The Limits of Causal Inference

How the distinction between causal dependence and production highlights deficiencies in the potential outcomes framework and the essential role of descriptive inference in the study of causal processes

Tess Wise[1]

Prepared for VIM 2016

Introduction

How do our notions of causality connect with the tests we use to attribute causality? Drawing on a distinction proposed by Hall (2004) between causal dependence and causal production, I argue that when we expand our notion of causality to include both these aspects, the potential outcomes framework is revealed to be powerful yet incomplete. When causal production occurs without causal dependence and visa versa, counterfactual tests will lead us to incorrectly conclude that there is no casual relationship between connected phenomena, or that causal relationships exists between events connected by omissions. While relatively few cases in political science exhibit complete dependence without production or visa versa, I will argue that many causal questions are at least somewhat influenced by these phenomena. In short, Hall (2004)'s distinction between dependence and production forces us to narrow the number of cases for which counterfactual dependence is a complete test of causality and highlights the essential role of descriptive inference in the study of causal processes.

---

[1] PhD Candidate, Harvard University Department of Government. Any comments, suggestions or critiques may be directed to wise@fas.harvard.edu. PLEASE NOTE THAT THIS PAPER IS A FIRST ATTEMPT TO FORMALIZE THESE IDEAS AND IS VERY MUCH A WORK IN PROCESS. CITE AT YOUR OWN RISK.

The article proceeds as follows: First, I describe Hall (2004)'s distinction between causal dependence and causal production and explore political science examples in which causal dependence and causal production come apart. Second, I explain that production without dependence constitutes a SUTVA violation; highlighting that SUTVA violations are more widespread than we generally think. The solution, I argue, is to leverage descriptive inference which, when combined with theory, allows us to assess causality outside of the potential outcomes framework. Third, I argue that in the case of causal dependence without production, descriptive inference is again crucial if we want to make useful policy recommendations.

A. Two Concepts of Causation

Hall (2004) proposes that causal relationships come in at least two fundamentally different varieties: dependence and production. Causal dependence is likely familiar to most social scientists as it is synonymous with counterfactual dependence. In general, we say that *Y* counterfactually depends on *X,* if *Y* would not have occurred without *X.* In statistical applications, we formalize counterfactual causal identification using the potential outcomes framework (sometimes referred to as the Rubin Causal Model, see Holland 1986).[2] In the potential outcomes framework, each unit has multiple potential outcomes, but only one actual outcome that depends on the treatment received. In order to use the potential outcomes framework for causal identification, certain assumptions must hold. First, we must assume that treatment status is independent of the potential

[2] Note that this is not the only way to formalize test of causality in statistics. Dawid (2000) uses a decision theoretic model and Pearl (2009) uses a structural equations framework. In political methodology, however, the potential outcomes framework is increasingly the predominant way to formalize tests of causal identification (Keele 2015).

outcomes. Second, the stable unit treatment value assumption (SUTVA) must hold (Rubin 1986). We will return to the second of these assumptions later as cases in which we see production without dependence and visa versa are often cases in which SUTVA is violated.

Unlike causal dependence, causal production is likely less familiar to social scientists. According to Hall (2004), causal production occurs when an event helps to "*generate* or *bring about* or *produce* another event" (Hall 2004, 1). In most cases production and dependence occur together, but sometimes they can come apart. While the potential outcomes framework allows us to test for causal dependence, tests of counterfactual dependence fail to attribute causality in the case of causal production. Familiarizing ourselves with cases in which causal dependence and causal production diverge can help us get a feel for the production aspect of causality that we usually ignore and observe why this divergence signals the limits of causal inference—at least as it is currently practiced in post quantitative applications.

Building up from toy examples to political science applications, I now consider examples of dependence without production and production without dependence.

