

The Future of Field Experiments in International Relations

By
SUSAN D. HYDE

Can international relations (IR) be studied productively with field experimental methods? The two most common existing empirical approaches in IR rely on cross-national data, detailed case studies, or a combination of the two. One as yet uncommon approach is the use of randomized field experiments to evaluate causal hypotheses. Applying such methods within IR complements other theoretical, case study, and observational research, and permits a productive research agenda to be built by testing the micro-foundations of theories within IR. This argument is illustrated by exploring how field experimental methods could be applied to two existing areas: how international institutions facilitate cooperation, and whether international actors can promote democracy in sovereign states.

Keywords: field experiments in international relations; democracy promotion; international cooperation

Can international relations be studied productively with field experimental methods? Within political science, the field of international relations (IR) focuses primarily on explaining international conflict and cooperation, and some of the more prominent research agendas center on war, peace, economic exchange, and the exercise of power at the international level. In part because the most prominent actors in these theories are states rather than individuals, there is a widespread perception that IR—or, at least, the most important questions within IR—cannot be studied systematically using field experimental methods. In this article, I evaluate this claim, discuss common objections to field experiments, and argue that field experiments

Susan D. Hyde is an assistant professor of political science and international affairs at Yale University. Her research interests include international influences on domestic politics, elections in developing countries, international norm creation, and the use of natural and field experimental research methods. Her current research explores the effects of international democracy promotion efforts, with a focus on international election observation.

DOI: 10.1177/0002716209351513

can play an important role in the empirical study of international relations. I illustrate this point by outlining several potentially productive areas for field experimentation. Although it is true that many features of international politics would be difficult, irresponsible, or even impossible to randomize at the international level, it does not follow that field experimentation cannot be fruitfully applied to the central questions within IR. On the contrary, many prominent theories imply micro-level behavioral implications that can be tested with field experimental methods.

The two most common existing empirical approaches in IR rely on cross-national data, detailed case studies, or combinations of the two. One as yet uncommon approach is the use of randomized field experiments to evaluate causal hypotheses.¹ Applying such methods within international relations complements other theoretical, case study, and observational research, and a productive research agenda can be built by testing the micro-foundations of IR theories. Although I optimistically explore the potential for using field experimental methods in IR, I do not wish to argue that this approach is the only way to go about answering questions within the field. Rather, field experiments are an underutilized method, and should be viewed as complementary to other research.

As the use of experimental methods expands in the social sciences, the ways in which experimental methods have been applied continues to grow. In this short article I focus my comments on field experiments, but there are several related methods with similar potential to contribute to the study of IR. Researchers may be able to take advantage of “natural” experiments in which the treatment variable is assigned in a manner that approximates randomization, but is not directly supervised by the researcher (Dunning 2008). Recent uses of natural experiments include the exogenous allocation of land titles to some but not all Argentinean squatters, allowing scholars to evaluate the causal effects of property rights on attitudes and behavior (Galiani and Scharfgrösky 2007); the “as-if” random assignment of international election observers to polling stations, allowing an evaluation of the effect of international monitors on election fraud (Hyde 2007); the use of arbitrary shelling by the Russian military to study the effects of indiscriminate violence on insurgent attacks (Lyll 2009); and the use of the unanticipated discovery of oil in São Tomé and Príncipe to evaluate the effect of natural resources wealth on public perceptions of corruption (Vicente 2006).

In addition to naturally occurring experiments, regression discontinuity designs, survey experiments, and laboratory experiments share many of the advantages of field experimentation. Regression discontinuity designs exploit an exogenous threshold or criteria in the assignment of the treatment variable, typically comparing outcome variables between units just above the predefined threshold to those just below the threshold.² Survey experiments embed randomized experiments within individual surveys in order to understand, for example, the effects of providing varying amounts of information to voters.³ Laboratory experiments typically involve the random assignment of a treatment to research subjects in the controlled world of university laboratories, although some innovative studies have moved the laboratory model into the field (Habyarimana et al. 2007, 2009).

Each of these methods utilizes randomization as an element of research design, and each has the potential to demonstrate causality in a way that other methods do not.

