Relationships among Rivals (RAR): A Framework for Analyzing Contending Hypotheses in Process-Tracing Research

Sherry Zaks
University of California, Berkeley
sherry.zaks@berkeley.edu

ABSTRACT

Methodologists and substantive scholars alike agree that one of process tracing's foremost contributions to qualitative research is its capacity to adjudicate among competing explanations of a phenomenon. Existing approaches, however, only provide explicit guidance on dealing with mutually exclusive explanations, which are exceedingly rare in social science research. I develop a tripartite solution to this problem. The Relationships among Rivals (RAR) framework (1) introduces a typology of relationships between alternative hypotheses, (2) develops specific guidelines for identifying which relationship is present between two hypotheses, and (3) maps out the varied implications for evidence collection and inference. I then integrate the RAR framework into each of the main process-tracing approaches and demonstrate how it affects the inferential process. Finally, I illustrate the purchase of the RAR framework by reanalyzing a seminal example of process-tracing research: Schultz's (2001) analysis of the Fashoda Crisis. I show that the same evidence can yield new and sometimes contradictory inferences once scholars approach comparative hypothesis testing with this more nuanced framework.

1. INTRODUCTION: THE CHALLENGE OF RIVAL HYPOTHESES

Process tracing is widely considered a powerful inferential tool in social science.[1] Of its many contributions to qualitative research, one of the most universally touted is its capacity to help scholars adjudicate among competing explanations of a phenomenon (Bennett 2010). Critical adjudication of rival hypotheses is important for fully assessing causal claims and guarding against confirmation bias (George & Bennett 2005, Hall 2006). As such, nearly every template for conducting process-tracing research enjoins researchers to begin by "casting the net widely for alternative explanations" (Bennett & Checkel 2015, 18).[2] However, those who advocate "casting the net widely," stop short of telling us how to deal with these alternatives once we reel them in. Consequently, despite its goals, the process-tracing scholarship suffers from a critical limitation: its treatment of rival hypotheses is fundamentally incomplete.

A deeper examination of the literature reveals the source of this shortcoming. Due either to the assumption that supporting one hypothesis necessarily weakens another (as in Collier (2011)), or to self-conscious bracketing in the context of providing more concise introduction to Bayesian inference (as in Bennett (2015, 278)), current process-tracing frameworks only provide explicit guidance on dealing with mutually exclusive explanations. Mutual exclusivity, however, is a strong modeling assumption; and empirically, it is more often the exception than the rule. Competing explanations may exhibit a variety of relationships to the main hypotheses, each of which has distinct implications for collecting evidence and drawing inferences. Although some methodologists have acknowledged the existence of non-exclusive relationships among rivals (Rohlfing 2012, Rohlfing 2014, Bennett 2014, Bennett & Checkel 2015), questions of both how to identify those relationships and what their consequences are for causal inference

[1]The scope of this article is limited to process tracing as it is used for theory-testing purposes. For discussions of other uses of process tracing see Kay & Baker (2015).

[2]See also (George & Bennett 2005, Brady 2006, Bennett 2010, Collier 2011).

have been side-barred in favor of advancing other aspects of the method.

The stakes of filling this gap are quite high. An unjustified assumption of mutual exclusivity can lead researchers to hastily rule out explanations that work in addition to the main hypothesis. Essentially, this error is the qualitative equivalent of omitted variable bias: it artificially inflates the importance of one explanation at the expense of another. Furthermore, if two explanations *are* mutually exclusive, researchers do not need to do much legwork to adjudicate between them. Once a researcher finds evidence in favor of one hypothesis, ruling out the other is only a matter of mathematical logic: if the two cannot work together and one is true, the other is invalidated by definition. In contrast, if two explanations can simultaneously bring about an outcome, evidence in favor of one does not automatically invalidate the other. Researchers must conduct a separate search for counter-evidence specific to the rival theory they seek to discredit. Thus, the processes of evidence gathering and inference change as a function of the relationship between rival hypotheses; and researchers are long overdue for a comprehensive guide to this end.

To fill this critical gap, this article develops the Relationships among Rivals (RAR) framework: a comprehensive method for evaluating competing hypotheses in process-tracing research. I begin by reviewing the incomplete handling of rivals in the existing literature. Section 3 then introduces the RAR framework by (1) providing a typology of the possible relationships among alternative explanations, i.e. mutual exclusivity, coincidence, congruence, and inclusiveness; (2) presenting a step-by-step procedure for identifying each relationship; and (3) deriving their corresponding implications for collecting and evaluating evidence. Section 4 integrates the RAR framework into existing process-tracing approaches and identifies key modifications needed to accommodate this new insight. This section pays special attention to Bayesian process-tracing and the mathematical implications of expanding Bayes' rule to accommodate non-exclusivity. Section 5 builds on the foundations of RAR to construct a pluralistic template for conducting process tracing. Finally, Section 6 illustrates the value of RAR by applying it to a seminal example of process tracing: Schultz's (2001) *Democracy and Coercive Diplomacy*. Another example, Tannenwald's (2007) *The Nuclear Taboo*, is available online in Appendix B.[3] I show that despite both Schultz's and Tannenwald's analytic rigor, the new framework reveals new—and sometimes contradictory—insights overlooked by both authors.

The RAR framework makes two global contributions to the process-tracing scholarship beyond the added nuance described above. First, by providing researchers with the tools for identifying and working with relationships among rival hypotheses, this approach finally brings the *procedures* of process tracing in line with one of its foremost *priorities*: critical adjudication of competing explanations. Second, the RAR framework seamlessly integrates into all existing process-tracing approaches: the use of analytic narratives (George & McKeown 1985, Van Evera 1997, Collier, Brady & Seawright 2010),[4] the crisp set-theoretic approach (Mahoney 2012, Goertz & Mahoney 2012, Blatter & Haverland 2012), and the procedures of Bayesian inference (Beach & Pedersen 2012, Bennett 2014, Bennett 2015).