A.1. Dependence without production

A.1.1 Toy Example: Banana peels cause broken windows

Suzy (a troublemaker!) likes to throw rocks at windows. Her friend Billy has been instructed to stop her from breaking any more windows. As Suzy grabs a rock and prepares to throw it at the nearest window Billy runs towards her in an attempt to stop her, but slips on a banana peel. Suzy throws her rock, breaking the window. Had Billy not slipped on the banana peel, he would have prevented Suzy from throwing her rock and thus prevented the window from breaking.

Hall (2004) calls this an example of "double prevention" because one event (Billy tripping on the banana peel) prevents another event (Billy stopping Suzy) that, had it occurred, would have prevented yet another event (Suzy breaking the window). Hall (2004) argues that the relationship between the banana peel and the state of the window is one of causal dependence without causal production.

The state of the window *causally depends* on the banana peel because had the banana peel not been there, Billy would have stopped Suzy and the window would still be intact. However, it is not the banana peel, nor Billy's subsequent failure to stop Suzy, that *produces* the broken window—it's Suzy's throw. Moreover, it would be strange to only blame Billy or the banana peel for the window's breaking.

To be clear, the banana peel does matter—it matters crucially! A counterfactual analysis of the effect of the banana peel on the window's integrity gives us useful information, but if we say the banana peel caused the window to break (and leave it at that) we would be technically correct (because we can show causal dependence), but we might be accused of missing something.

A.1.2 Consequences of dependence without production

Does it matter if we attribute blanket causality to the relationship between the banana peel and the window in the toy example? After all, there *is* a causal relationship (counterfactual dependence) at work. But there are drawbacks to attributing blanket causality in cases of dependence without production. For one thing, considering only the relationship of causal dependence (the banana peel breaking the window) might lead us to miss the potentially more important production relationship (Suzy breaking the window). This could matter greatly when it comes to the policy recommendations we might make. Armed with the "fact" that banana peels cause windows to break, we might recommend a banana peel removal campaign to keep neighborhood windows in tact. But even if the banana removal campaign works, we would still be missing the fact that it's actually Suzy who is the problem. When Suzy visits a neighboring town (without Billy to stop her) and the locals complain that all their windows are being broken, the banana peel removal campaign is unlikely to appease anyone.

A.1.3 Example of dependence without production: Democratic Peace

Imagine an experiment in which the assigned treatment doesn't directly produce an observed outcome but instead blocks a parallel causal process which, had it not been blocked, would have prevented the observed outcome from occurring. When this occurs, counterfactual analysis will lead us to attribute causal dependence to the relationship between the treatment and the outcome even though the treatment does not actually *produce* the outcome. When this happens, we may find a strong empirical regularity, but have trouble providing a compelling causal explanation for our findings. An example that fits this scenario is democratic peace theory.

Democratic peace theory claims that democracies rarely fight one another because they share common norms of live-and-let-live and domestic institutions that constrain the recourse to war (Rostano 2003). Democratic peace theorists have shown that the correlation between democracy and peace is very robust (Maoz 1998; Bueno de Mesquita et al 1999; Oneal and Russet 1999; Ray 1995; Russet 1993; Weart 1998), but other scholars argue that the finding is correlation, not causation (Faber and Gowa 1997;

Gartzke 1998; Layne 1994; Rostano 2003). If we expand our notion of causation to include production as well as dependence, we find that both sets of scholars may be correct. It is possible that democracy does not *produce* peace, but rather blocks other causal paths, which had they not been blocked, would be more likely to lead to war. For example, the advent of democracy may mean that authoritarian or totalitarian regimes don't prevent the development of constraining domestic institutions and thus peace prevails. Perhaps democracy doesn't cause the development of these institutions, but rather it sets the country on a path that avoids other forces that would block the development of these institutions. If this is true, democracy is like the banana peel in the toy example—it is part of the explanation for the observed peace, but it does not directly *produce* the peace.

