Classifying states: instrumental rhetoric or a compelling normative theory?

Mathew Coakley & Pietro Maffettone

Published online: 07 Jul 2017.

Submit your article to this journal

Article views: 1254

View related articles

View Crossmark data

Citing articles: 1 View citing articles

Routledge
Taylor & Francis Group

ORIGINAL ARTICLE

Classifying states: instrumental rhetoric or a compelling normative theory?

Mathew Coakley[a] and Pietro Maffettone[b]

[a]Politics and International Studies, Warwick University, Coventry, U.K; [b]Lecturer in Global Politics and Ethics, School of Government and International Affairs, Durham University, Durham, U.K

ABSTRACT
Many states use a classificatory approach to foreign policy: they put other states into particular categories and structure their engagement and relations partly as a result. There is one prominent modern international political theory – Rawls' Law of Peoples – that seems to adopt this approach as an account of justified state behaviour. But should we expect this type of theory ultimately to prove attractive, justified and philosophically distinct compared to more instrumentalist rivals? This paper explores the challenges generic to any such account, not merely those relating to Rawls' specific version, and surveys possible responses and their shortcomings.

Foreign policy and classifying states

Many modern states appear to rely on a 'classificatory approach' to their foreign policy making,[1] namely they identify other states as belonging to various categories – those respecting human rights, non-aggressors, dictatorships, democracies, allies, trade partners, 'friendly' autocracies, and many more[2] – and use membership of such categories to structure how they approach a very wide range of interstate interactions. But does it actually make sense – normatively and conceptually speaking – to have a foreign policy that is guided by such classifications?[3]

While the practice of classifying states as a way to structure foreign policy relations is reasonably widespread (Herrmann & Fischerkeller 1995), very few appraisals of the general approach exist in the global justice literature or in international ethics more broadly.[4] Furthermore, the few appraisals that do actually appear in the literature are centred on the question of how to operationalize state classificationism in a normatively credible way and have left the conceptual presuppositions of the approach largely

---

[2]For an excellent overview see (Hudson 2005).
[3]Readers familiar with the global justice literature will immediately relate this question to that of collective responsibility: see, for example, Miller (2007). Yet, while the two topics may be coterminous under some theories, they need not be. Members of a given political community need not necessarily be held responsible for how their state behaves internally and externally: the content of a foreign policy can be wholly directed at the governement of a state, or at its ruling elite or at different apsects of state institutions.
[4]On some of the issues related to operationalizing a specific version see Forster (2014) and Doyle (2006).

unaddressed. This paper aims to fill this theoretical gap by surveying the main challenges 'state classificationism' would face if it is to be a conceptually plausible and philosophically distinct theory of justified international engagement. To be very, very, very clear: the question being examined here is not one in empirical comparative politics or international relations – it is not the question of how states actually behave and whether we can model such behaviour using state-classifications. This a relevant concern – as discussed subsequently – because it gives us reason to think that state classificationism as an approach has appeal to key real-world actors. But the question we are asking is different in kind, and the huge literature looking at how states actually behave does not answer it. What we want to explore is whether the practice of classifying states can be theoretically justified as an attractive and distinct approach to international state-behaviour ethics without collapsing into well-known alternatives.

Furthermore, the value of discussing the approach's strengths and weaknesses in principle may also provide some clarity on the abuse of it in practice. Clearly existing states do classify other states for ideological reasons and for the pursuit of self-interest, this being done with categories that are sometimes unreflectively influenced by power imbalances and elite norms.[5] It would be helpful as such to know if the entire approach should be discounted or not, even if we recognize that it can be abused.

Within normative political philosophy, Rawls' (1999) *The Law of Peoples* (hereafter *LP*) potentially lends itself to be described as a form of state classificationism. How best to conceive of this work is an important issue, but to the extent possible we will try to be agnostic on this; thus the status of *LP* will vary depending on the reader's judgement. If they think it straightforwardly provides state or polity based conditions that determine a justified foreign policy, it will serve as a direct example. Alternatively, if the reader judges it as providing an ideal theory in which different kinds of political communities are conceptualized as elements of a constructivist approach to the justification of a conception of international justice, then it will serve as a referent, albeit one that may vary when applied. Furthermore, though we use the theory as an illustration, the overall argument doesn't depend on the very specific suggestions Rawls makes.

Thus one way to view the task of this paper is by the following thought experiment: imagine Rawls had proposed a different classification of political communities, one that nonetheless did not collapse into an instrumentalist, or some form of consequentialist, theory. Clearly, much of the criticism of his (counterfactual) proposal would be different to the criticism he did face. But some types of problems might be common to any reasonably compelling version of such a theory. This is our focus here.

The first part of this paper tries to produce a sympathetic reconstruction of the underlying assumptions that guide such classifications, notably relying on the idea that we can classify political communities given some of their 'defining features'. We then go on to show that, despite the use of such classifications appearing to be quite widespread, the approach *in general* suffers from a number of underlying structural conceptual difficulties, with the two most significant being: (a) completeness – that categories often comprise multiple criteria that need not correlate and often do not correlate, and as such do not necessarily yield threshold evaluations (that a state can be rightly regarded

---

[5]See here, for instance, Teti (2012) on post-democratization discourse, or Giannone (2010) on the ideological influences on the state rankings provided by Freedom House.

as a certain 'type' because it fits the features of that type); and (b) gradation – that adherence to categories is often by matter of degree not by the presence of a discrete feature.

We assess these objections and suggest that there are at least potential solutions or modifications that would blunt the aforementioned difficulties. The completeness problem can in principle be addressed by using some sort of lexical ordering of the moral principles underlying international relations, though the question of how to justify the precise ranking philosophically is a difficult one. And, in theory at least, if a state's divergent evaluations are based on its behaviour in different 'realms' then this might be addressed by having different spheres of engagement – that is, a polity might have a close relationship with another state in one sphere, and an antagonistic relationship with it in another. The gradation problem could in theory be overcome by replacing threshold conditions with judgements based on family resemblances or 'distance' to various ideal types.

We conclude, however, by noting that, while in principle addressing the difficulties, there are possible downsides concerning the reformulated approach's ability to be both theoretically compelling and realistically action guiding. The lexical solution notably generates a problem of 'mixed-incentives', namely that other states will face incentives to make small improvements in lexically decisive criteria even if causing a major worsening on some lower-ranked criteria – opening one's economy to international trade (say) may apparently 'excuse' widespread human rights violations or a crackdown on domestic dissent. The 'multiple realms' solution risks rendering foreign policy decidedly, and possibly unsustainably, 'schizophrenic' – being allies in some areas and strong antagonists in others is hard to maintain given the type of multi-area bargaining and engagement that characterizes state relations. Finally, the 'ideal types' approach to resolve the gradation problem could generate divergent assessments given its reliance on judgment rather than more clearly identifiable threshold evaluations. If the goal of the state classification approach is to act as a normative guide to foreign policy behaviour in the real world then the aforementioned problems would ultimately need to be addressed.