2. EXISTING CONCEPTUALIZATIONS OF RIVAL HYPOTHESES

Process-tracing scholars have made crucial inroads in the advancement of the method, from the codification of standards for historical case analysis (Van Evera 1997), to elucidating the underlying logic of process tracing by mapping it onto Bayesian inference (Bennett 2014). Notwithstanding the value of each innovation, scholars have side-stepped the development of guidelines for one of the most important priorities: critical examina-

[3] [Link to online appendix.]

[4] I use "analytic narratives" to refer to the classical approach to process tracing in the vein of George (1979), Van Evera (1997), and George & Bennett (2005), among others. This term should not be confused with the usage of "analytic narrative" in Bates et al. 1998.

tion of rival hypotheses. This section shows that despite some progress in conceptualizing the relationships among rival hypotheses, existing frameworks are fundamentally incomplete.

Collier (2011) presents the first process-tracing framework that implicitly acknowledges how relationships among rival hypotheses may affect inferences. His framework is an expansion of Bennett's (2010) 2×2 table, which sorts evidence for the main hypothesis (H_M) according to its uniqueness and certainty.[5] Collier proposes that with each piece of evidence bearing on H_M, researchers can evaluate the plausibility of both the main hypothesis and rival hypotheses (H_R). Collier's (2011) framework is reproduced in Table 1.

Table 1: The Four Process Tracing Tests†

	Sufficient to Affirm Causal Inference	
Necessary to Affirm Causal Inference	**1. Straw-in-the-Wind** **Passing:** Affirms relevance of hypothesis, but does not confirm it. **Failing:** Hypothesis is slightly weakened, though not eliminated. **Implications for Rival Hypotheses:** **Passing:** *slightly* weakens them. **Failing:** *slightly* strengthens them.	**3. Smoking-Gun** **Passing:** Confirms hypothesis. **Failing:** Hypothesis is somewhat weakened, though not eliminated. **Implications for Rival Hypotheses:** **Passing:** *substantially* weakens them. **Failing:** *somewhat* strengthens them.
	2. Hoop **Passing:** Affirms relevance of hypothesis, but does not confirm it. **Failing:** Eliminates Hypothesis. **Implications for Rival Hypotheses:** **Passing:** *somewhat* weakens them. **Failing:** *somewhat* strengthens them.	**4. Doubly-Decisive** **Passing:** Confirms hypothesis and eliminates others. **Failing:** Eliminates Hypothesis **Implications for Rival Hypotheses:** **Passing:** *eliminates* them. **Failing:** *substantially* strengthens them.

†Source: Collier (2011), who adapts the table from Bennett (2010).

While Collier takes a step towards recognizing relationships between H_M and H_R, his template is limited because it treats mutual exclusivity as the *only* relationship researchers

[5] Bennett's (2010) table draws on the criteria of Van Evera's (1997) four ideal-type tests based on the same dimensions. Uniqueness refers to evidence that is so specific to one theory that finding it is sufficient to confirm the hypothesis. Certainty refers to evidence that must be found for the theory to be true.

will encounter. To be sure, competing hypotheses are sometimes mutually exclusive. If, for example, one cannot observe both H_M and H_R in a given case, then—as per Table 1—evidence *in favor* of H_M will indeed *undermine* H_R. If, instead, both explanations together bring about the outcome, evidence validating H_M may have a diverse range of implications on H_R. Collier's (2011) framework only provides accurate inferences if H_M and H_R exhibit mutual exclusivity, and this approach is therefore seriously incomplete.

Rohlfing (2012; 2014) takes a step forward in conceptualizing a wider variety of relationships and potential problems that arise when mutual exclusivity is presumed without justification. His most recent work draws on the conceptual distinctions and implications previously discussed by Zaks (2012) to develop a framework of case-selection principles (Rohlfing 2014, 30). Although Rohlfing acknowledges varied relationships, his work is not intended as a full-scale discussion of relationships among rivals—rather, the relationships are only minimally developed in the service of a framework for case-selection. Thus, a full-scale solution to the incomplete treatment of rival hypotheses must push much farther on three fronts. First, Rohlfing's exclusive focus on case-selection only begins to tap the broad implications non-exclusivity has for research. Second, the framework leaves open the question of how to identify the range of relationships. Third, it lacks a discussion of how different relationships affect the inferences.

The challenges of conceptualizing and analyzing relationships among rival explanations also arise in Bayesian process tracing. According to its proponents, the recent move toward Bayesian inference represents a synthesis of the underlying logic of process tracing with a formal procedure codifying what researchers have been doing implicitly all along (Bennett 2009, Beach & Pedersen 2012, Bennett 2014, Fairfield & Charman 2015, Bennett & Checkel 2015). Relationships among rival hypotheses occupy a pivotal role in both the mathematical procedures and the interpretation of Bayesian analysis. However, since Bayesian process tracing is still in its infancy, scholars pioneering this approach have had to set aside explorations of non-exclusivity in order to provide more digestible introductions to Bayesian logic (Bennett 2014). As a result, examples used in the literature assume mutual exclusivity for the sake of illustrating the method, yet scholars have explicitly acknowledged that this relationship does not always hold (Bennett 2014, 47). While Bayes' rule would have to be expanded to accommodate non-exclusivity (2015), scholars in this tradition have not yet derived or discussed the expansion.

To varying degrees, the process-tracing literature acknowledges the existence of relationships among rival hypotheses. Whether due to an ontological framing (as in Collier 2011), or a simplifying assumption (as in Bennett 2014, 2015; Fairfield and Charman 2015), the literature only provides explicit guidance on dealing with mutually exclusive rivals. Stepping back, this assumption explains scholars' affinities for illustrating process tracing with detective stories (Van Evera 1997, Collier 2011) and epidemiological examples (Freedman 2010, Humphreys & Jacobs 2015), where comorbidity is the exception, not the rule. This picture of research is pleasant, tidy, and usually inaccurate.

The treatment of relationships among rival hypotheses is a pivotal, yet underdeveloped modeling assumption in process tracing. In light of this gap, three central questions demand attention: (1) What relationships exist beyond mutual exclusivity? (2) How can researchers identify the relationship between two hypotheses in a given case? (3) How does the relationship affect the process of inference?

