Registering Theory-Based Predictions in Political Science

Andrew Cesare Miller, *United States Naval Academy, USA*

ABSTRACT How can political scientists rigorously evaluate the predictive power of theories? Many peer-reviewed political science articles include predictions about future outcomes, and scholars make predictions on social media and other public forums. The prevalence of predictions suggests that scholars recognize the utility of leveraging theories for this purpose; however, the predictions often are not made in a manner that allows for rigorously evaluating their accuracy. Building on the increasing popularity of study preregistration in the social sciences, this article proposes "prediction registration" as a means for scholars to publish falsifiable, systematic, and verifiable theory-based predictions. Increasing the rigor of predictive theory testing can advance often-circular debates about accuracy and presents a "win-win" for scholars who aim to test the predictive power of theories. With a more rigorous approach, correct predictions would better demonstrate a theory's ability to forecast outcomes, and missed predictions would reveal information that can be used to calibrate the theory.

Theory-based predictions—that is, the application of a theory to estimate the likelihood of future outcomes—appear in political science scholarship and public commentary. The fact that political scientists make predictions suggests that many in the discipline view forecasting as an important use of theory; some might even agree with the economist Milton Friedman's (1966) proposition that predictive accuracy is *the* primary metric by which to judge a theory's utility. However, the rigor expected for testing claims of theories' predictive power lags considerably behind the rigor expected for testing claims of the explanatory power of theories. Not all political science theories aim to provide predictive power, but many of those that do have not undergone rigorous testing.

This article introduces "prediction registration" as a framework to facilitate the rigorous testing of predictive power. Prediction registration involves scholars posting theory-based predictions about outcomes that have yet to occur on OSF Registries (2022) or a similar registration site as part of the study publication process. Unlike conventional study registration in which scholars wait until an outcome is revealed to publish their findings, prediction registration calls for publishing the predictions about future outcomes—which may occur over a long time horizon—prior to the results being known. In all, the prediction registration framework (1) specifies the parameters required for making predictions falsifiable; (2) facilitates the systematic aggregation of predictions for a given theory; and (3) provides a process for establishing an externally verifiable prediction record.

Prediction registration builds on a growing political science literature about prediction that broadly falls into two categories. One category evaluates individuals' or groups' ability to predict political outcomes, focusing on identifying "superforecasters" (Horowitz et al. 2019; Tetlock and Gardner 2015). A second category leverages machine learning and other statistical techniques to develop primarily inductive predictive models from which it can be difficult to discern the underlying theory (Grimmer, Roberts, and Stewart 2021; Hegre et al. 2013).[1]

Superforecasting and machine-learning approaches focus on the accuracy of predictions; therefore, neither approach requires explicitly specifying the theories that underlie predictions. In contrast, the prediction registration framework does not evaluate predictive accuracy alone but rather evaluates the predictive accuracy of theories themselves. Superforecasters, scholars, and policy

Andrew Cesare Miller *is assistant professor of political science at the United States Naval Academy. He can be reached at millera@usna.edu.*

doi:10.1017/S1049096524000404

makers often rely on theories, so there is value in identifying those that offer predictive power. Unlike individuals engaged in prediction, theories can be applied, adapted, and improved over time, rendering them building blocks for aggregating knowledge. Theories also can fill gaps in machine-learning approaches that remain limited by the unavailability of machine-readable data related to many important political science topics (Cederman and Weidmann 2017; Montgomery and Sagan 2009).

Prediction registration involves scholars posting theory-based predictions about outcomes that have yet to occur on OSF Registries or a similar registration site as part of the study publication process.

This article proceeds in four sections. First, I discuss the prevalence of predictions in political science despite many of the predictions being unfalsifiable. Second, I lay out the process of prediction registration to provide a framework for making theory-based predictions more rigorous. This section also introduces the Prediction Registration Template, which specifies the parameters for boosting the rigor of predictions. Third, I address challenges to the evaluation of the predictive power of theories. Fourth, I discuss practical considerations for overcoming hurdles that might arise in implementing prediction registration by drawing on lessons from the broader adoption of preregistration for experimental studies.