In short: while almost everyone engaged in it recognizes that accurately classifying states involves difficult empirical questions (is someone really respecting human rights, are their democratic procedures really robust?), the whole approach also brings with it difficult normative choices, despite its apparent widespread use and genuine appeal to a great many people.

The existing literature on this topic mainly falls into two categories: discussions of Rawls' *particular* proposals in *LP*; and examinations of the descriptive aspects of state classification as practiced by foreign policy actors. The goal of this paper is to draw upon but equally step back from these debates and ask the broader question: is state classificationism plausible as a distinct general normative approach? If so, what problems will any such account have to overcome, and how optimistic should we be that it will succeed?[6]

---

[6]Of the two authors, one is reasonably skeptical of the likelihood of success, the other cautiously optimistic. We both agree however that such judgements can only be justified if we have a clear sense of the approach's challenges and resources.

State classification: real-world practice

From the perspective of normative theory, discussions of the 'state categorization' approach to foreign policy making have gained prominence most significantly following Rawls's *LP* and several subsequent extended discussions and critiques (a very good overview is Williams 2010, and for some of the canonical criticisms see Pogge 1994; Buchanan 2000; Beitz 2000; Caney 2002; and Shue 2002; with Wenar 2006 and Reidy 2007 offering some replies).[7] Some will immediately resist the idea that *LP* can reasonably be considered as a form of state classificationism. They will point out that *LP* is a realistic utopia (1999, 11), to use Rawls' expression, that the eight principles of *LP* and the companion idea of well-ordered peoples are part of ideal theory (1999, 4–5), and that Rawls explicitly rejects the use of the term 'states' (1999, 23–30) precisely because he wants to avoid relying on agents that mirror the actual behaviour of political communities in the international system.

In line with what has been argued above, our reply is deflationary. First, let us assume that the aforementioned understanding of Rawls is the correct one. Even if so, the basic project of this paper is not affected. *LP* is a (putative) illustration of the idea of state classificationism – its cogency (as an illustration) does not settle the question of whether the approach is itself cogent (the focus of our discussion). Furthermore, even if *LP* is a pure ideal theory, it will need to be somehow applied, and that might necessitate some form of state classificationism (see notably Doyle 2006; Forster 2014).

In a similar way, note that some have criticized Rawls' *LP* for using an excessively 'rigid' theoretical framework (see for example Caney 2002, 104, but also; Shue 2002). A theory that wants to be 'Realistically Utopian' (1999, 5–7) and of use as a guide to the foreign policy of liberal democratic regimes (1999, 8, 83), cannot be grounded in categorical distinctions ill-befitting the complexity of international politics. Yet, such criticisms often seem to rely on the assumption that applying *LP*'s framework would necessarily fall prey to what we describe below as the gradation and completeness problems (Caney 2002, 105). Hence it matters whether these concerns can be defused by reconceptualizing the way in which state classifications are approached.

Second, whatever the 'right' interpretation of Rawls, the broader approach is worth taking seriously because of a perhaps less-appreciated fact: that this is plausibly how significant numbers of powerful state actors approach the problem.

Indeed, there has been a steadily expanding body of research, going back to at least the 1960s, whereby foreign policy scholars, in attempting to understand how states construct their foreign policy, have looked to the type of representations that leaders and elites (which are often considered the prime movers of policy formation within state bureaucracies) have of their own polity and of those with whom they interact on the international scene (Hudson 2007). This approach now has a lot of nuance to it, but the general idea, broadly stated, is that states form both self- and other-oriented role representations and these 'self-conceptions' or 'images' are then used to structure and orientate a state's interaction with other states.

Such classifications can be relatively rich – witness, for example, an early classic in this literature, Holsti's *National Role Conceptions in the Study of Foreign Policy* (1970),

---

[7]For more recent discussions of the implications of Rawls' work for international theory and practice see (Williams 2011; Riker 2014, 2009; Neufeld 2013; Reinhardt 2012; Forster 2014; Lister 2012; Eckert 2015).

which identifies close to 20 different main ways in which countries can perceive their role in international society, ranging from 'bastion of revolution', to 'defender of the faith'.[8] What matters for present purposes, however, is that this now quite large literature is premised on the extensive evidence that many state actors do seem to use the qualitative classification of other states (and their own) as a significant guide to their foreign policy.

So why not simply use the way influential state actors classify states in the real world as the 'correct' way of conducting state classifications? Well, for two reasons: one pragmatic, one more fundamental. First, influential actors have mixed incentives and so may not apply the categories consistently or in line with what their judgements might be absent self-interested considerations. Take, for instance, the U.S.'s use of 'rogue state', a term relied upon in structuring foreign relations by the White House, the Pentagon, and the Department of State of several different administrations (see Hoyt 2000; O'Reilly 2007).[9] It's then possible to use such designations to identify statistically the policy dimensions that typically trigger 'rogue status': (a) the attempt to develop weapons of mass destruction; (b) involvement with terrorism; (c) posing a military threat; and d) challenging international norms (Hoyt 2000, 303).

However, clearly this does not fully explain the actual classificatory choices made by the U.S. government: though both India and Pakistan, for instance, have developed nuclear programs in open violation of international norms, they are not officially considered as potential rogues (O'Reilly 2007, 302–4). Even if partly adopting the 'state classification' approach, the U.S. does so with a certain degree of inconsistency, to say the least. States may also simply be cynically using the categories as a justificatory veneer for pre-determined behaviour.

Second, however, and most importantly from a normative theorist's perspective, we want to know how these categories should be *justifiably* applied and structured, and what actions this should guide. The fact that many state actors adopt the broad approach seems to elevate it to the status of 'the sort of approach that can be implemented in the real world' and as such raises the question of this paper – is there a variant that has real-world utility *and* can still be normatively both distinct and justified. Empirical practice cannot settle the latter question.

State classification: normative requirements

Clearly, a classificatory-based foreign policy is attractive to many, and seems to be part of the way a large number of states approach foreign policy making. Yet for any theory endorsing this to represent a *distinct* normative account it has, crucially, not simply to collapse into the following, very well-theorized, view:

> *Foreign-Policy Reductivism* (F.P.R.): Any set of principles by which we classify states and use to guide foreign policy are all ultimately justified by reference to one underlying value or combined scalar metric.

---

[8]A more recent updated elaboration of the approach is Krotz (2002), notably where how states classify *themselves* is argued to be an important determinant of foreign policy behaviour.
[9]It's often wrongly believed that this is simply an outgrowth of the Bush Administration's 'war on terror' but, in fact, it predates it. See, for example, Miles (2013 ch. 1 and ch. 2).

This, of course, is plainly the case for utilitarians and for most consequentialists generally, though the approach in practice need not at all always require the naïve direct pursuit of a single goal, potentially involving the use of principles as rules-of-thumb with allied dispositions (Hooker 2000), as self-commitment devices (Elster 1979), and as coordination tools or epistemic shortcuts (Hardin 1988). F.P.R. also captures the position of a range of nationalists or those advocating national self-interest as the supreme test of policies, and for that matter those aimed at perpetuating stable rule by the current executive. This also could be a coherent position of a range of libertarian positions where the underlying value was liberty as some sort of non-interference or non-coercion (Lomasky 2001).