3. RELATIONSHIPS AMONG RIVALS (RAR): TYPOLOGY AND APPLICATION

"Casting the net widely for alternative explanations" is important; but this exercise alone cannot enhance the quality of research without a comprehensive guide to identifying the relationships between rival hypotheses and assessing how those relationships affect our inferences. To fill this gap, this section presents what I call the Relationships among

Rivals (RAR) framework. The contribution of the framework is threefold. First, it develops a novel typology—depicted in Table 2—of four possible relationships among hypotheses: mutual exclusivity, coincidence, congruence, and inclusiveness. Second, it offers guidelines for identifying the relationship present for a given set of explanations. Third, it derives a comprehensive set of implications for drawing inferences on the basis of that relationship. I begin by grounding the logic of the typology in a foundational concept from probability theory, the sample space; then I introduce each component of the RAR framework in turn.

Table 2: Relationships Among Alternative Hypotheses

Relationship	**Description**
Mutually Exclusive	Hypotheses are completely disjoint. Corroborating one necessarily undermines the validity of the other. *Outcome* mutual exclusivity occurs when two two theories make divergent predictions on the outcome. *Evidentiary* mutual exclusivity occurs when one theory requires evidence that would undermine the other.
Coincident	Hypotheses are independent of one another. The validity of one neither corroborates nor undermines the other. Both may contribute to the phenomenon or one may give rise to the phenomenon in one case, while the other does in another case.
Congruent	Hypotheses are similar in the type of evidence they require and they may work also work together to bring about the outcome. Corroborating one hypothesis simultaneously lends support to the other.
Inclusive	Inclusive hypotheses are a special case of congruence in which evidence reveals that one hypothesis constitutes an extension of an alternative theory that was originally presented as a rival.

3.1. *"Rival Spaces" – The Logic of the Typology*

The exercise of "casting the net widely for alternative explanations" is analogous to a foundational concept in probability and set theory: building a sample space.[6] The difference is trivial: in probability theory, sample spaces are populated by a set of potential outcomes; in process tracing, sample spaces—or what we might call "rival spaces"—would be populated by a set of competing hypotheses explaining a phenomenon.

Examining the properties of the sample space reveals what process-tracing research is missing: in many cases, statisticians cannot accurately compute the likelihood of one event relative to another without specifying the relationship among events, which can take numerous forms (e.g. exclusivity, independence, and dependence).[7] If a statistician does not know the relationships between events (or, alternately, assumes that one relationship is true all of the time), any attempt to calculate probability is liable to generate inaccurate estimates. Recasting competing hypotheses in these terms makes clear that without proper specification of relationships, process-tracing conclusions are susceptible to the same errors. The following sections draw on this insight to elaborate the variety of relationships and derive the corresponding implications for inference.

3.2. *Mutually Exclusive Hypotheses*

The first relationship researchers may encounter is mutual exclusivity, in which two explanations cannot simultaneously be valid in a single case. Mutual exclusivity comes in three flavors. *Outcome* mutual exclusivity occurs when two theories make consistently di-

[6]The sample space, Ω, describes the set of all possible outcomes, and the corresponding probability space is the assignment of probabilities to each event ω_i. For example, the roll of a single fair die can be represented by a sample space with six disjoint outcomes $(\omega_1, \omega_2, \ldots, \omega_6)$ of equal probability, $p(\omega_i) = \frac{1}{6}$.

[7]Mutual exclusivity is defined as $P(A \cap B) = 0$ (i.e. the two events can never co-occur), independence describes is defined as $P(A \cap B) = P(A) * P(B))$ (i.e. the two events may co-occur, but neither affects the other), and dependence is defined as $P(A \cap B) = P(A) * P(B|A))$ (i.e. the likelihood of one event depends upon the occurrence of the other).

vergent predictions on the outcome, whereas *evidentiary* mutual exclusivity occurs when one theory requires evidence that would undermine the other (i.e. two theories operate by mutually exclusive paths). Finally, two theories may exhibit *conditional* mutual exclusivity, in which a scope condition or an idiosyncrasy of the case forces two explanations to be mutually exclusive in some instances but not in others.[8] This relationship is the only one in which evidence in favor of one hypothesis logically subverts an alternative hypotheses.

Mutual exclusivity is a strong modeling assumption for which the burden of proof is on researchers. In social science phenomena, mutual exclusivity is rare and difficult to establish. It is not the default state of the world, notwithstanding the wealth of substantive literature operating on this assumption. In the literature on rebellion, for instance, *greed* and *grievance* are often posited as "rival" explanations of insurgent participation (Collier & Hoeffler 2004), as though someone cannot be both sad and acquisitive.

The most straightforward case of mutual exclusivity is the null hypothesis. After all, if one's theory is that X affects Y, it would be impossible for X to simultaneously *not* affect Y. Brady's (2010) examination of Lott's (2000) argument regarding the effect of the "early call" of the 2000 U.S. presidential race is a clear case of using process tracing to evaluate one hypothesis against the (mutually exclusive) null. Lott argues that when the Florida media called the race in favor of Al Gore after the polls closed in the main peninsula, potential Republican voters in the Florida panhandle (where polls were still open) were dissuaded from voting. In response to Lott's claims that the media suppressed approximately 100,000 votes, Brady systematically performs a series of hoop tests, providing evidence against each one in turn. In effect, Brady's analysis lends

[8]The respective predictions of Newtonian and quantum mechanics illustrate this case well. The kinematic and dynamic predictions of Newtonian mechanics are accurate for a wide range of phenomena, but the laws break down at the atomic level. For objects smaller than 10^{-9}m, quantum mechanics makes different, and more accurate predictions about how particles behave. Thus, at the atomic level, Newtonian and quantum mechanics exhibit mutual exclusivity; yet, the predictions of both theories converge for larger phenomena, which suggests that their mutual exclusivity is conditional on size.

support to the null.