THE PREVALENCE OF PREDICTIONS

Theory-based predictions are a feature of political science (Schneider, Gleditsch, and Carey 2011). Political scientists make predictions about a range of outcomes from election results (e.g., Dassonneville and Tien 2020) to the duration of armed conflict (e.g., Pilster and Böhmelt 2014). In the international relations (IR) subfield, Fomin et al. (2021) find that of the 5,559 articles published in the top IR journals between 1992 and 2014, 817 (15%) included what the authors coded as a prediction.[2]

Some of these predictions are made with a rigorous approach that renders them falsifiable, systematic, and verifiable. At the same time, the majority of predictions in the discipline appear to be unfalsifiable.[3] At a minimum, falsifiability—that is, a prediction formulated such that it can be proven wrong—requires the specification of time frames within which the prediction is expected to manifest, measurable outcomes and independent variables, and scope conditions identifying the cases to which the theory applies. According to Fomin et al.'s (2021) data, only 27 of the 817 articles included a prediction of a specific event within a specified time frame. Moreover, from my analysis of these 27 articles (figure 1),

Figure 1
Falsifiability of Predictions in International Relations Articles

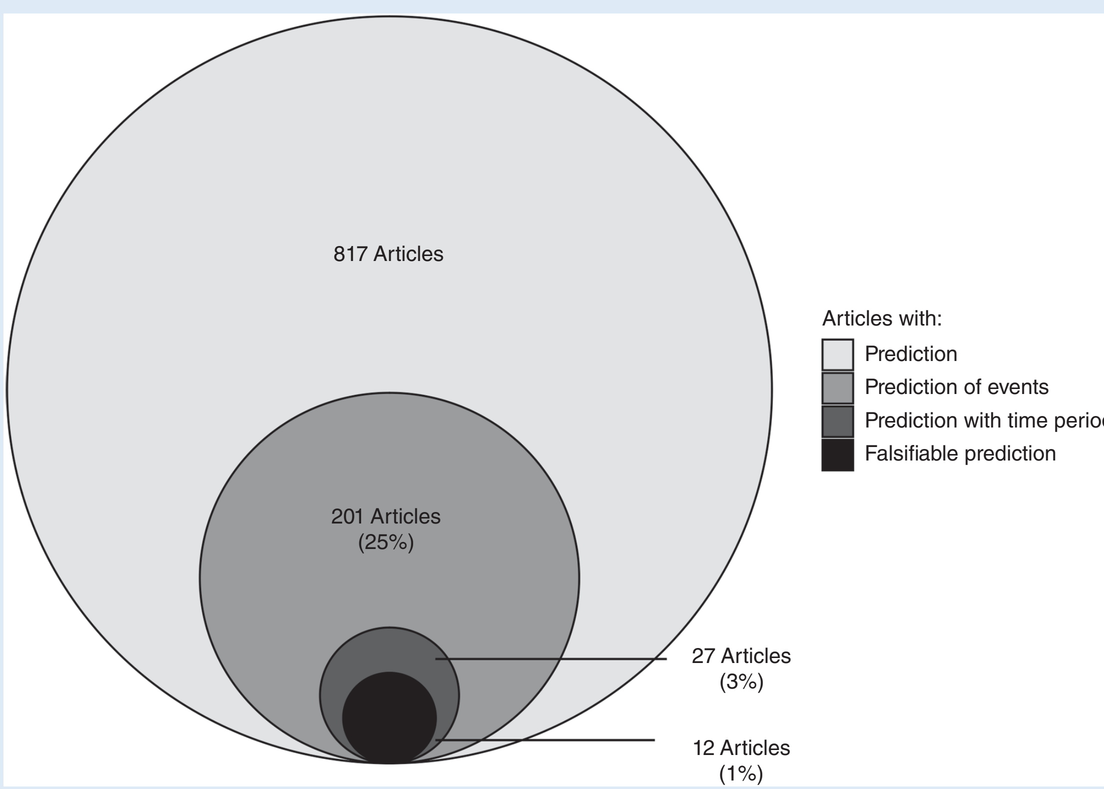

https://doi.org/10.1017/S1049096524000404 Published online by Cambridge University Press

only 12—approximately 1% of all articles with a prediction—contained what arguably could be considered at least one falsifiable prediction and thus could be evaluated for accuracy.[4]