Although it is not likely that the existence of states, the outbreak of war, the possession of nuclear weapons, the distribution of power, membership in international organizations (IOs), or military capacity will ever be randomized for the purposes of field experimentation, the behavior of individuals and design of institutions that influence important aspects of these international phenomena frequently take place on a smaller scale. Theories of international cooperation, conflict, or international pressure are primarily discussed in the abstract at very general levels of analysis, but many also have implicit or explicit micro-level implications.⁴ Comparative politics and development economics have already begun to move in this direction, as contributors to this volume emphasize. Although the scale is arguably different at the international level, the IR literature is already blurring the lines between comparative politics and IR, disaggregating the state in order to explain international political phenomena and explore how international politics influences politics within states. The possibilities for field experimentation increase as scholars move away from theories that treat states like “billiard balls,” defined primarily by their relative power and influence.

My assertion that scholars of IR can and should increase their use of field experiments relies on researchers’ ability and willingness to move between macro-level theories and micro-level implications. This macro-micro movement requires well-specified theories and definition of the mechanisms at work in each theory. It also requires that scholars work to connect their micro-level findings to the broader implications of their study. To illustrate, the simplified theoretical argument that international organizations facilitate international cooperation is not sufficiently precise (as I just stated it) to generate micro-level testable implications. It is unlikely that social scientists will ever be able to compare the current state of the world in which the United Nations exists with the counterfactual world in which all else is held equal but the United Nations does not exist. I therefore agree with prominent scholars of international institutions who argue that “rarely, if ever, will institutions vary while the ‘rest of the world’ is held constant” (Keohane and Martin 1995, 47). Yet this does not mean that implications of these theories of international cooperation cannot be tested using field experimental methods.

To clarify my expectations about what can be learned from such micro-level testing, I would note that under ideal conditions, field experiments test whether randomized variation in x causes change in variable y . Thus, if the experiment reveals that x causes a change in y , the experiment will have provided strong empirical support for the theoretical expectation that x *can* cause a change in y . In terms of micro-level tests of macro-level theories, a single field experiment will rarely be able to demonstrate that x *will always* cause a change in y . Similarly, if a single micro-level field experiment reveals that x did not cause a change in y , it does not prove that x will never cause a change in y , nor even whether x usually causes a change in y . Replication of field experiments is therefore important, as it is often possible that x causes a change in y only under certain conditions.

Although experiments have an unrivaled ability to demonstrate cause and effect (Druckman et al. 2006), the types of experiments proposed in this article are most likely to demonstrate that a causal relationship between the hypothesized variables *can* exist. As a result, replication and careful consideration of the conditions under which the expected relationships are most likely to hold are also important aspects of field experimentation, and may mean that some patience is required before field experiments yield influential and field-changing discoveries.

To facilitate debate on the subject and to illustrate my argument, I outline several ways in which experimentation could be applied within two prominent research agendas in IR: how international institutions facilitate international cooperation, and whether (or how) foreigners can promote democracy in sovereign states. Before outlining these examples, I address several common concerns about the application of field experiments in IR.

What about the “Big” Questions in IR?

An increasingly common objection to field experiments in other subfields of political science is that they are used to study only insignificant phenomena and do not contribute to the body of knowledge about the “big” questions that most social scientists are concerned with. For field experimentation in IR, this criticism underscores the need for theoretical precision, and for a broader research agenda in which complementary studies are carried out in tandem. As I stated above, experimental methods have the greatest potential when they are used to evaluate the micro-foundations of macro-level theories. Some theories lack sufficient theoretical precision and micro-level implications to generate such testable theoretical predictions. For example, Kenneth Waltz’s (1979) argument that the international system is most stable when it is characterized by balancing between two powers does not—at least as I read it—have clear micro-level implications that could be tested with field experimental methods.

In contrast, scholars of international institutions have defined a number of mechanisms—at least in theory—by which institutions facilitate cooperation, including reduced transaction costs and increased information, which make it possible to test the causal relationships outlined in the theory, as I propose below.

Additionally, if field experiments are used to test these micro-foundations, multiple experiments (combined with other research) should be conducted in order to evaluate the relevant components of any given theory. This would require a large and deliberately organized research agenda that would include multiple field experiments testing different implications of a specific theory, and may require greater cooperation among researchers than is currently the norm in the discipline. Although ambitious, the payoff of such a research agenda would be large: proven causal relationships that add up to a more complete understanding of the foundations of international politics.

Why Would States, IOs, NGOs, or Other Non-State Actors Cooperate with Researchers?