This interpretation seems to square with Rostano (2003)'s assessment of the causal mechanisms through which democracy theoretically *produces* peace. Rostano (2003) finds that none of the causal logics proposed by democratic peace theorists operate as stipulated by their proponents. While Rostano (2003) claims that the lack of evidence for the proposed causal mechanisms undermines the argument that democracy causes peace, I would amend his conclusion to state that the lack of evidence for the proposed causal mechanisms undermines the argument that democracy *produces* peace.

Here the distinction between production and dependence not only helps to reconcile the contradictory findings in the democratic peace literature but it also points us in the direction of a different research agenda. First, we might imagine that even if we could (somehow) run an experiment to test the counterfactual dependence between democracy and peace, this should not necessarily assure us that democracy *produces* peace. Second, if democracy is like the banana peel, we may have more luck understanding the nature of causal relationship between democracy and peace by considering the parallel causal pathways that democracy *blocks* as well as the ones it *opens*.

A.1.4 Another example of dependence without production: Education and voter turnout

Another case in which dependence without production may be at play is in the relationship between education and voter turnout. Analyzing two randomized experiments and one quasi-experiment, Sondheimer and Green (2010) find compelling evidence that voter turnout is *causally dependent* on education.

While causal dependence is clear, *how* education produces higher voter turnout is a little less clear. As Sondheimer and Green (2010) themselves note, education could produce higher voter turnout through many different paths. Education may impart skills that allow voters to better negotiate the bureaucratic barriers to voting such as registration (Wolfinger and Rosenstone 1980), increase general interest in and knowledge of politics (Delli Carpini and Keeter 1996; Hyman, Wright, and Reed 1975), or expand one's social network and thus political capital (Rolfe 2004)---And "[t]hese three explanations by no means exhaust the list of potential causal pathways connecting educational attainment to voting" (Sondheimer and Green 2010, 186).

Likely many of these explanations are correct, however, the existence of relationships of dependence without production means that there are many paths between education and voter turnout we are not considering because they have to do with omissions or double preventions. Consider, for example, that adolescents who stay in school are less likely to be involved in crime and therefore less likely to lose the right to vote because they commit a felony (which leads to disenfranchisement in many American states). In this case education causes voter turnout by *preventing* a separate process that, had it occurred, would have impeded voter turnout. This is still an example of causal dependence, and if this pathway were active, we would find a causal effect of education on voter turnout. However, if education were primarily causing higher voter turnout through these kinds of pathways, the *policy* recommendations we would provide are less likely to be effective.

Of course, our explanations are generally less effective when we do not explore the mechanisms connecting our explanatory and outcome variables (CITE yes, but what's the mechanism). When causal dependence occurs without causal production it exacerbates this problem because it is harder to trace processes triggered by the absence of an event (as compared to the presence of an event). Moreover, because omissions are much less likely to be manipulable, we cannot study these processes using existing quantitative methods for studying causal mediation (Imai et al 2013). In fact, if we maintain, as some statisticians do, that manipulation is necessary for causation, then we would not consider these relationship causal at all. I would argue, however, following Rubin (1986) that "no causation without manipulation" is a *motto* designed to guide us toward using the potential outcomes framework under appropriate conditions (when the counterfactuals are well defined) not a hard and fast rule.

This is not to say that studying these kinds of negative events is not worth our time. Many important studies in political science seek to answer questions about omissions. Hochschild's *What's Fair? American Beliefs about Distributive Justice* (1986) asks the question of why there is no socialism in American and Schlozman and Verba's *Injury to Insult* (1979) asks the question of why the unemployed don't demand more benefits. The reader might note, however, that most works in this vein do not rely on solely on quantitative methods for causal inference and this is no accident (a point to which I will return later).

A.1.5 Dependence without production summary

In sum, we see relationships of dependence without production when an independent variable is causally dependent on an outcome, but is connected to the outcome because it *prevents* a parallel process from taking place, which, had it unfolded, would have prevented the observed outcome. When this occurs, focusing solely on the causal dependence obscures the role of causal production and is more likely to result in policy recommendations that miss important causes.