Ultimately though, when F.P.R. adopts a set of principles to classify states these hinge for their justification on the empirical question of whether adopting the principles promotes the underlying value (be it overall utility, welfare, peace, stable rule, liberty, or whatever). As such, while F.P.R. is clearly one way of classifying states according to some of their defining features, it doesn't exhaust the approach as a whole and, if it were the only justifiable type of theory, the classificatory approach would lose its normative distinctiveness.

What is it to be normatively distinct to reductionist approaches? Two theories are normatively distinct, in the sense used here, to the extent that they disagree on two dimensions: (i) the type of factors that determine ethical evaluations and (ii) the content of these factors. To give an example, a deontological theory will be significantly normatively distinct to a consequentialist one – disagreeing on both criteria – whereas consequentialist theories typically disagree on the latter only (with exceptions). Thus when it comes to foreign policy, to be normatively distinct is to make such judgements on grounds that don't ultimately collapse into an instrumental metric. To avoid this, it needs to adopt some variant of the following:

> *Foreign Policy Non-Reductivism* (F.P.N.R.): The principles by which we classify states and decide foreign policy thereof are not themselves reducible to one underlying value or combined scalar metric.

The question philosophically, however, is how to set out and justify such an approach *so that it doesn't ultimately collapse into a single-value or single-metric reductivist view*. Why does identifying normatively distinct approaches matter? Well, for two main reasons. First, because identifying distinct approaches helps us be clear on the sources of disagreement. For foreign policy reductivists there are two, and only two, possible sources of disagreement: (a) what the underlying value should be; and (b) the empirical question of how actions affect the pursuit of (a). For non-reductivists there are minimally three: (a) what the different non-collapsible underlying values should be; (b) how conflicts between these values should be resolved; and (c) the empirical question of how actions affect the pursuit of (a) and (b). If we can work out which approach is better justified we can gain, in the process, a better understanding of the potential sources of disagreement (and convergence) between possible theories.

Second, because F.P.N.R. captures, relatively straightforwardly, the way many reflective individuals talk and approach foreign policy, specifically capturing the intuition that certain principles cannot be 'traded'; that, say, economic growth, human rights protection, territorial sovereignty and democratic governance draw upon different ideals and cannot, ultimately, be simply justified by reference to welfare (or well-being, or whatever the underlying metric). Foreign Policy Non-Reductivism, unlike

its conceptual rival, could for instance justify the sense that there is something decid-
edly distasteful about condoning widespread human rights violations if doing so yields
a sufficiently higher probability of a growth-enhancing trade relationship.[10] Put in a
slightly different way, F.P.N.R. is in line with many people's considered convictions
about the non-tradability of some moral values, and thus it is important, for those
committed to achieving reflective equilibrium, to explore whether a form of state
classificationism based on F.P.N.R. is normatively and conceptually plausible.

Some may immediately object[11] that the negative appearance of F.P.R. might be reduced
by the underlying value being suitably noble – such as, for example, world peace. This is true,
but only to an extent. While it's certainly the case that how 'distasteful' a single-value theory is
will depend on the content of that value, there still seems to be something troubling about
seeing the world in that way. Caring only about world peace – and not also about, say, liberty
and human health – would seem to many as ignoring morally important interests (with this
depending crucially on how encompassing the value is). Furthermore, note that, often, these
'ecompassing values' are internally complex and that their specification tends to uncover
particular conditions that contribute to make their realization worthwhile. To illustrate, 'world
peace' is underdetermined: it can be modelled along the lines of a 'pax Romana' or according
to a Kantian 'feudum pacificum'. Not only would these two conceptions of world peace be
different, but what makes the second more attractive is precisely the fact that its attainment is
made conditional on other values (e.g. mutual recognition between political communities)
that are conceptually distinct and morally defensible on their own terms.

How then might F.P.N.R. be different in practice? The most plausible option here
appears to be to appeal to certain 'defining features' which express a series of values that
themselves are not reducible to a single metric. (Why is this the most plausible option?
Well, the demonstration is by counter-positive: assume that the values underlying these
distinct categories are the same or can be compared by certain weights – in that case we
have implicitly adopted a single metric representing the underlying weighting or under-
lying ideal. Hence we must have a theory-structure where this doesn't hold.)

And indeed, it's no coincidence that Rawls' *LP* – arguably the only fully spelled out
systematic (normative) theory of F.P.N.R. – does exactly this: it sets out the defining
features of various political structures and behaviors, and uses these to determine justified
international ethical behaviour.

Rawls uses the term 'peoples' to describe what he calls liberal and decent political
communities with three further images, namely, burdened societies, benevolent absolut-
isms and outlaw states, with these categories based on how these agents behave in various
'realms' (domestic political affairs, the amount of social and political capital, interstate
conflict, adherence to various international norms and behaviours).[12] What is relevant here
is that the classifications then partly determine liberal peoples' justified foreign policy
stance. Decent peoples should be tolerated as equal members in good standing of what

---

[10]And the distaste here is not really addressed by hoping all-good-things ultimately go together, that for instance
protesting human rights abuses will always lead to better growth-enhancing governance. It might; but it might not.
[11]We would like to thank one of the anonymous referees for encouraging us to address this particular objection.
[12]Some will object that we are forgetting the role of the *LP*'s eight principles: Aren't they what guide foreign policy
behaviour for Rawls? Of course they are relevant, but (i) the principles are 'selected' after Rawls has described the
main features of the parties (including their basic interests) to the first international orginal position (i.e. liberal
peoples) and (ii) the principles themselves specify constraints for membership in the Society of Peoples, but do not
provide a full account of foreign policy.

Rawls calls the Society of Peoples; burdened societies should be helped through a duty of assistance to become well-ordered; benevolent absolutisms – though mostly neglected in *LP* – are to be considered as sovereign over their people and territory given that they respect human rights and are externally peaceful; and, finally, outlaw states should be the object of coercive measures, ranging from sanctions to military intervention.

What Rawls can be thought to provide, then, is an approach to state classification that doesn't collapse into an instrumental metric[13] (such as, for example, consequentialist accounts), and isn't undermined by the observation that in practice states often classify other states in a much more messy and inconsistent manner than their official rhetoric claims. The approach in general is attractive, it seems to capture a standard that states are (imperfectly) aiming for, and arguably coheres well with many people's considered judgements. This does not mean, however, that F.P.N.R. lacks philosophical and empirical problems.

Problems: completeness and gradation

Rawls' approach to international ethics is a potential variant of F.P.N.R. and as such is useful for illustrating its broader problems, in particular those relating to the degree of rigidity in the classificatory criteria and to the real-world implications that such rigidity could potentially have.