Scholars do not limit their predictions to academic outlets. They also leverage theories to predict events in public forums, as demonstrated by the public commentary surrounding Russia's invasion of Ukraine on February 24, 2022. Mearsheimer (2015), for instance, applied his theory of offensive realism to predict that "the West is leading Ukraine down the primrose path, and the end result is that Ukraine is going to get wrecked" in a YouTube video that received more than 29 million views. More recently, between January 1 and February 23, 2022, political scientists with a research focus on security and/or Europe posted on Twitter (now X) 96 tweets with predictive content about the Russian invasion.[5] As shown in figure 2, of those tweets, 45% were direct predictions; 26% were recommendations or assessments of policies with implied predictions; and, in 29% of the tweets, scholars evaluated prior predictions related to the conflict.[6]

The adoption of a common framework to make tests of predictive power falsifiable, systematic, and verifiable can advance positivist political science. From a Popperian perspective, falsifiability could mute often-circular debates about whether predictions are correct (Popper 2002). More broadly, a common framework provides a foundation for aggregating individual predictions into what the philosopher Imre Lakatos (1970) termed "research programs." Advances in science, according to Lakatos, turn on coordinated testing of a "protective belt" of theories centered around a "hard core" of foundational assumptions underpinning a research program. Increasing the rigor of predictive tests by specifying the predictions' theoretical bases through registration can help to aggregate predictions based on a given theory and to situate the predictions within their respective research programs.

Rigorously testing the predictive power of theories also can strengthen scholarly contributions to policy debates. When scholars make theoretically informed policy recommendations, they are either explicitly or implicitly predicting an outcome conditional on their recommendation being implemented. As political scientist Kristian Skrede Gleditsch (2022) posits, policy and prediction are closely related—"like love and marriage," in his words (see also Friedman 1966).

Rigorously testing the predictive power of theories also can strengthen scholarly contributions to policy debates. When scholars make theoretically informed policy recommendations, they are either explicitly or implicitly predicting an outcome conditional on their recommendation being implemented.

It can be argued that testing the ability of theories to predict future events is unnecessary given that scholars often make theory-based "predictions" about prior events using data from outside of the sample from which the theories were developed (e.g., King, Keohane, and Verba 1994). However, when out-of-sample data are retrospective in this way, it is difficult to verify that the data did not influence the theory's development. Moreover, even if retrospective predictions are consistent with fully out-of-sample prior events, they cannot account for temporal trends that might influence the outcomes of future events (e.g., Bowlsby et al. 2020).

Figure 2

Predictive Content about the Russian Invasion of Ukraine on Twitter

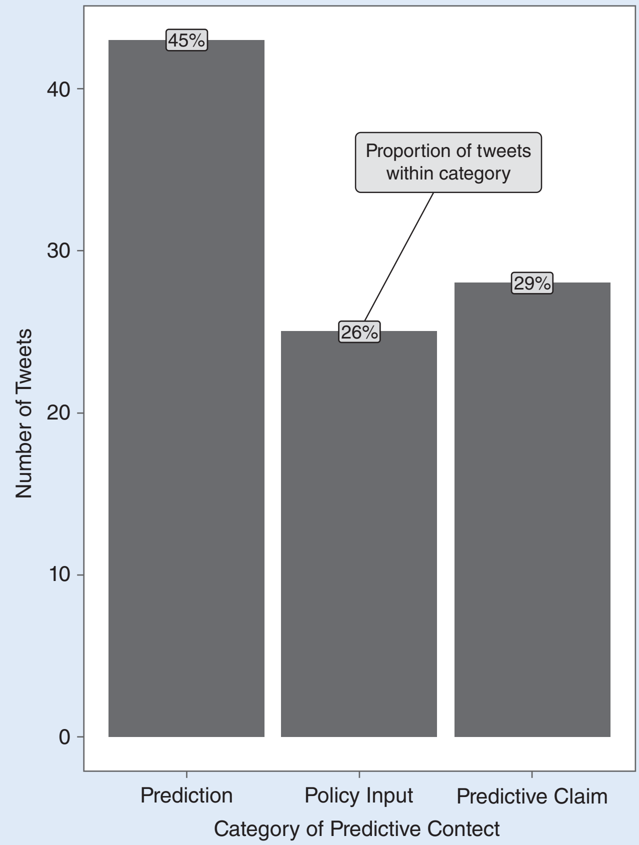

PREDICTION REGISTRATION

Registering predictions would render predictive theory testing more rigorous. Conventional study registration typically involves a researcher predicting an outcome before data collection, collecting the data, and analyzing the results, and only then publishing the findings in an academic outlet (Jacobs 2020; Nosek et al. 2018). Many political science theories, however, relate to events that might occur within wide temporal windows. Indeed, my coding of the 27 articles identified by Fomin et al. (2021) demonstrates that the median predicted event would occur eight years from publication, with the maximum time period being 78 years in the future.