Another common concern among scholars relates to the degree to which field experiments require the involvement of governments, IOs, NGOs, or other international actors. Bureaucracies are notoriously uninterested in high-quality evaluation of their own effectiveness. As scholars of international organizations have argued, evaluation of the work of IOs is often carried out by the same actors whose job performance is at stake (Barnett and Finnemore 1999; Easterly 2002), giving them a vested interest in “finding” that their programs and policies are effective in accomplishing their stated goals.⁵

This incentive structure arguably creates an environment in which the actors who could cooperate with researchers on field experiments may not wish to expose their programs to criticism. If learning that their actions are not having the intended effects would threaten their jobs, one can see why they might be hesitant to introduce randomization. Yet there is reason to believe that this incentive structure is not a serious barrier to randomization. Although practitioners tend to be wary of crusading academics of all stripes, field experimental methods appear to be attracting some momentum in fields that are closely related to IR theory. Organizations such as the World Bank have proven willing to adopt field experimental methods in the study of development programs (Duflo and Kremer 2004; Miguel and Kremer 2004; Olken 2007), and organizations such as USAID and the Millennium Challenge Corporation have begun to use these methods in evaluating the effectiveness of foreign aid (Millennium Challenge Corporation 2009; National Research Council 2008). Where the interests of researchers and policymakers overlap, field experimentation is most likely to be successful.

Anecdotally, it appears that field experimentation is the easiest sell when there is an existing demand among policymakers to identify which policies and practices work and which do not. Because field experiments can be used to adjudicate between competing theories in a transparent manner, the effects of various programs can be compared in a scientifically rigorous manner (see Olken [2007] for an excellent example).

Even lacking this type of scenario, many possible applications of randomization represent minimal change from existing practice and therefore do not require additional expenditures, such as the randomized phasing in of a development project over time in contrast to the planned phasing in of a development project based on some other arbitrary or non-random criteria. To the extent that these small changes would allow researchers to study important and interesting research questions at little cost to the partner organizations, there are large payoffs to persuading organizations that cooperation is worthwhile. I now propose several topics within IR that can be productively studied with field experimental methods.

Experiments, Information, Transaction Costs, and International Cooperation

One very broad research agenda within IR focuses on explaining international cooperation between states under anarchy. I do not summarize the extensive literature on the subject.⁶ Instead, I highlight the conclusions of two of the most prominent scholars in this tradition before exploring how we might begin testing this theory with field experimentation. Robert Keohane and Lisa Martin, building on other work on international cooperation, argue that international cooperation is facilitated because international “institutions can provide information, reduce transaction costs, make commitments more credible, establish focal points for coordination, and in general facilitate the operation of reciprocity” (1995, 42).

In evaluating this and related arguments, scholars have established correlations between the work of international organizations or state participation in international institutions and plausible positive or negative effects using cross-national data and detailed case studies, but these relationships are difficult to prove causally. In moving toward field experimental evaluations of this theory, it is helpful to evaluate the theory in terms of its component parts. Within international institutions, randomizing the provision of information in order to test its effect on international cooperation is a much more realistic goal than randomizing the existence of the international institutions.

How, exactly, does information facilitate international cooperation? In theory, one of the barriers to international cooperation is that states possess poor information about other states. Although there may be potential gains from cooperation that may be quite clear to all interested parties, states are wary of engaging in mutually beneficial cooperation with partners whose intentions, preferences, or capabilities are uncertain. Therefore, as the theory goes, international cooperation should be more likely as information provision increases. Both international cooperation and information provision are vague concepts, but they are useful because they apply across a variety of issue areas.

To make the case more concrete, I narrow the focus further to the effect of information provision on international cooperation within the World Trade Organization (WTO), which, according to its own public documents,

provides a forum for negotiating agreements aimed at reducing obstacles to international trade and ensuring a level playing field for all, thus contributing to economic growth and development. The WTO also provides a legal and institutional framework for the implementation and monitoring of these agreements, as well as for settling disputes arising from their interpretation and application. (World Trade Organization 2009)

According to its advocates, the WTO is useful because it helps prevent and resolve trade disputes, thus increasing trade overall and resulting in increased mutually beneficial cooperation. However, information provision within the WTO is already imperfect, with states varying in their ability to collect and distribute relevant information. Wealthy and influential countries are able to maintain a

full-time delegation at the WTO headquarters in Geneva and to provide information to their government on relevant issues and negotiations. They are also more likely to use the WTO's dispute resolution mechanisms. Because many states cannot afford to have the extensive staff that it would take to track every relevant piece of information within the WTO and to communicate it to the relevant stakeholders within their countries, and because it is well established that some countries use WTO dispute resolution mechanisms less than they "should," a plausible extension of Keohane and Martin's argument is that providing increased information and technical expertise (broadly defined) to its members should increase cooperation within the WTO.