First, note how Rawls seems to describe political communities through a set of necessary criteria. For example, a liberal people is such if: (i) it is organized on the model of a constitutional democracy; (ii) its citizens are united by 'common sympathies'; and (iii) it has a moral nature (1999, 23–4). As such, if any of the aforementioned conditions is not met, one should conclude that the object in question is not a liberal people. However, this would seem to create the apparent paradox that, for example, a political community institutionally organized as a constitutional democracy might not be considered a liberal people in the absence of, say, common sympathies shared by its citizens.[14] We are entitled to ask the broader question: what if a state does not meet one of the conditions that define it as a specific object (i.e. type of people), but meets all the other relevant conditions?[15] Sometimes, admittedly, it might be that failing one or some of the standards for membership of the relevant type can ipso facto push a people in a different group: for example, sticking to Rawls' categories for purposes of illustration, a benevolent absolutism[16] that is externally aggressive may simply become an outlaw state. However, this need not be the case for all relevant standards. Once again, consider a political community that only protects a subset of the basic liberal constitutional freedoms, or one that has a deeply liberal and secular value system but only an insufficient commitment to democratic voting mechanisms and procedures. How are we to classify it?

---

[13]There is a possible interpretation of *LP* that sees it as purely instrumentalist – such as to world peace. Neither author thinks this is the best interpretation, but if the reader accepts it then the references to *LP* throughout should be interpreted as referring to a hypothetical alternative where the justification was not meant as purely instrumentalist in the sense previously discussed.

[14]This is not as unlikely as might be initially thought. Many liberal democratic states in the real world contain more than one people (and thus, one might conjecture, more than one set of 'common sympathies' within the same institutional structure).

[15]Here we follow and expand on (Caney 2002, 104).

[16]Rawls defines benevolent absolutisms as political communities that respect the basic rights of their members, are externally peaceful, but fail to provide any meaningful form of political participation mechanisms to their citizens. See Rawls (1999, 90).

We can call this the *completeness problem*; namely, that even if we think of the relevant set of standards that comprise a state-classification as each being capable of being met or not, then it's still the case that a state might meet some and not others. This raises the pressing problem of how to classify a state when this happens.

As an example of the *completeness problem* consider China and the realm of human rights. Imagine, for the purposes of this example, that one is committed to a version of F.P.N.R. where respect for a clear set of basic human rights is central to determine the classification of a state. The Chinese state has what is widely regarded as a remarkably bad record on human rights (Foot 2000). If one were to focus on political rights, then China appears to have a sustained and well-thought through institutional apparatus the effect of which is to actively repress the political rights of its members. However, the latter picture is severely complicated by the fact that China's bureaucracy is also responsible for what is, according to many, potentially the largest successful exercise in lifting people out of desperate poverty in the history of humankind.[17] In so doing, the Chinese government can be credited with preventing a dramatically large amount of human suffering (and consequently credited for the partial fulfilment of a central class of basic human rights, namely, what Shue (1980) famously called subsistence rights).[18] Thus when it comes to human rights, how do we classify China? To classify it simply as an 'outlaw state' would be reductive of the complexity of the case in hand.

Now, one thought here might be that, in practice, states with poor behaviour on some human rights are likely to have poor behaviour on other such rights, so we might still have mostly complete classifications in this realm. This seems empirically plausible at least; but, as an observation, it understates the broader difficulties, for while many states do often have reasonably consistently good /bad /mediocre stances *within* a particular realm (e.g. domestic human rights, economic development and institutions, political structures, respect for borders, international cooperation, etc.) there is frequently really stark variance between realms. There are states that are peaceable and corrupt (Morocco),[19] those that are democratic yet seem to be constantly engaged in non-internationally authorized military action in other states (such as the U.S.A.), those that are highly economically developed but stifle dissent (such as Singapore),[20] and so on.

What these examples illustrate is that the completeness problem is likely to apply to a vast array of classificatory exercises based on widely accepted values such as external peacefulness, well-functioning internal institutions, and the representative character of government. More specifically, any criteria for classifying states will most plausibly cover different realms of activity and types of institution and, as a result, we should expect a failure of completeness to be very common. While conceptually this is still an issue of completeness, it deserves highlighting as the most empirically likely manifestation of this problem (call it, perhaps, the 'different realms case'); namely, if the criteria

---

[17]For an extended analysis of the complex and multifaceted reception of human rights norms in China see Peerenboorm (2005) and Svensson (2002).
[18]This example relies on accepting subsistence rights as key human rights, but the overall point can be applied to any such multiple-rights conception. In other words, nothing hinges on which rights one sees as human rights. Whatever the answer to that question it is possible to construct an alternative example that illustrates the same problem.
[19]For a useful international contextualization of Morocco's problems with corruption see https://www.transparency.org/cpi2014/results. Nothing hinges on accepting the validity of the classifications provided by transparency international.
[20]See, for instance, https://freedomhouse.org/report/freedom-world/2015/singapore. As with footnote 19, nothing in the argument hinges on accepting this specific illustrative case.

used to classify a state pertain to more than one realm (e.g. external peacefulness and respect for basic human rights at home), and, as is often the case, state behaviour varies very significantly between those realms (e.g. a state violates basic rights but is largely externally peaceful), then we will likely end up with a 'mixed' classification.

Second, as well as states often meeting some criteria and not others, it also seems usually to be the case in practice that most of the obvious criteria used to classify states can be met *to different degrees*.[21] It's not that there are simply some states which violate human rights and others which do not – many political communities violate some human rights, respect other human rights, and do a 'half-baked' job on yet more. The case of existing liberal democracies is particularly striking in this respect, at least if one adopts the Rawlsian framework in *LP*: given what Rawls considers to be a liberal people, the chances of an actual political community fully meeting these requirements are low, bordering the non-existent (see Doyle 2015).

Note, furthermore, that the issue here is not about the interpretation of the criteria and the difficulties don't merely stem from the uncertainty and perhaps indeterminacy inherent in general moral and political principles. It's true, for example, that different conceptions of human rights may individuate different central classes of such rights as a function of their differing philosophical groundings.[22] Yet, *even if* there was complete agreement on the content of the criteria, the problem would still stand insofar as meeting the criteria in question is often a matter of degrees. Call this the *gradation problem*; namely, that in a particular realm any standards or criteria may be met to a certain extent, not placing a state cleanly in one category or another.

As an example of the *gradation* problem consider the U.S.A. Here, for purposes of illustration, consider also a state classification approach in which the securing of a wide range of political and economic rights is central to being considered a liberal state. Other than exceptionally creative judgments, the U.S.A. is usually seen as a liberal democracy: most rights from the constitutional tradition are guaranteed, its citizens share at least some common sympathies, and the constitution requires that the political system should be organized on the model of a representative democracy. However, it is also fair to say that the U.S. political system does not guarantee the full array of basic subsistence rights of all its members,[23] while the role of capital in the political process does not seem to afford 'fair value' to the political liberties for all citizens (Bartels 2008). It secures some rights fully, some partially. How, then, do we classify the U.S.A.?[24]

---

[21]Once again we follow and expand on Caney (2002).