Assuming mutual exclusivity when it does not hold can prove detrimental for both present and future research. First, treating evidence in favor of H_M as though it necessarily counts against H_R can lead researchers to hastily rule out explanations that may work *in conjunction* with the main hypothesis. Scholars drawing on this flawed research are then susceptible to discounting the theoretical value of an explanation that was eliminated on the basis of an *assumption* rather than empirical evidence.[9]

3.3. *Coincident Hypotheses*

The second possible relationship between hypotheses is coincidence. Akin to statistical independence, explanations are coincident when the validity of one has no impact on that of the other. In substantive terms, coincident theories operate via sufficiently different mechanisms that both explanations could independently contribute to the outcome. As such, evidence in favor of one has no effect on the other. Consequently, the process of eliminating a coincident rival hypothesis involves collecting additional evidence that specifically invalidates the causal process of that explanation.

A prime example of coincident explanations comes from the democratic peace literature. In seeking an explanation for the phenomenon that democracies exhibit some immunity from fighting one another, scholars have proposed two models (Maoz & Russett 1993). The *normative* model posits that peaceful methods for *intra*state conflict resolution are externalized reliably enough to shape *inter*state conflict resolution. The *structural* approach contends that democratic institutions constrain leaders in a way that makes waging war difficult. Though both explanations arrive at the same outcome, they posit different causal stories. Indeed, as Maoz and Russett explicitly note "these two explana-

[9] In the Bayesian approach, for example, this problem would manifest as an artificially low prior on a given hypothesis.

tions are not mutually exclusive" (1993, 626). Thus, a piece of evidence corroborating a normative process in democratic decision-making does *not* preclude structural influences from playing a comparably strong role when explaining why two democracies did not go to war.

3.4. *Congruent and Inclusive Hypotheses*

The third possible relationship among "rival" explanations is *congruence*. Under congruence, hypotheses make similar predictions or operate via sufficiently similar mechanisms that evidence in favor of one also corroborates the other. In such cases, we would expect to find similar evidence under each theory. This is not to say that the two theories are indistinguishable—they may be congruent on some pieces of evidence, and just coincident on others. Alternately, the two theories may describe different levels of analysis, or one theory might describe an outcome that is sequentially prior to another. In cases of congruence, researchers must be especially careful to specify which pieces of evidence simultaneously support both theories, and what would be necessary to eliminate the rival from consideration.

Finally, inclusive hypotheses represent a special case of congruence in which one hypothesis represents a novel extension of an existing theory. Scholarship typically proceeds by first demonstrating the inadequacy of existing theories, and then proposing a novel explanation to fill the gaps. Sometimes, however, what begins as a separate theory may instead constitute a new dimension of an existing one—even when that framework is initially presented as a "competing" explanation. Either an assumption of mutual exclusivity or just a habit of dismissing alternative explanations can lead researchers to overlook cases in which new theories represent useful extensions of existing frameworks. The RAR approach provides guidance for finding evidence that suggests theoretical integration.

Researchers are potentially dealing with inclusive hypotheses when three conditions are met. First, the hypotheses must make similar or otherwise complementary predictions. In other words, the same input cannot generate contradictory output. Second, the hypotheses cannot operate via contradictory assumptions. If assumptions across two theories are dissimilar, scholars should ask if they can be reconciled; perhaps the original set of assumptions is less empirically valid than the new ones, so the latter should supplant the former. Third, hypotheses are potentially inclusive if the evidence presented in favor of one either supports or has no impact on the other and if no invalidating evidence is found for either hypothesis.

3.5. *Guidance for Identifying Relationships among Rivals*

The typology introduced above raises two important questions. How can scholars establish which relationship is present, and what are the implications for evaluating evidence? Since the relationship between two hypotheses affects both evidence gathering and interpretation, researchers need a clear guide to assess which relationship is present. Figure 1 presents a decision tree to guide researchers through the process of identifying relationships and drawing inferences. The first set of questions addresses broad theoretical considerations by asking whether the competing explanations make divergent predictions on the outcome. If one theory expects outcome Y, and another expects no Y (or high versus low levels of Y), they may be *candidates* for mutual exclusivity, but even predictive divergence is not sufficient to ensure it.

Answering "yes" to the first question then requires a scope check—asking whether the theories make different predictions only at distinct levels of analysis or under certain conditions. If divergence does not seem to be a function of other variables or scope