https://doi.org/10.1017/S1049096524000404 Published online by Cambridge University Press

Figure 3

Contents of the Prediction Registration Template

Study Information
- Study Title
- Contributors
- Version & Date
- Study Summary
- Keywords

Theory & Measurement
- Theory Background
- Scope Conditions
- Study Variables
- Variable Measurement
- Evaluation Metrics
- Alternative Baseline

Predictions
- Initial Predictions
- Update Criteria

Additional Information
- Conflicts of Interest
- Acknowledgments
- References

Prediction registration modifies the conventional registration process by calling on researchers to register and publish theory-based predictions—even if the outcome is not yet known—to test theories that are claimed to have predictive power. The prediction registration process has three principle benefits. First, the prediction registration framework specifies the parameters required for a prediction to be falsifiable. A Prediction Registration Template (available at https://github.com/miller-research/prediction-template) lists the parameters to be included in a registration. Figure 3 shows the template's table of contents. Second, the framework facilitates the aggregation of predictions related to a theory over time so that the predictions (and thus the theory's predictive power) can be systematically evaluated together rather than as disparate "one-off" predictions. Third, it provides a transparent record of predictions so that they can be verified easily by external parties.

To engage in prediction registration, the researcher undertakes a three-step process (figure 4). *First, the researcher makes predictions related to the universe of known cases within a theory's scope conditions* (Step 1). The researcher identifies as many cases as possible within the theory's scope conditions—that is, cases relevant to the theory—to avoid "cherry-picking" intentionally or unintentionally easy-to-predict cases. It is important that the predictions include the parameters specified in the Prediction Registration Template with both the values on the theory's independent variables and the outcome values. Researchers also can include conditional variables in their registration, stating in which scenarios the independent variable is expected to influence the outcome. These conditional variables might be general by stating the effect on predictions if a given condition occurs in any case or applied to particular cases by stating the effect on predictions for the case of interest.

Second, the researcher registers their predictions on OSF Registries or a similar registration site (Step 2). OSF Registries provides an easily accessible record that logs any changes made to posted content. The log of changes allows researchers to make updates transparently to the registered predictions (Step 2b). The researcher may want to make updates if a new case enters into the scope of the theory. Or, if an independent-variable value of an existing prediction changes in the real world, the researcher may update the predictions and specify a newly predicted outcome (all while maintaining a record of the original prediction).

Third, the researcher and/or a third-party analyzes the prediction results (Step 3). Researchers and/or third parties including other scholars or policy makers calculate the results using the pre-specified metric for scoring predictions, which could be—among other metrics—a raw count for binary predictions or Brier Scores for probabilistic predictions. It is advised that the results are compared to an alternative baseline such as other theories, predictions from large language models, or 50–50 chance. When independent-variable values manifest as different from those specified in the registration, it should be documented in the results analysis but *not* counted for or against the theory's predictive power. Theories are contingent on independent-variable values; therefore, this says little about their predictive power if the predictions are based on different values than those that manifested.

The evaluation stage also offers the opportunity to account for outcomes that did not occur as predicted (Step 3b). Having a verifiable record of the prediction incentivizes scholars to not simply "look the other way" at incorrect predictions, and it gives the researcher three options: calibrate the theory, leave the theory unmodified, or discard the theory. Researchers might calibrate the theory to add or omit variables that are deemed influential for explaining the missed predictions. They might conclude that leaving the theory unmodified is the best approach if calibrating the theory is believed not worth the tradeoff of potentially making it more complex. Similarly, the theory could be left unmodified if missed predictions are the result of measurement issues, in which case the researcher would adjust their measurement strategy rather than the theory. In some cases, it might be necessary to discard or shelve a theory if, for example, data do not exist to measure the additional variables needed to improve the theory's accuracy.