[22]Witness, say, the difference between traditional views such as those of Griffin (2008) and so-called political approaches such asthose of Beitz (2010).

[23]For a classic, and in many respects unsurpassed, analysis of the deficiencies of the American health care system, and the importanc of health provisions for basic justice, see Daniels (1985).

[24]Again, it's possible to argue with the contemporary case, but we can instead think of the U.S.A.'s history of civil rights progress: it seems plausible that there wasn't any single moment when it became a full liberal democracy but, rather, that there was a series of partial improvements on access to the franchise, protection from egregious race-based state violence, and a somewhat weakening of institutional discrimination. Even if one (controversially) thinks that the U.S. A. does currently fully meet all the relevant criteria, there was clearly a point in time where it met very few of them, and many points in-between where they were partially met. The same would presumably be true of the history of almost all countries currently described as liberal democracies (European states with histories of colonialism, for instance, would have been historically evaluated in a very mixed manner). Gradation, in other words, can be taken to be both a synchronic and/or diachronic problem, and even those confident that all modern developed democracies are *fully* liberal would have to concede that this was not always the case, and may not be the case for various states in transition in the future.

The completeness problem and gradation problem are importantly different. The completeness problem arises because if we have some set of thresholds to meet a classification in a particular realm (such as domestic affairs or trade or respect for international norms etc.) then it is possible that a state might meet some and not others, and it is likely that there will be quite a few states that meet the criteria in one realm and fail them in another (the different realms case). The gradation problem is that in the real world states often only meet a threshold to a certain degree or in a fractured manner. That is, schematically, if meeting an 'A' classification requires doing/being a1, a2 and a3:

(i) One state could, for example, meet a1 and a2 but not a3 (a problem of completeness); and

(ii) One state could meet a1 fully, a2 partially, and a3 weakly (a problem of gradation).

While these seem like abstract problems, they are not difficulties anyone sympathetic to the approach should happily leave unresolved, for one of the distinctive features of state classificationism is that it should be able to provide (indirectly, at the very least) *actual guidance* with respect to the attitudes and foreign policy engagement of a given state with respect to the international community and foreign polities. Rawls' work on international ethics, for example, promises to be a normatively salient guide to the foreign policy of existing liberal democratic peoples. However, if the latter promise is to be kept, any framework that relies on 'defining features', given the problems of completeness and gradation, will need to be adapted at least to some extent.

This action-guidance also helps highlight the scope of these challenges – for example, do they apply to all such classifications, such as that of Freedom House or of the Democracy Index? This depends on how the classifications are normatively used, because there are two different options here: one, as a way of summarizing some data; or, two, as an implied basis for deciding upon foreign policy. Officially both indexes are only the first, but they are sometime framed implicitly as the second, and so serve as a useful example. Strictly speaking, though, it is only classifications which then partly determine foreign policy actions that count as an example of the normative approach being discussed. Hence, the focus on foreign policy guidance.

One final matter of note. A reviewer has suggested that the extent of these difficulties should in principle depend on the number of variables or defining features. This seems to us an important observation and a potential practical argument for theoretical parsimony, since the theory here is meant to be also action guiding.

Some answers: counting principles and ideal types

Section four took the most obvious and straightforward theory structure of state classificationism – that where categories have a set of defining features, the structure is indeed that which Rawls seems to adopt – and raised a series of problems any such

theory seems likely to encounter. In what follows we discuss some possible changes that in principle could overcome the difficulties, and in the subsequent section we note a few of the downsides each of these brings.

The first sympathetic change is directed at the completeness problem, which, recall, is that objects of classification defined by a set of necessary conditions may only meet some of the relevant criteria and thus not fall into any one camp. However, note that this arises due to a specific way of conceiving of the definition of the categories. One could hold instead that states should be classified through some appropriately weighted good-making features with a partial lexical ranking (yielding quite a diverse set of sufficient conditions and ranking of different states within the category).

For example, we can take the following four parameters[25] for triggering rogue status and make the criteria of 'rogueness' ordered as follows, where the sign '>' indicates the level of importance of the criterion for ascribing the label: (a) the attempt to develop weapons of mass destruction > (b) involvement with international terrorism > (c) posing military threats > (d) challenging international norms. The lexical rankings would then allow us to be more confident in assessing and comparing cases where some, but not all, criteria are satisfied. For example, one could define a rogue state as a state that meets at least one of the first two criteria by order of importance ((a) and (b) in this case) or, alternatively, satisfies any three of them. In the same way, this approach would also provide for more easily discernible comparisons of 'rogueness' levels (e.g. country X adopting behaviour described by (a), (b), and (c) is 'roguer' than country Y adopting behaviour described by (b), (c), and (d)).

Needless to say, one would have to provide a justification for the precise lexical ordering, one that does not rely on some version of F.P.R. In other words, to retain the classificatory approach as normatively distinct, the ranking of the criteria should not be justified according to a single value or combined metric. The latter, however, is not insurmountable. While all versions of F.P.R. can potentially generate a lexical ordering of criteria, not all lexical orderings of a set of criteria need to rely on a single metric approach.

To illustrate, the ranking of the four features of 'rogueness' we have just suggested could be developed by citing different types of moral considerations such as the importance of peace and stability in the international system (so that weapons of mass destruction should not proliferate), the relevance of the principle of non-combatant immunity (which would disqualify support for terrorism), and the idea of the general obligation to follow valid legal norms (which would support respect for at least some international norms). In turn, their ranking could be developed by conducting pairwise comparisons between different ordering using Scanlon's reasonable rejectability contractualist methodology or, alternatively, by relying on the construction of a suitably defined evaluative standpoint (Scanlon 2000).[26]

---

[25]Here we follow Hoyt (2000) for the content of the parameters and their discussion.

[26]More broadly, and without delving too much into the suggested solution, the relevant point is that F.P.R.N. approaches seem to be able to avail themselves of broadly constructivist methodologies to produce a lexical ordering without necessarily having to accept some form of F.P.R. Rawls' original position is perhaps the most famous illustration of the claim that lexical orderings *do not require* the use of a single metric approach but can still be determined from a unified evaluative standpoint.

A second option for overcoming, or at least blunting, the completeness problem – and in particular its most prominent manifestation, the 'different realms case' – is to hold that we should simply classify states in area specific ways and *our foreign policy should be very different depending on our type of engagement*. We could, for example, maintain good economic links with those states that respect international treaties and yet seek to exclude them from a range of multinational bodies if they engage in domestic suppression. States as such would be evaluated in each 'realm' – economic structures, human rights, domestic political institutions and behaviour, respect for borders, and so on – and this would define our relations for a range of interactions germane to that realm. This is clearly an approach that has been adopted in several actual foreign policy scenarios, including, most notably, the relationship between Iran and some of its European commercial partners (at the time of writing – 2016/17).

This – as we note below – is not an option without problems, but it does at least seem to defuse some of the difficulties of completeness (and incidentally reflect how many states try to prevent 'trade missions' with non-liberal states from talking about human rights or border disputes).