A.2. Production without dependence

A.2.1 Toy Example: The broken bottle that doesn't depend on the bottle breaker

Suzy and Billy are throwing rocks at a bottle, trying to break it. Suzy throws her rock a split second before Billy throws his. Suzy's rock hits the bottle, breaking it, then Billy's rock whizzes through the air in the space where the bottle was. When asked, "What caused the bottle to break?" most of us would answer, "Suzy's rock." However, the breaking of the bottle does not *counterfactually depend* on Suzy's rock because had Suzy missed, Billy's rock would have broken the bottle. Causal production accounts for the fact that Suzy's throw *produced* the bottle's breaking even if it doesn't *depend* on it. The fact that the bottle can only break once pre-empts the causal dependence that would have occurred between Billy's throw and the state of the bottle had it not already been broken by Suzy and thus example is referred to "late pre-emption."

To illustrate how counterfactual dependence fails to correctly attribute causality in this case, imagine running an experiment in which Suzy's presence is randomly assigned. Assuming both Billy and Suzy have perfect aim, comparing the state of the bottle under "treatment" (when Suzy is there) to the state of the bottle under "control" (when Suzy isn't) will lead us to conclude that Suzy's presence does *not* have a causal effect on the bottle's state.

A.2.2 Consequences of production without dependence

Does it matter that we fail to attribute causality to Suzy's throw even though it broke the bottle? Again, if all we care about is whether or not the bottle is broken, perhaps it's fine that we don't give Suzy causal credit since the bottle's breaking is overdetermined and it will end up broken either way. However, if we go on to assume that Suzy doesn't break bottles (or windows, for that matter!) we'd likely find ourselves eating our words. The lack of a causally dependent relationship shouldn't be taken as a lack of any causal relationship. When we don't find a causal effect we're expecting (or even not expecting), it may be due to a backup process that causes the outcome to be overdetermined, not the absence of any relationship between our treatment and the outcome.

A.2.3 Example of production without dependence: Opposition to immigration

In general terms, when might we see production without dependence? Say we run an experiment and find no statistically significant difference in the outcome between the treatment and the control conditions. In certain cases of overdetermination, the treatment may still be *producing* the outcome even though the outcome does not *counterfactually depend* on the treatment.

Consider the well-known study by Hainmueller and Hiscox (2010) on attitudes towards immigration. Using a survey experiment, Hainmueller and Hiscox (2010) randomly assign respondents one of the following questions: Do you agree or disagree that the US should allow more **highly skilled/low-skilled** immigrants from other countries to come and live here? [emphasis added] Respondents answer on a scale from "Strongly disagree" (1) to "Strongly Agree" (4).

Explanations that emphasize economic self-interest argue that natives should think about their own economic position when considering whether or not to allow more immigration (Kessler 2001; Mayda 2006; Scheve and Slaughter 2001). According to these theories, more highly skilled Americans should prefer low-skilled immigration (to prevent competition with their own jobs) while low-skilled Americans should prefer highly skilled immigration for the same reason. Contrary to the predictions of these theories, Hainmueller and Hiscox (2010) do not find systematic variation in the premium attached to highly skilled immigrants across respondents' skill level. Rather, Hainmueller and Hiscox (2010) find that both low-skilled and highly skilled natives prefer highly skilled immigration over low-skilled immigration, and this preference is not decreasing in natives' skill levels.

Is it possible that economic self-interest *produces* the premium placed on highly skilled immigration even though this premium is not *causally dependent* on natives' skill levels? If attitudes towards immigration are overdetermined there could be causal production without counterfactual dependence.

Imagine that economic self-interest works just as its proponents claim—leading to opposition to low-skilled immigration and support for highly skilled immigration among low-skilled natives and the opposite among highly skilled natives. If economic self-interest were the only process in play we would expect to see opposition to immigration that matches the respondent's skill level and support for immigration that's less economically competitive.