CHALLENGES TO PREDICTION

There are four main counterarguments to testing the predictive power of theories. First, scholars have long argued that prediction of complex human behaviors is too difficult empirically (Bernstein et al. 2000; Ward 2016). However, the fact that superforecasters consistently make more accurate predictions than non-superforecasters suggests that predictive accuracy can be learned (Horowitz et al. 2019; Tetlock and Gardner 2015). For his part, Miller (2022) finds that the US Government became highly accurate at predicting nuclear proliferation, eventually achieving correct assessments in 80% of cases. An evaluation of US presidential-election predictions also found that they have become increasingly accurate over time (Cuzán 2020). From the 12 Fomin et al. (2021) articles with at least one falsifiable prediction, I extracted 23 predictions in which it was plausible to infer whether they were correct or incorrect; from those predictions, 18 were correct—a success rate of 78%.[7]

https://doi.org/10.1017/S1049096524000404 Published online by Cambridge University Press

Figure 4
Prediction Registration Process

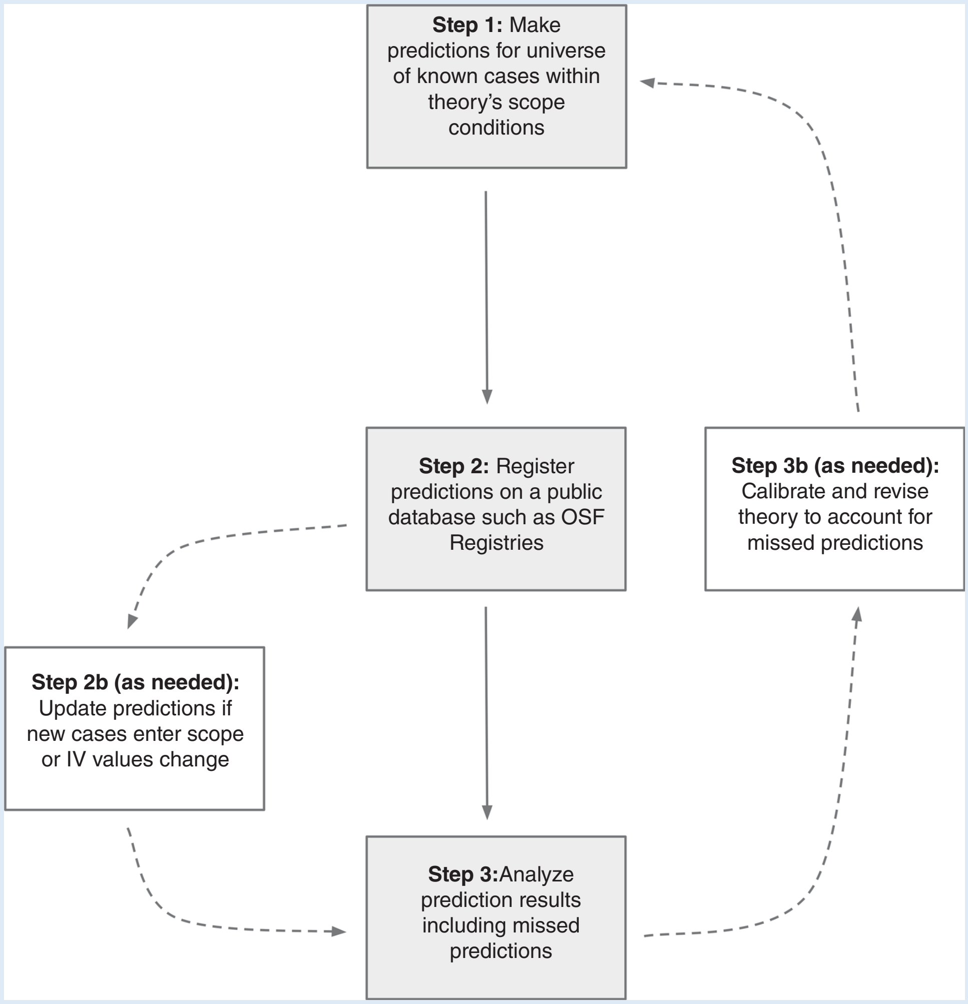

A second counterargument is that many theories aim to predict rare events, which makes it difficult to obtain statistical traction on predictive power. If a scholar posits that a theory predicts, for example, a 70% chance of an outcome, we would need predictions of many cases to identify whether the probabilities inferred from the theory are accurate. Although the predictive power of a theory would be difficult to evaluate in such cases, predictions about highly consequential cases (e.g., war onset) remain useful from a policy perspective. The fact that requests for such predictions appear on "snap polls" of IR scholars suggests demand for their forecasts around these events (TRIP 2022). Prediction registration also mitigates the difficulty of evaluating the predictive power of theories related to rare events because researchers can add other predictions for a given theory to its registry site when a new case comes within scope. Thus, all of their predictions using the theory would be aggregated and could be evaluated together rather than as disparate, one-off predictions.

An incorrect prediction is an opportunity to identify new variables, more precise scope conditions, and other factors that can strengthen a theory. In this sense, registered predictions are a win-win for those who seek to build predictively powerful theories: either the predictions prove to be correct or incorrect predictions reveal paths to improve the theory.

Third, many predictions have policy relevance (Schneider, Gleditsch, and Carey 2010); thus, predictions that policy makers

https://doi.org/10.1017/S1049096524000404 Published online by Cambridge University Press

subsequently act on could hypothetically change the outcome of events that the theory attempts to predict and mask a given theory's predictive accuracy. This problem would be a manifestation of Goodhart's Law, which states that "when a measure becomes a target, it ceases to be a good measure" (Stumborg et al. 2022). A theory shown to have predictive power could incentivize policy makers to either work toward or avoid the outcome based on the prediction. The potential for this confounder, however, is limited by the fact that prediction registration involves scholars specifying the values on the independent variables that drive a predicted outcome. If the value of the independent variable changes, including due to policy makers leveraging a prediction, the given case would not be scored for or against the theory's predictive power.