The possibilities for overcoming the gradation problem are different. Recall that the gradation problem arose because it is extremely unlikely that any political community could realistically meet *in full* the defining features of a given category; more specifically, that their compliance with a particular requirement is often going to be a matter of degree rather than a 'yes/no' scenario.

Here again it is important to note that the gradation problem relies on a specific conceptualization of the nature of the defining features of a given category, namely being based on thresholds or conditions. Empirically, though, these seem likely to be often partially met, or met in fragmented ways.

However, a different but sympathetic approach would be to take the 'defining features' to represent that of an ideal type to which an instance can be closer or further away from that based on some sort of family resemblance.[27] The conceptual shift seems to relieve, at least to some degree, the need to view actual existing states as solely meeting the requirements or not.[28]

One reasonable objection here would be that considering state classifications as ideal types creates a problem of indeterminacy. The basic issue is that if a given category stands for an ideal type, and as such is not something for which we can simply take its defining features as conditions for normatively labelling real-world instances, then we still need to decide the extent to which an object is 'closer' or 'further' from the ideal type; that is, how well it has to match the relevant features that define a category in order to fall under it. To illustrate, let us assume the category 'liberal democracy' is an ideal type constituted by, say, the following two features: (i) respect for the basic freedoms that are familiar from the constitutional tradition; and (ii) a system of governance that embodies political equality by giving each adult citizen one vote, and by adequately protecting the fair value of the political liberties. The question, then, is to

---

[27]This seems to be the approach implictly used by (Herrmann and Fischerkeller 1995).

[28]Again, one further intrepretation of Rawls suggested by one of the reviewers was that this might be how *LP* is best conceived. If so, one could then start with the subsequent problems, and try in addition to provide an account of how the *LP*'s guidance would concretely work in real-world contexts.

what extent a country can be considered a liberal democracy given that we are prepared to see the latter category as an ideal type to which no country will perfectly conform.

Our answer is that, without rejecting the whole approach, such sorts of difficult assessments are probably unavoidable. There simply is no escape from the fact that real-world classifications according to ideal types are subject to some degree of judgment and sometimes will not be clear-cut. However, while unavoidable the problem is not utterly disabling. To say that a classificatory exercise is subject to some level of indeterminacy is not to say that it is groundless.[29] The point, rather, is to develop justificatory strategies in order to provide reasons for a specific type of classificatory choice. One strategy could be to start from those central cases in which we feel more confident about the classification. For example, for the ideal type 'liberal democracy' as defined above, one could start from specific examples – such as the Scandinavian countries – in order to determine the extent to which the relevant criteria are met by what we normally consider to be the best case(s) of real-world exemplification(s) of the ideal type. One could then move on to examine other countries that are standardly considered as liberal democracies according to ordinary and historically repeated political judgements, such as the U.K. and France, in order to compare them to the cases where we are more confident and establish if they are sufficiently similar. The method could then be reiterated by conducting pairwise comparisons between what are believed to be the central cases and those that are less clear-cut.

Note also that, so far, we have discussed the conceptual presuppositions of conducting classificatory exercises as a static and self-contained task. In other words, the analysis we have conducted relied on two assumptions: (1) that the objects of classification do not change as we are trying to classify them; and (2) that their features cannot be explained in terms of factors that are outside of their control. Both assumptions can be challenged. However, weakening them actually allows us to make progress on the gradation problem by providing a more sophisticated way of structuring our judgements.

As an example of a challenge to assumption (1), consider the role of intentions.[30] Thus far we have assumed that a state's 'intentions', very broadly understood, are irrelevant to how we classify them. The assumption is controversial. To illustrate, imagine that you are trying to classify a moderately autocratic country with a relatively decent human rights record.[31] Does it matter that such country has been (ex-hypothesis) reforming its constitutional structures in the direction of a much more participatory role for its citizens? Many would be inclined to think that it should. One way to find a place for the latter idea is to make clear that it can affect the way in which we structure our judgments in the context of the gradation problem. For example, in the context of pairwise comparisons, intentions may prove to be useful ways of assessing a 'direction of travel' that includes various judgements of what the state is likely to be like subsequently. This might notably also be

---

[29]While not in principle disabling from a particular agent's evaluative standpoint – indeterminacy is, after all, a feature of much judgement – it does risk conflict between different agents. On this see section six. The solution we are suggesting does not however collapse into the gradation problem because instances within a category can be further from or closer to the 'boundaries', with this being affected by partial or fractured compliance.

[30]For an excellent discussion of this issue see Forster (2014). She uses the idea of 'aspiration' to convey the importance of attitudinal features.

[31]For a real life example, consider Singapore.

something that impacts the foreign policy measures, in effect intentions indicating both a destination *and* a process that we may wish to buttress (or, if negative, to undermine).

Now consider one way to challenge assumption (2). States are placed in an environment (international society) that routinely affects their internal make-up and external behaviour. Some may thus be sceptical that classificatory exercises can have any meaningful point unless they also reflect the fact that membership in international regimes and organizations and the wider international context significantly affects how states behave both externally and internally. The latter claims are often seen as central to both liberal and institutionalist approaches to international relations (Moravcsik 1997; Legro and Moravcsik 1999). To illustrate, compare the proportion of a state's budget that goes towards defence with that allocated to the fulfilment of basic rights to subsistence. The relative weight of these expenditures may be affected by both membership in international organizations and the geopolitical environment in which a country is placed. For example, a state facing repeated challenges to its territorial integrity by its neighbours may commit more resources to its defence budget which, in the context of moderate scarcity, may affect its ability to protect subsistence rights. Should such external constraints affect the way in which we classify the state in question? Intuitively, many would be inclined to think that it should. What is relevant, for the purposes of our argument, is that, once again, challenging the assumption that we can classify states without looking at the context in which they operate is useful in identifying further sources of input for the gradation judgments. It signals that we should take into account (as one plausible source of input) the kinds of constraints that international society, broadly construed, places on its members.

To summarize: in section four we highlighted how the classificatory approach to foreign policy suffers from two main problems, that of completeness and of gradation. In section five we suggested that categories could be defined by counting principles which are partially lexically ordered, that they could be realm-specific, and that the overall classifications could be viewed as constituting ideal types closer to or further from which an instance can be. The use of counting principles to define categories avoids the rigidity of necessary conditions, and combining this with some lexical rankings justified through a broadly constructivist method ensures we are not (implicitly) collapsing everything to one underlying metric and thus adopting F.P.R. Making things realm-specific allows us to tailor our relations by area and thus to respond to states acting markedly differently in particular realms. Furthermore, conceiving of the defining features of a category as ideal types potentially allows us to make judgements of closeness or distance from the category and thus to respond to the fact that real-world states conform unevenly and partially to most qualitative standards.