Now imagine that backup processes occur that have *differential effects* on high skill vs. low skill natives such that for low-skill natives these processes *reinforce* the effects of economic self-interest while for high-skill natives these processes *swamp* some of the effects of self-interest relationship. In the case of low-skill natives, economic self-interest is like Suzy's throw that breaks the bottle even though without it, we would still see the same result due to the backup processes.

This is not to say that Hainmueller and Hiscox (2010) are incorrect to believe that self-interest has been overstated as an explanation for opposition to immigration. The experimental results, however, only provide evidence that opposition to immigration is not causally *dependent* on economic self-interest. Economic self-interest could still be *producing* opposition to immigration, especially among low-skill natives.

In fact, in Hainmueller and Hiscox (2010)'s review of explanation for their results (that both highly skilled and low-skill natives prefer highly skilled immigration) they find that in states with high fiscal exposure (high levels of welfare spending and a high ratio of immigrant to native households) there is more opposition to immigration among poor natives, suggesting that concerns about access to or overcrowding of public services contribute to anti-immigrant attitudes among poorer native citizens. This suggests that for poor natives, economic self-interest could well be *producing* opposition to low-skill immigration, even if the relationship between native skill-level and opposition to immigration does not *depend* on economic self-interest because this opposition is over determined.

A.2.4 Production without dependence summary

Production without dependence occurs when an outcome is overdetermined and can be achieved through multiple paths. Even though one path may be activated in a particular case, the existence of redundant pathways means that the outcome does not causally depend on its producers. We cannot identify these causal effects using the potential outcomes framework, because, as I will explain now, production without dependence violates SUTVA—one of the key assumptions we must make to use the potential outcomes framework for causal identification.

B. Production without dependence violates SUTVA

SUTVA is one of the key assumptions needed to use the potential outcomes framework for causal identification of counterfactual dependence. SUTVA is the a priori assumption that the value of the outcome for a particular unit receiving a particular treatment will be the same no matter what mechanism is used to assign treatment and no matter what treatments the other units receive (Rubin 1986). Sometimes this assumption is called consistency (Keele 2015). SUTVA is actually an incredibly strong assumption. SUTVA is violated when there are multiple versions of the treatment, when there is interference or spillover between units, and when the potential outcome under treatment depends on whether the unit received the treatment at a particular time. The overdetermination I described in the previous section can be thought of as multiple versions of the treatment and therefore violates SUTVA. By pursuing a traditional analysis of causal dependence Hainmueller and Hiscox (2010) have implicitly assumed that SUTVA holds and that this kind of scenario is not occurring.

Many other forms of production without dependence would be easy to recognize as SUTVA violations. Consider, for example, a scenario in which Suzy pushes Billy while she is throwing her rock, causing Billy to miss (interference). Without Suzy's presence, Billy would have hit the bottle and thus the bottle's state does not *causally depend* on Suzy's throw.

Knowing that we cannot adequately capture causal production using tests of causal dependence means that SUTVA violations are far more common and widespread than we usually consider. Despite this, there is very little guidance on what to do when SUTVA does not hold. In the case of spillover effects, the suggestion is to model the spillover and there is increasingly a literature on how to do this (Sinclair 2012). Yet in the case of heterogeneous treatment effects the ways to proceed are less clear. Here I will make an informal case that when there are heterogeneous treatment effects, the way forward is descriptive inference.

B.1 Leveraging descriptive inference to study causality when SUTVA is violated

One approach to studying heterogeneous treatment effects is to allow for a multi-valued treatment variable. Doing this, of course, increases the number of potential outcomes enormously and we must find data to serve as a proxy for the unobserved counterfactual states that we need to estimate to calculate our causal effects. When dealing with more complex causal relations quantitatively we are forced to find larger and larger samples and spread our data more and more thinly.

The problem is compounded when studying causal processes such as mediation in which there are multi-part counterfactuals that require even very strong (unit-level!) assumptions such as sequential ignorability and consistency in order to be estimated (Imai et al 2011, 2013). Elsewhere, I have argued that these assumptions are sufficiently implausible as to render most mediation analysis hopeful at best (Wise 2014).