Fourth, the scoring of prediction results can become the subject of debate that could complicate the assessment of the predictive power of theories (e.g., Caplan 2018). Prediction registration, however, limits researchers' degrees of freedom in this regard because they specify *ex-ante* how they will score their predictions. Thus, researchers cannot select the metrics that cast their theory's predictive power in the light most desirable to them after the outcome is known.

PROMOTING PREDICTION REGISTRATION

What incentives do scholars have to adopt prediction registration? The "fuzziness" of many predictions in political science to date suggests that a temptation exists to make predictions that cannot be falsified. It could be argued that the expected professional gains from demonstrating that a theory has predictive power are much lower than the expected losses in credibility from inaccurate predictions. The asymmetry in incentives, however, also rings true with respect to explanatory theory testing—and the rigor of explanatory theory testing in political science has increased markedly in recent decades. Testing predictive power can follow suit. Previously, few scholars used registration for experiments; however, once they were introduced into the discipline, registrations have become standard practice, especially for experimental studies.

Even a prediction that turns out to be incorrect is much preferred to making no prediction at all if a theory is claimed to have predictive power. As Lakatos (1970) observed, theories subjected to more scrutiny are more likely to encounter data that do not fit the theories than those not subjected to any scrutiny. An incorrect prediction is an opportunity to identify new variables, more precise scope conditions, and other factors that can strengthen a theory. In this sense, registered predictions are a win-win for those who seek to build predictively powerful theories: either the predictions prove to be correct or incorrect predictions reveal paths to improve the theory.

Similar to encouraging preregistration of experimental studies, professional incentives could accelerate the adoption of rigorous evaluations of predictive power. The strongest professional incentive likely would be increased potential for publication. To this end, journal editors might encourage registration for studies that claim to have theories with predictive power. The demands on journal editors already are substantial; therefore, any registration standards should align with existing journal processes on study preregistration that many editorial boards already have in place for experimental studies. Procedurally, prediction registration would not be a radical departure from existing editorial processes but rather would only expand what has become known as the "preregistration revolution" that already has been established in the social sciences (Nosek et al. 2018).

ACKNOWLEDGMENTS

For helpful input, I thank Cullen Nutt, discussants and panelists of a 2022 American Political Science Association panel, the *PS: Political Science & Politics* editorial team, and three anonymous reviewers. I also thank Kathleen Murphy for research assistance.