New problems: judgment conflicts, incentives and 'schizophrenia'

The aforementioned solutions are, in our view, the obvious ways of making progress in order to make the state classification approach conceptually more defensible and, specifically, able to overcome what we have highlighted as its main weaknesses. Before concluding, however, it's worth noting a few of the difficulties raised by the very reconceptualization we have put forward. Such difficulties pertain to the effective action-guiding potential of the approach. We believe that any fully specified account of state classification would ultimately need to overcome (or deny) them. The three problems we want to highlight are: (a) the problem of conflicts of judgment; (b) the

problem of mixed incentives; and (c) the problem of 'schizophrenic' diplomacy. These problems are raised by, respectively, the use of ideal types, the use of lexical rankings, and the use of realm-specific classifications.

First, consider the problem of conflicts of judgment. One of the central moves in resolving the gradation problem was to conceive of the classifications as ideal types. Doing so partly relieved the criteria from having to be fully met in practice. However, as we noted above, there is likely to be more than one way to then perform the relevant exercise of deciding whether a specific state falls under a particular ideal type and they will, in all likelihood, significantly rely on judgment. In this picture, it seems probable that different state actors would be likely to end up with different sets of actual classifications. While the latter is not a decisive theoretical reason to abandon the use of ideal types, it at least puts pressure on its action-guiding potential. This is because, even assuming that all states accept exactly the same state classification ideal types (by itself a very optimistic assumption), merely implementing the approach could still generate large disagreements.

Second, consider the use of lexical rankings. One key factor that separates our somewhat reformulated state classification from being an instance of F.P.R. is that the criteria which define a classification can be given a lexical ranking that does not reflect an underlying single metric. However, providing a lexical ordering can also potentially generate a problem of mixed incentives. In what follows we use the example of what Rawls calls 'benevolent absolutisms' to illustrate concretely both how the lexical ranking approach could work and why it would generate this difficulty.

Recall that Rawls defines benevolent absolutisms as political communities that respect human rights and are externally peaceful (1999, 90). One way to operationalize the lexical ordering approach we have suggested above is to give priority (for the purposes of the classification) to a specific and easily observable subset of human rights and to a restricted understanding of external peacefulness. For instance, it could be that to be considered a benevolent absolutism one would at a minimum have to respect the borders of one's neighbours and not be engaged in the mass killing or enslavement of one's own people and the violent persecution of religious minorities. In this picture, the lexical ordering could be justified by appealing to the centrality of physical integrity for the basic interests of individuals, and by recalling the importance of territorial integrity for collective self-determination. In turn, a liberal state's foreign policy towards a political community that is classified as a benevolent absolutism could, say, envisage full trade and diplomatic relations and a degree of internal non-interference.

The downside of this sort of lexical approach, however, is that it creates the incentive to make progress on the lexically decisive criteria but to disregard those that affect where one lies within the category. For example, a state with a small border dispute and a relatively free press and civil society could as such gain full relations with liberal states were it to settle the border dispute and crack down on press freedoms, freedom of association and political dissent (assuming unchanged performance for what concerns the aforementioned basic human rights). The latter problem is that within the category of benevolent absolutisms as defined above there would, presumably, be wide variance in terms of respect for the full array of domestic human rights over and above the minimal idea that these rights preclude enslavement and mass murder. In the same way, there would likely be a great deal of practical and conceptual space between

respecting the territorial integrity of other states and a peaceful foreign policy. This creates the potential for problematic incentives in situations where minor gains with respect to a specific criterion can become more beneficial (for the would-be benevolent absolutism) than widespread large progress on lexically non-decisive criteria.

Furthermore, the latter problem becomes particularly striking if relations are different based on the realm. If interstate economic relations, for instance, are *only* based on adherence to a set of international economic norms and standards then there is very little self-interested incentive for state actors to make progress in other realms such as human rights, positive political reform or the protection of vulnerable minorities.[32]

Now perhaps this is all acceptable: we make the key criteria the ones that are most important. However, state classification does seem at a minimum here to need to provide some reassurance that non-lexically decisive criteria are not implicitly being jettisoned. The fact that under F.P.R. every major action of a state may affect how they are treated by foreign liberal states does seem to incentivize progress better than having certain criteria that count in reaching a threshold of treatment and others that – while judged very important – do not. This is the problem of lexical rankings and the stark incentives they can create.

Finally, the third difficulty accrues to the suggestion that states might be evaluated very differently in different realms. Someone could, as such, be a close economic ally, for example, and a security foe due to their support for various insurgent regimes and lack of respect for international law.[33] In principle we might have both a warm trade relationship and decidedly conflictual engagements in international bodies. In practice, however, this seems as though it would be sometimes too 'schizophrenic' to be sustainable: when we engage with states we usually do so simultaneously in a range of areas, and based on a relatively consistent evaluation of the state. To be foes in one area and allies in another is theoretically possible, but presumably practically difficult, and the mixed messages (this is a state we approve of vs. this is a state we regard as behaving very badly) might be hard to sustain.

Thus at a minimum any state-classificationism would need to address how one would maintain a consistent policy stance towards a foreign state if it was an ally in one realm and a foe in another (or some other strongly conflicting categories). States do do this – they can have strong economic relations even while engaged with each other in significant territorial or value-based disputes. But it does complicate the picture. It also, from the perspective of the non-liberal state, might produce quite a bit of antipathy even in the areas where there was strong cooperation: it is hard to be part of a government being strongly condemned by a foreign power and still to maintain real cooperation and trust in another area. One can fake it, but that undercuts some of the value of close relations.

The problems we have highlighted so far are not necessarily insuperable. What they do indicate, however, is that if we want the state classification approach to be not only normatively attractive but also able to provide actual guidance in the real world then further progress is required. It is one thing to describe some category of polity based on its defining features – as Rawls' does – and another to argue for how we should

---

[32]Note how under F.P.R., by contrast, these different areas are somewhat 'traded', so that external recognition and productive inter-state relations, in principle at least, can be improved by making progress in a range of areas. However, if there are lexical conditions in one area that bestow such relations then the other areas are de-incentivized.

[33]This type of situation is far from implausible. As noted above, it seems to be a feature of the relationship between European countries and Iran. For a specific illustration, think of the relationship between post-revolution Iran with Germany (see The Jerusalem Post, 28 June 2015). http://www.jpost.com/Middle-East/Analysis-Why-Germany-is-the-weakest-Western-link-in-nuclear-talks-with-Iran-407350.

structure our relations with those that conform to it. Yet to be of pragmatic import we need to know how to use this to relate to states that don't cleanly meet the classifications, and at the same time we need to do so and try to avoid perverse outcomes, overly 'schizophrenic' foreign policy relations and persistent evaluative disagreement.