Thus, assuming that increased information would benefit a well-defined subset of states, it is possible to randomize the provision of such potentially valuable information. The effect of information provision could then be tested on a variety of outcome measures related to international trade, such as aggregate levels of imports and exports, trade between member states, or the use of WTO dispute resolution mechanisms. Increased information provision within the WTO could also make it more likely that individual exporting or importing businesses within countries could better identify favorable conditions for economic exchange. The variables that may be affected by information provision are diverse, and depend on the type of information provided, but it should be theoretically possible to exploit this unmet demand for information within the WTO to test whether increased provision of information causes increases in international cooperation. It would be left to the researcher to determine the exact mechanism by which this information was distributed and the type of information that would be most likely to facilitate international cooperation. For example, it is well known that the least developed countries within the WTO, or LDCs, are likely to have small or non-permanent delegations at the WTO headquarters, and are also likely to export commodities, such as agricultural products or textiles, that face high barriers to international markets. A number of scholars, policymakers, NGOs, and the WTO itself have recognized the challenges faced by LDCs in accessing international markets and in effectively utilizing the WTO.

A permanent delegation at the WTO is expensive to maintain and represents a serious barrier to LDCs in maximizing their potential gains from WTO membership. Improving LDC access to international markets and to the WTO is a relatively widely agreed upon goal within the organization, and one that I take as given in outlining possible field experiments. If the argument is that information provision within the WTO—as provided to many WTO members by their permanent delegations in Geneva—helps to facilitate international cooperation, one form of information provision that could be randomized is the capacity of some LDCs to fund a permanent delegation, which would increase the country's access to information. Because such funding is expensive, LDCs could gain access to (periodically rotating) grant money in order to support a permanent delegation through a lottery system. This would be fair in that all LDCs would have access to the funding, although not all would receive it. Because such grants would target LDCs that do not already have sufficient permanent

delegations, they would be providing more services than would exist in the absence of such a program.

This process would generate a randomly selected subset of LDCs, and would allow both the organization and scholars to learn about the effects of information provision on international trade. Alternatively, expert information could be provided to a random subset of LDCs by external organizations, such as the Agency for International Trade Information and Cooperation (AITIC), an organization with the stated objective of providing LDCs “with information and policy advice to help them integrate into the multilateral trading system and the work of the WTO and other international trade-related organisations in Geneva (AITIC 2009).” If AITIC, or a similar organization, was interested in learning how to best accomplish its stated goals of facilitating trade, it might be willing to randomize some aspects of its work.

Similarly, but more narrowly, if WTO negotiations pertaining to a specific issue area or sector—such as hand-woven textiles—were of particular interest to manufacturers within LDCs, but their access to information about the negotiations was limited in a manner that hurt their ability to trade internationally, relevant information could be provided to a random sample of interested manufacturers of hand-woven textiles, and the average effects of this information on their behavior could then be tested. Although such information could have differential effects on manufacturers in the short term and might not be perceived as fair, learning about how to best help these specialized manufacturers take advantage of international markets could have much greater benefits in the long term.

Similar field experiments could be designed for the other components of theories of international cooperation, such as reducing transaction costs, providing focal points, or encouraging reciprocity, as outlined by Keohane and Martin. I now turn to the related topic of democracy promotion in which field experiments are somewhat more common, but in which the link between existing micro-level findings and macro-level theories of international relations are rarely made explicit.

Experiments and Democracy Promotion

Democracy promotion, or efforts by external actors to encourage the development of democratic political institutions in other countries, is an overtly stated goal of many international organizations, including the Organization of American States, the Organization for Security and Cooperation in Europe, the European Union, and the United Nations (Youngs 2001; Newman and Rich 2004; Rich 2001; Pevehouse 2005). Democracy promotion is also an important topic in foreign policy, as a number of states link democracy promotion to foreign aid and long-term strategic goals, and some argue that increasing the number of democracies in the world will improve interstate relations on a variety of fronts (Cox, Ikenberry, and Inoguchi 2000; Smith 1994). Given the increased focus on democracy promotion since the end of the cold war, scholars and practitioners agree upon

surprisingly little about the effectiveness of democracy promotion activities. This is true even though the effectiveness of democracy-promotion programs is of interest to many audiences within political science and within the democracy promotion industry.⁷