Rather than stretch the potential outcomes framework to its limits, where it is likely on shaky ground, we could accept that more complicated causal accounts in which there is production without dependence (SUTVA violations) *inevitably* place the researcher in a place where the potential outcomes framework does not apply. Luckily, descriptive inference is still entirely valid in these situations and can add leverage that is missing in purely causal analyses.

Quantitative tools for descriptive inference mostly include modeling of various varieties—experimental results are often lauded because they allow us to nonparametically identify causal effects (Keele 2015), but this nonparametic identification is only possible under limited conditions (when SUTVA isn't violated). Instead of encouraging researchers to remove models (and modeling assumptions) from their research, we should encourage them to include well-theorized models that allow them to study causal situations that are not captured by the potential outcomes framework.

Another option is to use qualitative tools such as process tracing for descriptive inference (Collier 2011). The descriptive component of process tracing involves taking "good snapshots" at a series of specific moments so that the key steps in the process can be characterized (Collier 2011, Mahoney 2010). To see how crucial these snapshots are in cases of production without dependence, think back to the toy example in which Billy and Suzy are breaking bottles. It is a clear the *description* of what happens at the snapshot in time when Suzy's rock hits the bottle first, then Billy's whizzes through the air right where the bottle was is crucial to understanding the nature of causality in this case. Without a good description of this moment we are stuck with a messy overdetermined situation in which two rocks are thrown, a bottle is broken and we are not sure who caused what.

C. Leveraging descriptive inference when we have dependence without production

The role of descriptive inference is also crucial in cases of dependence without production. When we have a good description, it allows us to ask more focused causal questions. For example, thinking back to our toy example involving the banana peel, we might ask: Why did the window break, *given that Billy had been tasked with stopping Suzy from breaking windows?* In this case, answering: "Billy tripped on a banana peel" makes complete sense as a *description* of a causal event. Because our question is more precise, our situation of dependence without production is less problematic.

Moreover, when we have dependence without production more detailed *descriptions* of the counterfactual states can help us figure out which aspects of the counterfactual world differ and can point us towards the production relationships that might be obscured by our initial analysis. For example, if we simply compare the window's state with and without the banana peel, we find that it's shattering depends on the banana peel. Yet if we *describe* the process by which the window broke when the banana peel was present and failed to break when the banana peel was not present, we learn that that banana peel matters because it allowed a production relationship (Suzy's throw) to proceed unimpeded. Good description can help us avoid recommending banana removal to prevent windows from breaking.

Conclusion

Using Hall (2004)'s distinction between production and dependence helps us explore the limits of causal inference using the potential outcomes framework and highlights the essential role of descriptive inference in the study of causal processes. In cases in which we have dependence without production, we can identify causal effects using tests of counterfactual dependence, but these causal effects will not give us insights into the production relationships that are often at the heart of the causal relationship. We need insight into these production relationships in order to make good policy recommendations. Luckily, these relationships that are impossible to nail down with causal dependence are often easy to describe, therefore scholars should use the tools of descriptive inference to explore these relationships, keeping in mind that it causal dependencies are just as likely to emerge from *blocked* paths as they are from opened ones.

In the case of production without dependence, our tests of counterfactual inference fail to identify the salient causal relationship because SUTVA is violated. Rather than give up on studying these cases, we must use other tools to explore causal relationships that are production based. In this case theory, models, and descriptive inference are our best tools. It is also likely that these cases will require extensive qualitative research because they often involve complex treatments that inevitably place researchers in a small-N situation that is most amenable to qualitative tools such as process tracing.

REFERENCES

Bueno de Mesquita, Bruce, et al. 1999. "An Institutional Explanation of the Democratic Peace." American Political Science Review 93 (December): 791–812.

Collier, D., 2011. Understanding process tracing. *PS Political science and politics*, *44*(4), p.823.