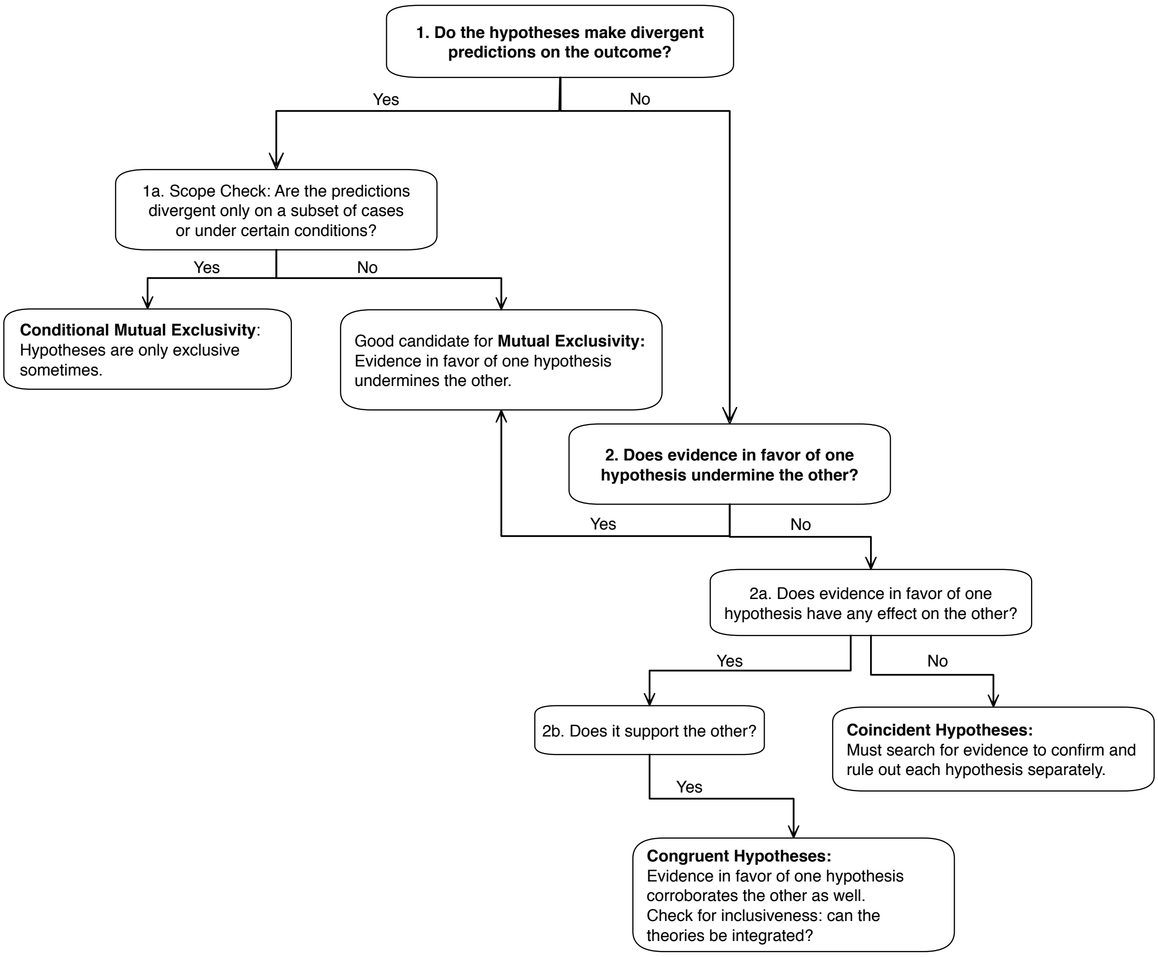

Figure 1: Identifying Relationships Among Rivals

conditions, then the two explanations are good candidates for outcome mutual exclusivity. If, however, a scope condition is driving the divergence, the two theories may exhibit conditional mutual exclusivity, in which they are exclusive under some conditions, but not others. Since mutual exclusivity may be contingent on scope, researchers should search for additional variables or conditions that affect a theory's validity in a given case, and they must be cautious when generalizing about the overall relationship between theories. Finally, if two hypotheses are conditionally mutually exclusive, researchers should note the driving factor and the boundaries of each theory.

The next set of questions concerns the nature evidence required under each explanation. Researchers should start by asking whether the two theories call for divergent types of evidence. Crucially, this question is not about asking whether researchers would expect to find *different* pieces of evidence from one theory to another, but rather whether they would expect to find *incompatible* pieces of evidence.[10] For example, if one theory is supported by the presence of E, yet another requires its absence, they would also be good candidates for mutual exclusivity (even if they make the same predictions about the outcome).

If two theories do not have contradictory requirements, researchers should then ask whether evidence in favor of one hypothesis has *any* effect on the validity of the other. If not, the two theories are likely coincident: they require different pieces of evidence both for corroboration and invalidation. Explanatory factors A and B may together bring about an outcome; consequently, verifying one does not invalidate the other. As such, to rule out a coincident alternative (say, theory B) researchers must search for additional evidence that specifically addresses B. Alternately, if the evidence in favor of one theory *does* have an effect on the other theory, researchers must then ask whether the effect is undermining or corroborating. In the former case, the two hypotheses become good candidates for evidentiary mutual exclusivity; in the latter case, the two hypotheses likely exhibit congruence (at least on that piece of evidence).

Finally, if two theories are congruent, researchers should ask whether they are candidates for inclusiveness. Does integrating them into a single theoretical framework provide more analytic leverage than either one separately? If two explanations make predictions about the same outcome and one explanation represents a new dimension or channel through which the prior may operate, the theories are good candidates for inclusiveness.[11] If their predictions or mechanisms are too different to reconcile into a single

[10] The distinction between expecting *different* and *divergent* evidence is potentially the source of many errors in process tracing. Two congruent theories may nonetheless need different types of evidence from one another to be verified. Too often, however, are researchers inclined to rule out one hypothesis because they found the unique kind of corroborating evidence in favor of the other.

[11] For an example of congruence in practice, see the Tannenwald analysis in the online appendix

framework, then they remain congruent.

The framework introduced in this section provides researchers with a tool for explicitly modeling and testing relationships among competing hypotheses. By disaggregating competing theories into their constituent mechanisms and assumptions, researchers can identify the precise dimensions on which two theories differ, thereby providing a more precise guide to the type of evidence needed to support or undermine a given explanation. While this framework represents a crucial advancement to the method, many previous innovations have given rise to internecine debates over how process tracing is best conducted—some of which go so far as to advocate one approach to the exclusion of others. Thus, contributions notwithstanding, the introduction of the RAR framework raises two key questions: (1) What are the implications for existing process-tracing approaches? and subsequently, (2) how should this framework be incorporated into the research process more broadly? The next two sections address these questions in turn.

4. IMPLICATIONS FOR EXISTING APPROACHES TO PROCESS TRACING

One of the most beneficial features of the RAR framework is that it is not approach-specific: it transcends debates over whether process tracing is best done with analytic narratives[12] (George & McKeown 1985, Van Evera 1997, Collier, Brady & Seawright 2010), crisp-set theory (Goertz & Mahoney 2012, Blatter & Haverland 2012), or Bayesian inference (Beach & Pedersen 2012, Bennett 2014, Bennett 2015). The RAR framework can be seamlessly integrated into each approach. This section first explores the brief implications of integrating the RAR framework into the analytic narrative approach and the crisp-set theoretic approach. The bulk of this section, however, is dedicated to integrating RAR

(Appendix B).

[12] It is worth noting that Bayesian scholars argue that the construction of analytic narratives is merely an implicit use of Bayesian inference. The purpose here is not to create sharp (let alone false) divides among approaches, but rather to examine the implications for practitioners.

into the mathematical procedures of formal Bayesian process tracing.