DATA AVAILABILITY STATEMENT

Research documentation and data that support the findings of this study are openly available at the *PS: Political Science & Politics* Harvard Dataverse at https://doi.org/10.7910/DVN/ZGDHOU.

CONFLICTS OF INTEREST

The author declares that there are no ethical issues or conflicts of interest in this research. ■

NOTES

1. A third category involves using rational-choice game-theory models (Bueno de Mesquita 2011; Schneider, Gleditsch, and Carey 2011).
2. Fomin et al. (2021, 8) define predictions as "any research that has statements about future states of affairs as its key findings." They coded articles included in the TRIP Journal Article Database (Maliniak et al. 2018; TRIP 2020).
3. With important exceptions (e.g., Hegre 2013), predictions also rarely are followed by systematic evaluations of their accuracy.
4. Given that my claim is that the predictions tend to lack falsifiability, I err toward coding predictions as falsifiable in ambiguous cases to avoid overstating the claim. The replication data include the coding for these articles (Miller 2024).
5. Eighty-two scholars were identified as having this research focus through a search for relevant terms in the Twitter biographies of the 1,236 political scientists listed in Bisbee, Larson, and Munger (2020).
6. The replication data include the coding for all tweets (Miller 2024).
7. This success rate would be inflated if scholars selected easy-to-predict cases rather than predicting outcomes for the universe of cases within the theories-scope conditions—a shortcoming that prediction registration would mitigate.

REFERENCES

Bernstein, Steven, Richard Lebow, Janice Gross Stein, and Steven Weber. 2000. "God Gave Physics the Easy Problems: Adapting Social Science to an Unpredictable World." *European Journal of International Relations* 6 (1): 43–76.

Bisbee, James, Jennifer Larson, and Kevin Munger. 2020. "#polisci Twitter: A Descriptive Analysis of How Political Scientists Use Twitter in 2019." *Perspectives on Politics* 20 (3): 1–22.

Bowlsby, Drew, Erica Chenoweth, Cullen Hendrix, and Jonathan D. Moyer. 2020. "The Future Is a Moving Target: Predicting Political Instability." *British Journal of Political Science* 50 (4): 1405–17.

Bueno de Mesquita, Bruce. 2011. *The Predictioneer's Game: Using the Logic of Brazen Self-Interest to See and Shape the Future.* New York: Random House.

Caplan, Bryan. 2018. "Franck on the French Terrorism Bet." *EconLog* [blog], May 30. www.econlib.org/archives/2018/05/franck_on_the_f.html.

Cederman, Lars-Erik, and Nils B. Weidmann. 2017. "Predicting Armed Conflict: Time to Adjust Our Expectations?" *Science* 335: 474–76.

Cuzán, Alfred G. 2020. "The Campbell Collection of Presidential Election Forecasts, 1984–2016: A Review." *PS: Political Science & Politics* 54 (1): 99–103.

Dassonneville, Ruth, and Charles Tien. 2020. "Introduction to Forecasting the 2020 US Elections." *PS: Political Science & Politics* 54 (1): 47–51.

Fomin, Ivan, Konstantin Kokarev, Boris Ananyev, Nikita Neklyudov, Anzhelika Bondik, Pavel Glushkov, Aliya Safina, et al. 2021. "International Studies in an Unpredictable World: Still Avoiding the Difficult Problems?" *European Journal of International Relations* 27 (1): 3–28.

Friedman, Milton. 1966. "The Methodology of Positive Economics." In *Essays in Positive Economics*, ed. Milton Friedman, 3–43. Chicago: University of Chicago Press.

https://doi.org/10.1017/S1049096524000404 Published online by Cambridge University Press

Gleditsch, Kristian Skrede. 2022. "One Without the Other? Prediction and Policy in International Studies." *International Studies Quarterly* 66 (3): 1–8.

Grimmer, Justin, Margaret E. Roberts, and Brandon Stewart. 2021. "Machine Learning for Social Science: An Agnostic Approach." *Annual Review of Political Science* 24: 395–419.

Hegre, Håvard. 2013. "Observed vs. Predicted Conflict, 2012." *Håvard Hegre* [blog], November 7. https://havardhegre.net/forecasting/how-well-do-we-predict/observed-vs-predicted-conflict-2012.

