Alm, James. "Measuring, Explaining, and Controlling Tax Evasion: Lessons from Theory, Experiments, and Field Studies." *International Tax and Public Finance* 19, no. 1 (February 2012): 54–77.

Baier, Annette. "Trust and Antitrust." *Ethics* 96, no. 2 (1986): 231–60.

Bartels, Larry M. *Unequal Democracy: The Political Economy of the New Gilded Age*. Second edition. New York: Russell Sage Foundation, 2016.

Bobek, Donna D., Amy M. Hageman, and Charles F. Kelliher. "Analyzing the Role of Social Norms in Tax Compliance Behavior." *Journal of Business Ethics* 115, no. 3 (July 2013): 451–68.

Brandeis, Louis D. *Other People's Money, and How the Bankers Use It*. New York : Stokes, 1914.

Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. *Explaining Norms*. Oxford University Press, 2013.

Brock, Gillian. "Global Justice." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, 2017. https://plato.stanford.edu/archives/spr2017/entries/justice-global/.

———. "Taxation and Global Justice: Closing the Gap between Theory and Practice." *Journal of Social Philosophy* 39, no. 2 (June 2008): 161–84.

Buckley v. Valeo, 424 U.S. 1 (January 29, 1976).

Bullough, Oliver. "How Britain Can Help You Get Away with Stealing Millions: A Five-Step Guide." *The Guardian*, July 5, 2019. https://www.theguardian.com/world/2019/jul/05/how-britain-can-help-you-get-away-with-stealing-millions-a-five-step-guide.

———. *Moneyland: Why Thieves & Crooks Now Rule the World & How to Take It Back*. London: Profile Books, 2018.

Chen, Kevin. *Political Alienation and Voting Turnout in the United States, 1960-1988*. San Francisco: Mellon Research University Press, 1992.

Christiano, Thomas. "Democracy." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University, 2018. https://plato.stanford.edu/archives/fall2018/entries/democracy/.

———. "Money in Politics." In *The Oxford Handbook of Political Philosophy*, edited by David Estlund. Oxford University Press, 2012.

Cialdini, Robert B., and Melanie R. Trost. "Social Influence: Social Norms, Conformity, and Compliance." In *The Handbook of Social Psychology*, edited by Daniel Todd Gilbert, Susan T. Fiske, and Gardner Lindzey, 4th ed. Boston: McGraw-Hill, 1998.

Citizens United v. Federal Election Comm'n, 558 U.S. 310 (January 21, 2010).

Confessore, Nicholas. “How to Hide $400 Million.” *The New York Times Magazine*, November 2016.

DeCew, Judith. “Privacy.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, 2018. https://plato.stanford.edu/archives/spr2018/entries/privacy/.

Dietsch, Peter. *Catching Capital: The Ethics of Tax Competition*. New York, NY: Oxford Univ. Press, 2015.

Does de Willebois, Emile van der, Emily M. Halter, Robert A. Harrison, Ji Won Park, and J.C. Sharman. *The Puppet Masters: How the Corrupt Use Legal Structures to Hide Stolen Assets and What to Do about It*. Washington, DC: World Bank, 2011.

Dworkin, Ronald M. “The Model of Rules.” *The University of Chicago Law Review* 35, no. 1 (1967): 14–46.

“Edelman Trust Barometer.” Edelman Trust, January 2019. https://www.edelman.com/trust-barometer.

Etizioni, Adam. “The Limits of Transparency.” In *Transparency, Society and Subjectivity: Critical Perspectives*, edited by Emmanuel Alloa and Dieter Thomä, 179–202. Cham: Palgrave Macmillan, 2018.

Faulkner, Paul. “Finding Trust in Government.” *Journal of Social Philosophy* 49, no. 4 (December 2018): 626–44.

———. “The Attitude of Trust Is Basic.” *Analysis* 75, no. 3 (July 2015): 424–29.

Faulkner, Paul, and Thomas Simpson, eds. *The Philosophy of Trust*. First edition. Oxford: Oxford University Press, 2017.

Ferrarello, Molli. “One Year after the Panama Papers Leak, Starting a Shell Corporation in the US May Be Easier than Getting a Library Card.” Brookings Institute, April 7, 2017. https://www.brookings.edu/blog/brookings-now/2017/04/07/one-year-after-the-panama-papers-leak-starting-a-shell-corporation-in-the-us-may-be-easier-than-getting-a-library-card/.

Findley, Michael G., Daniel L. Nielson, and J.C. Sharman. *Global Shell Games: Experiments in Transnational Relations, Crime, and Terrorism*. Cambridge Studies in International Relations. New York: Cambridge University Press, 2014.

Fischer, Brendan, and Blair Bowie. “Elections Confidential: How Shady Operators Used Sham Non-Profits and Fake Corporations to Funnel Mystery Money into the 2012 Elections.” US PIRG Education Fund, January 2013.

Global Witness. “Congo’s Secret Sales,” May 2014. https://www.globalwitness.org/en/campaigns/oil-gas-and-mining/congo-secret-sales/.