4.1. *Analytic Narratives*

The classical approach to process tracing—constructing analytic narratives on the basis of systematically collected evidence—must be expanded beyond Collier's (2011) framework to reflect the diverse implications of evidence when hypotheses are non-exclusive. Thus, while the general principles of this approach hold, when adjudicating among competing explanations, researchers should follow the inferential procedures laid out in Figure 1. Other than incorporating a more nuanced conception of rivals into the research procedure, no specific changes to the approach itself are required.

4.2. *Crisp Set-Theoretic Process Tracing*

The set-theoretic approach understands process tracing as a search for "necessary, sufficient, and jointly sufficient factors for an outcome" (Blatter & Haverland 2012, 24). Consequently, Goertz and Mahoney maintain that inferences must follow the rules of set-theoretic logic (2012, 13). Integrating RAR into this approach adds inferential nuance in two key areas. First, insights about the rarity of mutual exclusivity suggest that researchers should not be looking for the ideal "doubly-decisive" evidence in all cases. Doubly-decisive tests entail finding a single piece of evidence that simultaneously confirms H_M while undermining H_R. For doubly-decisive evidence to exist, three rare conditions must align at once: (a) the two theories must be mutually exclusive; (b) researchers must find a piece of evidence sufficiently strong to confirm the main hypothesis; (c) the piece of evidence must be on the same dimension that defines the exclusivity of the two theories.

The second implication for the crisp set-theoretic approach is that the non-exclusive relationships among rivals (coincidence, congruence, and inclusiveness) dovetail with the logic of INUS (jointly sufficient) conditions. Set-theoretic scholars routinely acknowledge that multiple conditions may work jointly to bring about an outcome (Mahoney & Goertz 2006, Mahoney 2008), yet this insight remains analytically separate from discussions of rival hypotheses.[13] The diverse set of relationships among rival explanations can help scholars in the set-theoretic camp further parse out the relationships among INUS conditions, which are simply non-exclusive rivals. RAR thus provides those using set-theoretic process tracing with a guide to identifying jointly sufficient conditions and a more nuanced way of describing how these conditions work together.

4.3. *Bayesian Process Tracing*

The Bayesian innovation represents the frontier of process-tracing research. Bayesian process-tracing asks a simple analytic question for each piece of evidence: *What is the probability that our main hypothesis (H_M) is correct, given that we searched for and found evidence K: $P(H_M|K=1)$?*[14] Formally, this question takes the following form:

$$P(H_M|K=1) = \frac{P(H_M)P(K=1|H_M)}{P(H_M)P(K=1|H_M) + P(\neg H_M)P(K=1|\neg H_M)}, \text{ where,} \tag{1}$$

- $P(H_M)$ represents the estimated probability that our hypothesis is correct before new evidence is found (i.e. the prior),
- $P(K=1|H_M)$ represents the probability of observing that piece of evidence when our hypothesis is correct (i.e. the likelihood),

[13] Mahoney (2012), for example, provides a detailed description of identifying joint sufficiency, and then separately poses the question, *"How are rival hypothesis eliminated?"*. This separate treatment suggests that scholars lack a conceptual framework for integrating alternative explanations that may work in conjunction with the main hypothesis.

[14] This notation draws on Humphrey and Jacobs (2015), which represents the successful search of a piece of evidence as $K=1$.

- $P(\neg H_M)$ represents the probability that our hypothesis is incorrect (this quantity is computed as $(1 - H_M)$),
- $P(K = 1|\neg H_M)$ represents the probability of finding that piece of evidence under a different hypothesis, and finally
- $P(H_M|K = 1)$ represents the updated (i.e. posterior) probability that our hypothesis is correct, given that we found evidence K.

Thus, for each piece of evidence (K), scholars must first surmise and justify these probabilities, and then input them into Bayes' rule to compute the updated confidence in H_M in light of the evidence.[15]

Due in part to its pivotal role in the refinement of qualitative inference and in part to the formalized procedures on which Bayesian process tracing relies, this approach warrants special consideration in light of the RAR framework. Specifically, since Bayes' rule includes terms that presume a disjoint sample space (i.e. comprising N mutually exclusive possibilities), it requires extensive modification when dealing with non-exclusive theories. Although Bayesian inference is sensitive to misspecification of relationships among competing explanations, the challenge of expanding Bayes' rule to accommodate non-exclusivity is practical, rather than conceptual.

Expanding Bayes' Rule to Accommodate Non-Exclusivity. This section explores the mathematical implications of computing posterior probabilities in cases of non-exclusivity. To motivate this discussion, consider the two rival spaces depicted in Figure 2. Computing the posterior probability $P(H_1|K = 1)$ in Figure 2(a)—where the rival space is clean and disjoint—follows the procedures outlined above. Complications arise, however, when researchers move into the more common, non-exclusive territory illustrated in Figure 2(b). If, for example, H_1 and H_2 are congruent (i.e. the same evidence is expected under both),

[15] Methodologists are engaged in a debate about the benefits and drawbacks of adopting this explicit mathematical approach. The debate itself is beyond the scope of this article. For a comprehensive introduction to the Bayesian approach and the debates surrounding it see, Bennett (2009, 2014, 2015), Beach & Pedersen (2012), Humphreys & Jacobs (2015), Fairfield & Charman (2015).

researchers must now evaluate the probability that the case is explained by the dark grey area, which represents the intersection $H_1 \cap H_2$. Using the equation above to compute $P(H_1|K = 1)$ when H_1 and H_2 are congruent would compromise every term in Bayes' rule. First, the analysis would be performed on the wrong area, thus rendering the prior incorrect. Second, the likelihood of finding K under H_1 is most likely lower than that of finding K in the intersection $(H_1 \cap H_2)$, thus rendering the likelihood incorrect. Third, the probability of finding K under H_2, $P(K = 1|H_2)$, wrongfully acts as a penalty term, thus rendering the denominator incorrect.