Some existing cross-national research addresses the question of whether money spent on democracy promotion is effective. Scholars in this vein compare aggregate totals of aid or money spent on democracy assistance to a country's relative score on aggregate indices of democratic institutions or political rights. The results are mixed, and the expected cross-national empirical patterns are not clear (Knack 2004; Burnell 2008, 2000; Finkel, Pérez-Liñán, and Seligson 2007). At a theoretical level, the various ways in which democracy and foreign aid can be linked represent a serious barrier to cross-national empirical research. It is possible that donors give more money to the countries with the furthest to go in bringing about democratization, as donors could view these countries as areas in which they can achieve more progress for their money. It is also possible that donors give aid to reward governments that have already moved toward democracy and are relatively close to full democratization. In the first model, cross-national correlations would show a relationship between more aid and less democracy. In the second model, cross-national correlations would show a positive relationship between aid and democracy. These cross-cutting logics make it more difficult to interpret these cross-national patterns and lead to a mixed picture of the relationship between democracy assistance and democratization.

Another part of the problem in evaluating the effects of democracy promotion is that democracy is a widely contested concept, with little agreement on how to best measure progress toward democratization. Although the relationship between aid and democracy scores can be informative, these studies cannot reveal that democracy assistance or aid tied to democracy causes democratization. A potentially more productive line of research focuses on the effects of specific democracy-promotion activities, such as political party training, international election monitoring, support for nonpartisan domestic observers, improving the accuracy of voter registration, training journalists and supporting the development of a free and independent media, democratization of local governance structures, and a variety of other programmatic activities tailored to the challenges of individual countries.

Because many of these democracy-promotion activities are implemented at the subnational level and democracy promoters have an interest in understanding the effects of their work, field experiments are likely to be feasible and useful. The challenge lies in linking micro-level effects of democracy promotion programs to a macro-level theory and conclusion.

For example, democratic elections are an important and fundamental element of democracy. Elections are frequently one of the first steps toward democratization, and international monitoring of elections has spread in part because international actors and leaders of democratizing countries argue that their presence brings about cleaner elections by reducing election fraud and increasing voter confidence in the electoral process. The proposition that observers reduce election

fraud is widely asserted, is frequently met with skepticism by academics, and, until recently, was untested. By randomly assigning international election observers to polling stations, it is possible to compare the group of unmonitored polling stations to the group of monitored polling stations and test whether international observers can reduce election day fraud. The implication of the study is that IOs can reduce such fraud and help bring about cleaner elections. Randomly assigning international election observers to polling stations does not prove that democracy promotion causes democratization, but it does test whether the specific forms of democracy promotion can have the intended effect of deterring fraud or improving the quality of the electoral process.

As with testing theories of international cooperation, democracy promotion represents an ideal application of experimental methods for two primary reasons. First, some researchers are uncomfortable with field experiments because they represent direct intervention in the “real” world. Applying these methods to the study of international cooperation or democracy promotion would most often entail—at least as I envision it—randomizing components of existing programs. Because the work of IOs already represents a direct intervention by international actors in the “real” world, introducing randomization to some components of this work implies little additional researcher-driven interference. Second, the relevant organizations should have a vested interest in understanding the effects of their actions and the conditions under which the work of the organization actually succeeds in accomplishing their goals. The WTO should, according to its own publications, be interested in determining the best ways to increase international trade among its member states. Organizations and states engaged in democracy promotion should want to understand the conditions under which their actions are most effective and whether their programs actually accomplish their intended objectives.

Similar to other areas in which field experimentation has already been successful, the introduction of randomized field experiments could be used to improve understanding of the effects of international organizations across a variety of other issue areas. Identifying areas in which there is clear overlap between the interests of researchers and international organizations should be a priority for scholars interested in applying field experiments to empirical questions relevant to IR.

Challenges Inherent in This Approach?

One of the central challenges in the use of field experiments relative to other methods is their timing. Unlike the majority of research in political science, the bulk of work on a given project takes place before the event under study occurs. This requires that researchers invest significant amounts of time in a project before they know whether or not it will be successful. Many field experiments in IR will require well-developed connections within international organizations or well-funded research projects. Developing these connections within IOs is not an impossible task, but may require significant investments of time before a research

project is ensured, an incentive structure that may be incompatible with the professional demands on academics.