———. “Learning the Lessons from the UK’s Public Beneficial Ownership Register,” October 2017. https://www.globalwitness.org/en-gb/campaigns/corruption-and-money-laundering/learning-lessons-uks-public-beneficial-ownership-register/.

Guex, Sébastien. “The Origins of the Swiss Banking Secrecy Law and Its Repercussions for Swiss Federal Policy.” *The Business History Review* 74, no. 2 (2000): 237–66.

Hamilton, Alexander, James Madison, and John Jay. “Federalist No. 68.” In *The Federalist Papers*. New York: Palgrave Macmillan, 2009.

Hardimon, Michael O. “Role Obligations.” *Journal of Philosophy* 91, no. 7 (1994): 333–363.
Hardin, Russell. *Trust and Trustworthiness*. The Russell Sage Foundation Series on Trust, v. 4. New York: Russell Sage Foundation, 2002.
Harding, Matthew. “Responding to Trust.” *Ratio Juris* 24, no. 1 (March 2011): 75–87.
Hawley, Katherine. *Trust: A Very Short Introduction*. OUP Oxford, 2012.
———. “Trust, Distrust and Commitment.” *Noûs* 48, no. 1 (March 1, 2014): 1–20.
Henry, James. “Let’s Tax Anonymous Wealth!” In *Global Tax Fairness*, edited by Thomas Pogge and Krishen Mehta, First edition. Oxford: Oxford University Press, 2016.
———. “The Price of Offshore Revisited: New Estimates for ‘Missing’ Global Private Wealth, Income, Inequality, and Lost Taxes.” Tax Justice Network, July 2012.
Hobbes, Thomas. *Leviathan*. Reprint. Pelican Classics. Harmondsworth: Penguin Books, 1974.
Jones, Karen. “Trust and Terror.” In *Moral Psychology: Feminist Ethics and Social Theory*, edited by Peggy DesAutels and Margaret Urban Walker. Feminist Constructions. Rowman & Littlefield Publishers, 2004.
———. “Trusting Interpretations.” In *Trust: Analytic and Applied Perspectives*, edited by Pekka Mäkelä and Cynthia Townley, 15–28. Value Inquiry Book Series. Amsterdam: Rodopi Press, 2013.
Katz, Leo. *Why the Law Is So Perverse*. Chicago ; London: University of Chicago Press, 2011.
Korobow, Adam, Chris Johnson, and Robert Axtell. “An Agent–Based Model of Tax Compliance with Social Networks.” *National Tax Journal* 60, no. 3 (September 2007): 589–610.
Krishnamurthy, Meena. “(White) Tyranny and the Democratic Value of Distrust.” *The Monist* 98, no. 4 (October 2015): 391–406.
Lenard, Patti Tamara. “The Political Philosophy of Trust and Distrust in Democracies and Beyond.” *The Monist* 98, no. 4 (October 2015): 353–59.
———. “Trust Your Compatriots, but Count Your Change: The Roles of Trust, Mistrust and Distrust in Democracy.” *Political Studies* 56, no. 2 (June 2008): 312–32.
Lessig, Lawrence. *Republic, Lost: How Money Corrupts Congress - and a Plan to Stop It*. New York, NY: Twelve, 2011.
Lever, Annabelle. “Feminism, Democracy and the Right to Privacy.” *Minerva* 2005, no. 9 (2005): 1–31.
Levitsky, Steven, and Daniel Ziblatt. *How Democracies Die*. Kindle Book. First edition. New York: Crown, 2018.
Luhmann, Niklas. *Trust and Power*. Chichester: Wiley, 1979.
Mackenzie, Catriona, and Natalie Stoljar, eds. *Relational Autonomy: Feminist Perspectives on Automony, Agency, and the Social Self*. Oxford University Press, 2000.
MacKinnon, Catharine A. *Toward a Feminist Theory of the State*. Cambridge, MA.: Harvard University Press, 1989.

McLeod, Carolyn. “Trust.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2015. Metaphysics Research Lab, Stanford University, 2015. https://plato.stanford.edu/archives/fall2015/entriesrust/.

Michel, Casey. “The U.S. Is a Good Place for Bad People to Stash Their Money.” *The Atlantic*, July 13, 2017. https://www.theatlantic.com/business/archive/2017/07/us-anonymous-shell-companies/531996/.

Mullin, Amy. “Trust, Social Norms, and Motherhood.” *Journal of Social Philosophy* 36, no. 3 (September 2005): 316–30.

Nissenbaum, Helen. “Respecting Context to Protect Privacy: Why Meaning Matters.” *Science and Engineering Ethics* 24, no. 3 (June 2018): 831–52.

Obermaier, Frederik, and Bastian Obermayer. “Oligarchs Hide Billions in Shell Companies. Here’s How We Stop Them.” *The Guardian*, April 3, 2018. https://www.theguardian.com/world/commentisfree/2018/apr/03/public-registries-shell-panama-papers.

O’Neill, Onora. *A Question of Trust*. The BBC Reith Lectures. Cambridge: Cambridge University Press, 2002.

Ornstein, Norman J., Thomas E. Mann, and Michael J. Malbin. *Vital Statistics on Congress 2008*. Washington, DC: Brookings Institution Press, 2008.