Dawid, A. Philip. "Causal inference without counterfactuals." *Journal of the American Statistical Association* 95, no. 450 (2000): 407-424.

Farber, Henry S., and Joanne Gowa, "Common Interests or Common Polities? Reinterpreting the Democratic Peace." Journal of Politics 59 (May): 393–417.

Hall, N. (2004). Two concepts of causation. In: J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals*, chapter 9.

Hainmueller, Jens, and Michael J. Hiscox. "Attitudes toward highly skilled and low-skilled immigration: Evidence from a survey experiment." *American Political Science Review* 104, no. 01 (2010): 61-84.

Holland, Paul W. "Statistics and causal inference." *Journal of the American statistical Association* 81, no. 396 (1986): 945-960.

Imai, Kosuke, Luke Keele, Dustin Tingley, and Teppei Yamamoto. "Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies." *American Political Science Review*105, no. 04 (2011): 765-789.

Imai, K., Tingley, D. and Yamamoto, T., 2013. Experimental designs for identifying causal mechanisms. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *176*(1), pp.5-51.

Keele, Luke. "The Statistics of Causal Inference: A View from Political Methodology." *Political Analysis* (2015): mpv007.

Kessler, A. 2001. *Immigration, Economic Insecurity, and the "Ambivalent" American Public*. Working Paper. LaJolla, CA: Center for Comparative Immigration Studies.

Layne, Christopher. 1994. "Kant or Cant: The Myth of the Democratic Peace." International Security 19 (Fall): 5–49

Lewis, David. "Causation." *The journal of philosophy* 70, no. 17 (1974): 556-567.

Mahoney, James. "After KKV: The new methodology of qualitative research."*World Politics* 62.01 (2010): 120-147.

Maoz, Zeev. 1998. "Realist and Cultural Critiques of the Democratic Peace: A

Theoretical and Empirical Re-assessment." International Interactions 24: 3–89.

Mayda, A., and D. Rodrik. 2005. "Why Are Some People (and Countries) More Protectionist than Others?" *European Economic Review* 49 (6): 1393–430.

Hochschild, Jennifer L. *What's fair?: American beliefs about distributive justice*. Harvard University Press, 1986.

Oneal, John R., and Bruce Russett. 1999. "The Kantian Peace: The Pacific Benefits of Democracy, Interdependence, and International Organizations, 1885–1992." World Politics 52 (October): 1–37.

Pearl, J., 2009. *Causality*. Cambridge university press.

Ray, James Lee. 1995. Democracy and International Conflict: An Evaluation of the Democratic Peace Proposition. Columbia: University of South Carolina Press

Rosato, Sebastian. "The flawed logic of democratic peace theory." *American Political Science Review* 97, no. 04 (2003): 585-602.

Russett, Bruce. 1993. Grasping the Democratic Peace: Principles for a Post Cold War World. Princeton, NJ: Princeton University Press

Scheve, K., and M. Slaughter. 2001. "Labor Market Competition and Individual Preferences over Immigration Policy." *Review of Economics and Statistics* 83 (1): 133–45.

Schlozman, Kay Kehman, and Sidney Verba. *Injury to insult: Unemployment, class, and political response*. Harvard University Press, 1981.

Sinclair, B., McConnell, M. and Green, D.P., 2012. Detecting spillover effects: Design and analysis of multilevel experiments. *American Journal of Political Science*, *56*(4), pp.1055-1069.

Sondheimer, Rachel Milstein, and Donald P. Green. "Using experiments to estimate the effects of education on voter turnout." *American Journal of Political Science* 54.1 (2010): 174-189.

Spiro, David E. 1994. "The Insignificance of the Liberal Peace." International Security 19 (Fall): 50–86.

Weart, Spencer R. 1998. Never at War: Why Democracies Will Not Fight One Another. New Haven, CT: Yale University Press

Wise, Tess. 2014. "How shall we open the black box? Alternative approaches to the study of causal mechanisms in the social sciences." *Political Methodology Conference Proceedings.*