Field experiments are also characterized by uncertain external validity. Although this method allows the demonstration of cause and effect, generalization to other contexts is frequently difficult. Some of this problem can be addressed by replicating experiments in multiple contexts. Replication of similar experiments across varying contexts will allow researchers to identify the contextual variables that influence the hypothesized relationships across experiments.

Finally, some IR scholars may be tempted to dismiss field experimentation based on the fact that the method has not (yet) generated a field-changing finding. Such a discovery would undoubtedly lead to a dramatic increase in the use of field experiments in IR, although it is perhaps worth noting that field experiments can already be applied to a number of important questions in international relations; a big discovery in the field will not change what we already know about field experiments, but simply will increase awareness of their possibilities.

Conclusion

The goal of this article was to motivate discussion about the most productive application of field experimental methods to international relations. Combined with other research, field experiments have the potential to test existing theories by evaluating causal relationships at the micro-level. Clearly, randomization cannot be applied at all levels of analysis and will not be applicable to all relevant questions within IR. However, there remain many opportunities, particularly in areas that are important to both policymakers and scholars. As the IR literature currently stands, these methods could be employed in a number of issue areas to better understand the causal effects of attempts by international actors to influence international and domestic politics. The willingness of international actors to cooperate with researchers has already been demonstrated in the field of development economics and is beginning to be demonstrated in the field of democracy promotion. Because there are a number of substantive areas in which the interests of practitioners and researchers overlap, there is significant potential for this type of research, although it will require researchers to reprioritize how they invest their time in new research projects. It may also require practitioners and policymakers to put increased emphasis on learning about the effectiveness of their policies.

The stakes are high in international relations, and a strong argument can be made that failing to understand causal relationships within IR is more risky than attempting to apply field experimentation in new and consequential areas of international politics. The next step is to identify those areas most likely to fit within this research agenda, potentially focusing on topics in which the interests of scholars and practitioners overlap, and carefully outline the micro-level implications and testable implications of these theories. A number of other scholars are already progressing in this direction, including ongoing (unpublished) field

experiments related to the evaluation of post-conflict recovery programs, methods by which international actors can improve local governance and economic development, the political effects of development assistance, and tests of how the provision of public goods relates to violence.

Notes

1. For a recent review of the use of field experiments in the political economy of development, see Humphreys and Weinstein (2009).
2. For a non-IR example, see Angrist and Lavy (1999).
3. See, e.g., Tomz (2007).
4. See Levy (1997) for a relevant discussion about the challenges inherent in linking the behavior of individuals in laboratory experiments to the behavior of state leaders.
5. These incentives are not exclusive to international organizations.
6. See, e.g., Axelrod and Keohane (1985); Oye (1985); Abbott and Snidal (1998); Keohane and Martin (1995); Martin and Simmons (1998).
7. One might argue that democracy promotion is a topic of comparative politics as well, and therefore, it is not a good example of field experimentation in IR. Although democracy promotion is also relevant to comparative politics, it has already been highlighted as a subject of interest to a number of prominent scholars of international politics, history, and law, and represents an area of productive overlap between IR and comparative politics (Burnell 2008; Carothers 2006; Cox, Ikenberry, and Inoguchi 2000; Youngs 2001; Smith 1994).

References

- Abbott, Kenneth W., and Duncan Snidal. 1998. Why states act through formal international organizations. *Journal of Conflict Resolution* 42 (1): 3-32.
- AITIC. 2009. About us: AITIC in short. Accessed October 24, 2009. Available at: <http://www.acici.org/aitic/aboutus/index.html>.
- Angrist, Joshua D., and Victor Lavy. 1999. Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics* 114 (2): 533-75.
- Axelrod, Robert, and Robert O. Keohane. 1985. Achieving cooperation under anarchy: Strategies and institutions. *World Politics* 38 (1): 226-54.
- Barnett, Michael N., and Martha Finnemore. 1999. The politics, power, and pathologies of international organizations. *International Organization* 53 (4): 699-732.
- Burnell, Peter J. 2000. *Democracy assistance: International co-operation for democratization*. London: F. Cass.
- Burnell, Peter J. 2008. From evaluating democracy assistance to appraising democracy promotion. *Political Studies* 56 (2): 414-34.
- Carothers, Thomas. 2006. The backlash against democracy promotion. *Foreign Affairs* 85 (2): 55.
- Cox, Michael, G. John Ikenberry, and Takash Inoguchi, eds. 2000. *American democracy promotion: Impulses, strategies, and impacts*. Oxford: Oxford University Press.
- Druckman, James N., Donald P. Green, James H. Kulinski, and Arthur Lupia. 2006. The growth and development of experimental research in political science. *American Political Science Review* 100 (4): 627-35.
- Duflo, Esther, and Michael Kremer. 2004. Use of randomization in the evaluation of development effectiveness. In *Evaluating Development Effectiveness*, ed. George K. Pitman, Osvaldo Feinstein, and Gregory Ingram, 205-232. World Bank Series on Evaluation and Development. New Brunswick, NJ: Transaction Publishers.
- Dunning, Thad. 2008. Improving causal inference: Strengths and limitations of natural experiments. *Political Research Quarterly* 61 (2): 282-93.