Hegre, Håvard, Joakim Karlsen, Håvard Mokleiv Nygård, Håvard Strand, and Henrik Urdal. 2013. "Predicting Armed Conflict, 2010–2050." *International Studies Quarterly* 57 (2): 250–70.

Horowitz, Michael, Brandon M. Stewart, Dustin Tingley, Michael Bishop, Laura Resnick Samotin, Margaret Roberts, Welton Chang, Barbara Mellers, and Philip Tetlock. 2019. "What Makes Foreign Policy Teams Tick: Explaining Variation in Group Performance at Geopolitical Forecasting." *Journal of Politics* 81 (4): 1388–404.

Jacobs, Alan M. 2020. "Pre-Registration and Results-Free Review in Observational and Qualitative Research." In *The Production of Knowledge: Enhancing Progress in Social Science*, ed. Colin Elman, John Gerring, and James Mahoney, 221–64. New York: Cambridge University Press.

King, Gary, Robert Keohane, and Sidney Verba. 1994. *Designing Social Inquiry*. Princeton, NJ: Princeton University Press.

Lakatos, Imre. 1970. "Falsification and the Methodology of Scientific Research Programmes." In *Criticism and the Growth of Knowledge: Proceedings of the International Colloquium in the Philosophy of Science, London*, 1965, ed. Alan Musgrave and Imre Lakatos, 91–196. Cambridge, UK: Cambridge University Press.

Maliniak, Daniel, Susan Peterson, Ryan Powers, and Michael J. Tierney. 2018. "Is International Relations a Global Discipline? Hegemony, Insularity, and Diversity in the Field." *Security Studies* 27 (3): 448–84.

Mearsheimer, John J. 2015. "Why Is Ukraine the West's Fault? Featuring John Mearsheimer." University of Chicago [YouTube], September 25. www.youtube.com/watch?v=JrMiSQAGOS4.

Miller, Andrew Cesare. 2024. "Replication Data for 'Registering Theory-Based Predictions in Political Science.'" *PS: Political Science & Politics*.. DOI:10.7910/DVN/ZGDHOU.

Miller, Nicholas L. 2022. "Learning to Predict Proliferation." *International Organization* 76 (2): 487–507.

Montgomery, Alexander H., and Scott D. Sagan. 2009. "The Perils of Predicting Proliferation." *Journal of Conflict Resolution* 53 (2): 302–28.

Nosek, Brian A., Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. 2018. "The Preregistration Revolution." *Proceedings of the National Academy of Sciences* 115 (11): 2600–606.

OSF Registries. 2022. "Open Science Foundation." https://osf.io/registries.

Pilster, Ulrich, and Tobias Böhmelt. 2014. "Predicting the Duration of the Syrian Insurgency." *Research & Politics* 1 (2). https://journals.sagepub.com/doi/epub/10.1177/2053168014544586.

Popper, Karl. 2002. *The Logic of Scientific Discovery*. London: Routledge.

Schneider, Gerald, Nils Petter Gleditsch, and Sabine Carey. 2010. "Exploring the Past, Anticipating the Future." *International Studies Review* 12 (1): 1–7.

Schneider, Gerald, Nils Petter Gleditsch, and Sabine Carey. 2011. "Forecasting in International Relations: One Quest, Three Approaches." *Conflict Management and Peace Science* 28 (1): 5–14.

Stumborg, Michael F., Timothy D. Blasius, Steven J. Full, and Christine A. Hughes. 2022. "Goodhart's Law." *Center for Naval Analysis*. www.cna.org/reports/2022/09/goodharts-law.

Tetlock, Philip E., and Dan Gardner. 2015. *Superforecasting: The Art and Science of Prediction*. New York: Crown Publishing.

TRIP. 2022. "Snap Polls." *Teaching, Research, and International Policy Project*. https://trip.wm.edu/data/our-surveys/snap-polls.

TRIP. 2020. "TRIP Journal Article Database." *Teaching, Research, and International Policy Project*. https://trip.wm.edu/dashboard/journal-articles.

Ward, Michael D. 2016. "Can We Predict Politics? Toward What End?" *Journal of Global Security Studies* 1 (1): 80–91.

https://doi.org/10.1017/S1049096524000404 Published online by Cambridge University Press