Oshana, Marina. “Trust and Autonomous Agency.” *Res Philosophica* 91, no. 3 (2014): 431–447.

Pettit, Philip. “Republican Theory and Political Trust.” In *Trust and Governance*, edited by Valerie Braithwaite and Margaret Levi. Russell Sage Foundation, 1998.

Pilkington, Ed, and the Guardian US interactive team. “Scott Walker, the John Doe Files and How Corporate Cash Influences American Politics.” *The Guardian*. Accessed September 25, 2019. http://www.theguardian.com/us-news/ng-interactive/2016/sep/14/john-doe-files-scott-walker-corporate-cash-american-politics.

Putnam, Robert D. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster, 2000.

Quentin, David. “Risk-Mining the Public Exchequer.” *Journal of Tax Administration* 3, no. 2 (December 6, 2017). http://jota.website/index.php/JoTA/article/view/142.

Rachels, James. “Why Privacy Is Important.” *Philosophy and Public Affairs* 4, no. 4 (1975): 323–333.

Rawls, John. *A Theory of Justice*. Revised Edition. Cambridge, Mass: Belknap Press of Harvard University Press, 1999.

———. “Two Concepts of Rules.” *The Philosophical Review* 64, no. 1 (1955): 3–32.

Riddell, Roger. *Does Foreign Aid Really Work?* Paperback ed. Oxford: Oxford University Press, 2008.

Roemer, John E. “Ideology, Social Ethos, and the Financial Crisis.” *The Journal of Ethics* 16, no. 3 (September 2012): 273–303.

Rosenstone, Steven J, Keith Reeves, and John Mark Hansen. *Mobilization, Participation, and Democracy in America*. New York: Longman, 2003.

Savage, James D. “Corruption and Virtue at the Constitutional Convention.” *The Journal of Politics* 56, no. 1 (February 1, 1994): 174–86.
Schapiro, Tamar. “Compliance, Complicity, and the Nature of Nonideal Conditions.” *Journal of Philosophy* 100, no. 7 (2003): 329–355.
Shapiro, Scott J. *Legality*. Cambridge, Massachusetts: Harvard University Press, 2013.
———. “The ‘Hart-Dworkin’ Debate : A Short Guide for the Perplexed.” In *Ronald Dworkin*, edited by Arthur Ripstein, 22–49. Cambridge University Press, 2007.
Story, Louise, and Stephanie Saul. “Stream of Foreign Wealth Flows to Elite New York Real Estate.” *The New York Times*, February 7, 2015. https://www.nytimes.com/interactive/2019/admin/100000003460356.embedded.html?
Strawson, Peter F. “Freedom and Resentment.” *Proceedings of the British Academy* 48 (1962): 1–25.
Tavani, Herman T. “Philosophical Theories of Privacy: Implications for an Adequate Online Privacy Policy.” *Metaphilosophy* 38, no. 1 (January 2007): 1–22.
Tax Justice Network. “Tax Avoidance.” Accessed May 22, 2019. https://www.taxjustice.net/faq/tax-avoidance/, https://www.taxjustice.net/faq/tax-avoidance/.
Teachout, Zephyr. *Corruption in America: From Benjamin Franklin’s Snuff Box to Citizens United*. Cambridge, Massachusetts: Harvard University Press, 2014.
Townley, Cynthia, and Jay L. Garfield. “Public Trust.” In *Trust: Analytic and Applied Perspectives*, edited by Pekka Mäkelä and Cynthia Townley, 95–108. Value Inquiry Book Series. Amsterdam: Rodopi Press, 2013.
Uslaner, Eric M. “The Study of Trust.” In *The Oxford Handbook of Social and Political Trust*, edited by Eric M. Uslaner, Vol. 1. Oxford University Press, 2017.
———. “Trust and Corruption Revisited: How and Why Trust and Corruption Shape Each Other.” *Quality & Quantity* 47, no. 6 (October 2013): 3603–8.
“Voices from Iceland as PM Resigns.” *BBC News*, April 5, 2016, sec. Europe. https://www.bbc.com/news/world-europe-35967320.
Wilson-Chapman, Amy, and Douglas Dalby. “Panama Papers Helps Recover More than $1.2 Billion around the World.” *ICIJ* (blog), 2019. https://www.icij.org/investigations/panama-papers/panama-papers-helps-recover-more-than-1-2-billion-around-the-world/.
Wintour, Patrick. “MPs Attack Ministers over Delay to Tax Havens’ Public Registers.” *The Guardian*, January 11, 2019. https://www.theguardian.com/world/2019/jan/11/mps-attack-ministers-over-delay-to-tax-havens-public-registers.
Wolff, Jonathan. “Method in Philosophy and Public Policy : Applied Philosophy versus Engaged Philosophy.” In *The Routledge Handbook of Ethics and Public Policy*, edited by Anabelle Lever and Andrei Poama, 1st Edition., 13–24. London: Routledge, 2018.
Zucman, Gabriel. *The Hidden Wealth of Nations: The Scourge of Tax Havens*. Translated by Teresa Lavender Fagan. Chicago: The University of Chicago Press, 2015.