- Easterly, William. 2002. The cartel of good intentions. *Foreign Policy* 131:40-49.
- Finkel, Steven E., Aníbal Pérez-Liñán, and Mitchell A. Seligson. 2007. Effects of US foreign assistance on democracy building. *World Politics* 59 (3): 404-39.
- Galiani, Sebastian, and Ernesto Scharngrodsy. 2007. Property rights for the poor: Effects of land titling. Unpublished manuscript.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein. 2007. Why does ethnic diversity undermine public goods provision? *American Political Science Review* 101 (4): 709-25.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the dilemmas of collective action*. New York: Russell Sage Foundation.
- Humphreys, Macartan, and Jeremy M. Weinstein. 2009. Field experiments and the political economy of development. *Annual Review of Political Science* 12:367-78.
- Hyde, Susan D. 2007. The observer effect in international politics: Evidence from a natural experiment. *World Politics* 60 (1): 37-63.
- Keohane, Robert O., and Lisa L. Martin. 1995. The promise of institutionalist theory. *International Security* 20 (1): 39-51.
- Knack, Stephen. 2004. Does foreign aid promote democracy? *International Studies Quarterly* 48 (1): 251-66.
- Levy, Jack S. 1997. Prospect theory, rational choice, and international relations. *International Studies Quarterly* 41 (1): 87-112.
- Lyll, Jason. 2009. Does indiscriminate violence incite insurgent attacks? Evidence from Chechnya. *Journal of Conflict Resolution* 53 (3): 331-362.
- Martin, Lisa L., and Beth A. Simmons. 1998. Theories and empirical studies of international institutions. *International Organization* 52 (4): 729-57.
- Miguel, Edward, and Michael Kremer. 2004. Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72 (1): 159-217.
- Millennium Challenge Corporation. 2009. Impact evaluation at MCC. <http://www.mcc.gov>.
- National Research Council, Committee on Evaluation of USAID Democracy Assistance Programs. 2008. *Improving democracy assistance: Building knowledge through evaluations and research*. Washington, DC: National Academies Press.
- Newman, Edward, and Roland Rich, eds. 2004. *The UN role in promoting democracy: Between ideals and reality*. Tokyo: United Nations University Press.
- Olken, Benjamin A. 2007. Monitoring corruption: Evidence from a field experiment in Indonesia. *Journal of Political Economy* 115 (2): 200-249.
- Oye, Kenneth A. 1985. Explaining cooperation under anarchy: Hypotheses and strategies. *World Politics* 38 (1): 1-24.
- Pevehouse, Jon C. 2005. *Democracy from above: Regional organizations and democratization*. Cambridge: Cambridge University Press.
- Rich, Roland. 2001. Bringing democracy into international law. *Journal of Democracy* 12 (3): 20-34.
- Smith, Tony. 1994. *America's mission: The United States and the worldwide struggle for democracy in the twentieth century*. Princeton, NJ: Princeton University Press.
- Tomz, Michael. 2007. Domestic audience costs in international relations: An experimental approach. *International Organization* 61 (4): 821-40.
- Vicente, Pedro C. 2006. Does oil corrupt? Evidence from a natural experiment in West Africa. SSRN eLibrary. <http://ssrn.com/paper=829865>.
- Waltz, Kenneth. 1979. *Theory of international politics*. New York: McGraw-Hill.
- World Trade Organization. 2009. WTO | What is the WTO?—About the WTO—A statement by the Director-General. http://www.wto.org/english/thewto_e/whatis_e/wto_dg_stat_e.htm (accessed April 30, 2009).
- Youngs, Richard. 2001. *The European Union and the promotion of democracy*. Oxford: Oxford University Press.