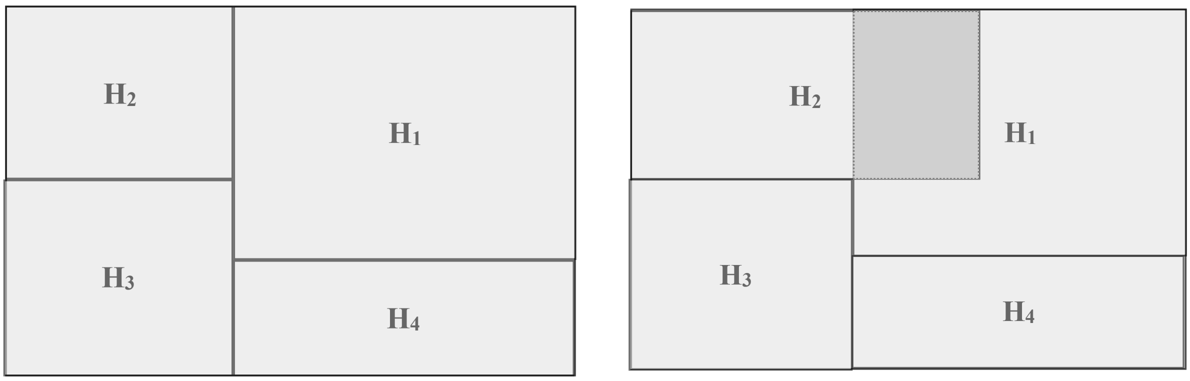

(a) Exclusive (Disjoint) Rival Space (b) Non-exclusive Rival Space

Figure 2: The Intuition behind Non-exclusivity

Thus, if a researcher believes a case is best explained by H_1 and H_2 together, Bayes' rule must be modified. For each piece of evidence, she must evaluate $P(H_1 \cap H_2|K = 1)$. Each term in the expanded equation requires numerous additional considerations, which I address in turn. The modified equation takes the following form:

$$P(H_1 \cap H_2|K = 1) = \frac{P(H_1 \cap H_2)P(K = 1|H_1 \cap H_2)}{[P(H_1 \cap H_2)P(K = 1|H_1 \cap H_2)] + \left[\left(1 - P(H_1 \cap H_2)\right)P(K = 1|\overline{H_1 \cap H_2})\right]}. \tag{2}$$

Rival Space and Priors. First, researchers are tasked with specifying a more complex rival space than the examples in the introductory literature. In addition to specifying priors $P(H_1)$ and $P(H_2)$, the burden is on the researcher to to specify the probability of cooccurrence: $P(H_1 \cap H_2)$ (which is not necessarily $P(H_1) \times P(H_2)$).[16] Rather, this quantity depends upon the researcher's beliefs about the frequency with which the two hypotheses work together. $P(H_1 \cap H_2)$ could be as small as 0, if the two are mutually exclusive; or it could be as large as the smaller probability if one is completely subsumed under the other. As with any quantity, the point the researcher chooses for $P(H_1 \cap H_2)$ must be justified theoretically and is subject to debate. The range is defined as follows:

$$0 \leq P(H_1 \cap H_2) \leq \min\big(P(H_1), P(H_2)\big). \tag{3}$$

Likelihoods. After deriving the priors—$P(H_1)$, $P(H_2)$, and $P(H_1 \cap H_2)$—researchers must specify three likelihoods. First, they must estimate $P(K = 1|H_1)$: the probability of finding evidence K if H_1 is true. Second, they must estimate the same quantity for H_2. Third, they must estimate the probability of finding K in the intersection, $P(K = 1|H_1 \cap H_2)$. How $P(K = 1|H_1 \cap H_2)$ is calculated depends on the brand of non-exclusivity. If H_1 and H_2 are coincident, and a given piece of evidence is likely under H_1, but effectively zero in H_2, $P(K = 1|H_1 \cap H_2)$ is given by

$$P(K = 1|H_1 \cap H_2) = P(K = 1|H_1)P(H_1 \cap H_2). \tag{4}$$

If, however, the two hypotheses are congruent and evidence K is expected with some positive probability under both H_1 and H_2, $P(K = 1|H_1 \cap H_2)$ is given by,

$$P(K = 1|H_1 \cap H_2) = 1 - \big[P(K = 0|H_1)P(K = 0|H_2)\big]. \tag{5}$$

Finally, the last term in the denominator, $P(K = 1|\overline{H_1 \cap H_2})$, represents the probability of observing K anywhere in the rival space except for the intersection of H_1 and H_2. Computing this term is cumbersome in its own rite. Building on our intuition from

[16] When estimating the coincidence of two statistically independent events, $P(A \cap B) = P(A) \times P(B)$.

Figure 2(b), we must compute the probability of observing K under each other alternative in the rival space and weight each by the area of the respective hypothesis (i.e. by the probability that hypothesis is true). Thus, $P(K = 1|\overline{H_1 \cap H_2})$ is given by,

$$\begin{aligned} P(K = 1|\overline{H_1 \cap H_2}) = &\left[P(H_1) - P(H_1 \cap H_2)\right] P(K = 1|H_1) + \\ &\left[P(H_2) - P(H_1 \cap H_2)\right] P(K = 1|H_2) + \\ &\left[1 - P(H_1 \cup H_2)\right] P(K = 1|\overline{H_1 \cup H_2}). \end{aligned} \tag{6}$$

For n hypotheses, the final term in Equation 6 is computed as,

$$[1 - P(H_1 \cup H_2)]P(K = 1|\overline{H_1 \cup H_2}) = \sum_{i=3}^{n} P(K = 1|H_i)p(H_i).^{17} \tag{7}$$

The key takeaway from this discussion is that Bayesian procedures can be expanded to accommodate non-exclusivity, but researchers will encounter a tradeoff between rigor and transparency on the one hand, and added complexity on the other. The Bayesian approach still has the benefit of forcing researchers to be explicit about their assumptions (Bennett 2015, Humphreys & Jacobs 2015). However, the number of quantities the researcher must posit and justify theoretically grows considerably as more hypotheses are thrown into the mix. The insights derived here will enable researchers to assess the tractability of Bayesian process tracing for a given project and to employ this approach more accurately should they choose to.