Concluding remarks

State classificationism seems to be a widespread empirical practice, and yet both conceptually and as a type of normative theory it is oddly under-theorized. The aim of this paper has been to explore whether the approach in general is attractive and viable, not just whether the detailed and specific judgements underlying different views about how to classify states are the right ones.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

Bartels, L. M. 2008. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton, NJ: Princeton University Press.
Beitz, C. 2010. The Idea of Human Rights. Oxford: Oxford University Press.
Buchanan, A. B. 2000. "Rawls's Law of Peoples: Rules for a Vanished Westphalian World." *Ethics* 110 (4): 697–721. doi:10.1086/233370.
Caney, S. 2002. "Cosmopolitanism and the Law of Peoples." *Journal of Political Philosophy* 10 (1): 95–123. doi:10.1111/1467-9760.00144.
Daniels, N. 1985. *Just Health Care*. Cambridge: Cambridge University Press.
Doyle, M. W. 2006. "One World, Many Peoples: International Justice in John Rawls's the Law of Peoples." *Perspectives on Politics* 4 (1): 109–120. doi:10.1017/S1537592706060117.
Doyle, T. E. 2015. "When Liberal Peoples Turn into Outlaw States: John Rawls? *Law of Peoples* and Liberal Nuclearism." *Journal of International Political Theory* 11 (2): 257–273. doi:10.1177/1755088215571648.
Eckert, A. E. 2015. "States as Social Entities: Re-Examining the Assumption of Mutual Disinterest in Rawls' Law of Peoples." *Journal of International Political Theory* 11 (2): 1–15.
Elster, J. 1979. *Ulysses and the Sirens: Studies in Rationality and Irrationality*. Cambridge: Cambridge University Press.
Foot, R. 2000. *Rights beyond Borders: The Global Community and the Struggle for Human Rights in China*. Oxford: Oxford University Press.
Forster, A. 2014. *Peace, Justice and International Order: Decent Peace in John Rawls' the Law of Peoples*. London: Palgrave McMillan.
Giannone, D. 2010. "Political and Ideological Aspects in the Measurement of Democracy: The Freedom House Case." *Democratization* 17 (1): 68–97. doi:10.1080/13510340903453716.
Griffin, J. 2008. *On Human Rights*. Oxford: Oxford University Press.
Hardin, R. 1988. *Morality within the Limits of Reason*. Chicago, IL: University of Chicago Press.
Herrmann, R. K., and M. P. Fischerkeller. 1995. "Beyond the Enemy Image and Spiral Model: Cognitive?Strategic Research after the Cold War." *International Organization* 49 (3): 415–450. doi:10.1017/S0020818300033336.
Holsti, K. J. 1970. "National Role Conceptions in the Study of Foreign Policy." *International Studies Quarterly* 14 (3): 233–309. doi:10.2307/3013584.

Hooker, B. 2000. *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Oxford University Press.

Hoyt, P. D. 2000. "The 'Rogue State' Image in American Foreign Policy." *Global Society* 14 (2): 297–310. doi:10.1080/13600820050008494.

Hudson, V. M. 2005. "Foreign Policy Analysis: Actor-Specific Theory and the Ground of International Relations." *Foreign Policy Analysis* 1: 1–30. doi:10.1111/j.1743-8594.2005.00001.x.

Hudson, V. M. 2007. *Foreign Policy Analysis: Classic and Contemporary Theories*. Second ed. Lanham, MD: Rowan & Littlefield.

Krotz, U. 2002. "National Role Conceptions and Foreign Policies: France and Germany Compared." Program for the Study of Germany and Europe, Minda de Gunzburg Center for European Studies, Harvard University, Working Paper 02.1. doi:10.1044/1059-0889(2002/er01).

Legro, J. W., and A. Moravcsik. 1999. "Is Anybody Still a Realist?" *International Security* 24 (2): 5–55. doi:10.1162/016228899560130.

Lister, M. 2012. "There Is No Human Right to Democracy. but May We Promote It Anyway?" *Stanford Journal of International Law* 48: 257–277.

Lomasky, L. 2001. "Toward a Liberal Theory of National Boundaries." In *Boundaries and Justice: Diverse Ethical Perspectives*, edited by D. Miller and S. H. Hashmi, 55–78. Princeton, NJ: Princeton University Press.

Miles, A. 2013. *US Foreign Policy and the Rogue State Doctrine*. Oxford: Routledge.

Miller, D. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.

Moravcsik, A. 1997. "Taking Preferences Seriously: A Liberal Theory of International Politics." *International Organization* 51 (4): 513–553. doi:10.1162/002081897550447.

Neufeld, B. 2013. "Liberal Foreign Policy and the Ideal of Fair Social Cooperation." *Journal of Social Philosophy* 44 (3): 291–308. doi:10.1111/josp.2013.44.issue-3.

O'Reilly, K. P. 2007. "Perceiving Rogue States: The Use of the 'Rogue State' Concept by U.S. Foreign Policy Elites." *Foreign Policy Analysis* 3: 295–315. doi:10.1111/j.1743-8594.2007.00052.x.

Peerenboom, R. 2005. "Assessing Human Rights in China: Why the Double Standard?" *Cornell International Law Journal* 38: 71–172.

Pogge, T. 1994. "An Egalitarian Law of Peoples." *Philosophy and Public Affairs* 23 (3): 195–224. doi:10.1111/j.1088-4963.1994.tb00011.x.

Rawls, J. 1999. *The Law of Peoples: With the Idea of Public Reason Revisited*. Cambridge: Harvard University Press.

Reidy, D. 2007. "On Global Economic Justice: In Defense of Rawls." *Journal of Ethics* 11 (2): 193–236. doi:10.1007/s10892-005-7985-9.

Reinhardt, K. 2012. "No Migration in a Realistic Utopia? Rawls's the Law of Peoples and the Topic of Migration." In *Proceedings from the 49th Societas Ethica Annual Conference*, edited by G. Collste. http://www.ep.liu.se/ecp/097/015/ecp12097015.pdf.

Riker, W. 2009. "The Democratic Peace Is Not Democratic: On Behalf of Rawls' Decent Societies." *Political Studies* 57 (3): 617–638. doi:10.1111/j.1467-9248.2008.00739.x.

Riker, W. J. 2014. "Human Rights without Political Participation?" *Human Rights Review* 15 (4): 369–390. doi:10.1007/s12142-014-0341-3.

Scanlon, T. 2000. *What We Owe to Each Other*. Cambridge: Harvard University Press.

Shue, H. 1980. *Basic Rights: Subsistence, Affluence, and U.S. Foreign Policy*. Princeton, NJ: Princeton University Press.

Shue, H. 2002. "Rawls and the Outlaws." *Politics, Philosophy & Economics* 1 (3): 307–323. doi:10.1177/1470594X02001003002.

Svensson, M. 2002. *Debating Human Rights in China: A Conceptual and Political History*. Oxford: Rowman and Littlefield.

Teti, A. 2012. "Beyond Lies the Wub: The Challenges of (Post)Democratization." *Middle East Critique* 21 (1): 5–24. doi:10.1080/19436149.2012.658495.

Wenar, L. 2006. "Why Rawls Is Not a Cosmopolitan Egalitarian." In *Rawls' Law of Peoples: A Realistic Utopia?* edited by Martin and Reidy, 95–113. Oxford: Blackwell.

Williams, H. L. 2011. *On Rawls, Development and Global Justice: The Freedom of Peoples*. Basingstoke: Palgrave.