5. A PLURALISTIC TEMPLATE ON RAR FOUNDATIONS

The process-tracing literature is brimming with standards and ingredients of "good process tracing." Lacking, however, is a set of foundational guidelines that both direct the

[17] And this definition does not include the additional complication of overlap among *other* alternatives.

research process and allow researchers to adjudicate among existing approaches for the purposes of testing across various contexts. To motivate the need for a consistent, yet pluralistic process-tracing template, imagine telling a group of people the ingredients of a good baguette, allowing them to taste some good baguettes, informing them that the best baguettes have a crunchy exterior and a chewy interior, and then turning them loose in kitchens around the world. The results would not be delicious. But even a good recipe—painstakingly measured down to the grain of flour—that produces a perfect baguette in Paris, will not produce favorable results in Denver. Good chefs balance guidelines with context, and the best methodologists do the same.

This section provides a set of guidelines that should form the foundation of any process-tracing research, but encourages researchers choose the approach best suited to answering their question. The RAR framework is an apt foundation of a pluralistic template for a three reasons. First, it is not approach-specific. The importance of considering how hypotheses relate to one another transcends debates about specific inferences should be made, which might change based on the research question—or even from one piece of evidence to the next. Second, the framework provides a nuanced guide for gathering and evaluating evidence in each approach. Finally, the insights from the RAR framework reveal where challenges may arise in a given approach, which in turn provides researchers with a more comprehensive understanding of trade-offs.

5.1. *Step 1: Build the Rival Space*

Researchers should begin process tracing by constructing the rival space: an exhaustive set of alternative explanations for the phenomenon.[18] This suggestion echoes Platt's (1964) exposition on the procedure of *Strong Inference*, in which he urges researchers to

[18]Though Kay & Baker (2015) encourage researchers to begin research in a similar way, their recommendation omits a consideration of relationships among competing hypotheses.

"devise and write down our alternatives every day" (348). Whether represented graphically (as something akin to a Venn diagram as in Figure 3) or as a list, working explicitly with rival spaces prior to analysis can reveal potential problems[19] and can guide evidence collection and interpretation in a more directed way. In accordance with the RAR framework, the rival space should include (1) the main explanation and the alternatives, (2) the constituent hypotheses and assumptions that comprise each theory, and (3) the relationships among them. In a graphical depiction of the rival space, this task involves asking, *Do H_M and H_R overlap? If so, by how much? And on which pieces of evidence?*

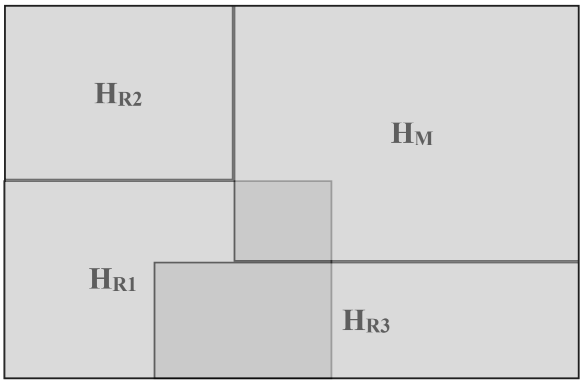

Figure 3: Graphical Representation of the Rival Hypothesis Space

5.2. *Step 2: Devise Tests*

In light of the RAR framework, researchers must devise two categories of tests. First, researchers must think about the type of evidence required to corroborate or invalidate each competing explanation. On this point, Waldner (2014) makes a compelling case for standardizing the use of directed causal graphs to specify mechanisms and to help

[19] In Appendix A (online), I demonstrate that process tracing conclusions (especially—though not limited to—conclusions in Bayesian process tracing) are particularly susceptible to bias when a piece of evidence is expected under two different hypotheses in the rival space, but is rare overall. Having a visual construction of the rival space populated by the competing explanations and the evidence expected under each can help reveal where these problems may arise. The appendix also suggests a modified test to overcome this issue.

researchers identify the constituent pieces of a theory for which they must find evidence.[20] This particular exercise forces researchers to think about (and then test) the underlying mechanisms and assumptions that motivate a theory's predictions.

Second, researchers should devise tests of the posited relationships among explanations. What sort of evidence would suggest that H_M and H_R are working in conjunction? In the context of Figure 3, we might ask how we knew we were operating in the darker overlap areas. Alternately, we might ask what sort of evidence would suggest mutual exclusivity. Thinking a priori about how to test relationships among competing explanations can both guide evidence collection and promote more nuanced interpretation.

5.3. *Step 3: Collect & Evaluate Evidence*

In this template, the process of evidence collection and evaluation is explicitly catholic in nature: researchers may opt to construct analytic narratives based on the evidence, they may draw on the crisp-set theoretic approach, or they may use explicit Bayesian inference. Each tool represents a modeling choice, the appropriateness of which is left to the researcher to justify. Furthermore—and in the spirit of avoiding assumptions of mutual exclusivity on all fronts—researchers may use different approaches for different pieces of evidence within the same project.[21] For example, if for one part of a project, a researcher is adjudicating between two exclusive hypotheses for which the overall balance of evidence is unclear, Fairfield and Charman (2015) argue that Bayesian computation is especially useful. Yet for another hypothesis, the balance of evidence may be so clearly tipped in one direction that the benefits of computation become negligible, in which case a simple narrative presentation may be preferable.

[20] Thus, one should not just write $A \to Z$, but instead, she should specify $A \to B, B \to C$, etc.

[21] I am grateful to one of my reviewers for pointing out that every piece of evidence need not be evaluated using the same approach.

6. THE FASHODA REDUX: APPLYING RAR TO SCHULTZ'S (2001) ANALYSIS

To demonstrate the value of the RAR framework, this section (re)analyzes a substantive example that is widely considered exemplary of process tracing: Schultz's (2001) analysis of the Fashoda Crisis.[22] Schultz argues that the structure of democratic institutions (specifically, the visibility of public and opposition party opinions) curtails democratic leaders' abilities to bluff, thus enhancing the credibility of threats. Although his exploration of the Fashoda incident illustrates convincingly the role of unified British public opinion in prompting France to surrender, the study's overall strength would be greater with a more nuanced template for evaluating evidence conditional on the relationship between rival explanations. Specifically, I show that his treatment of alternative explanations as mutually exclusive gives rise to two faulty conclusions.

In this section I evaluate Schultz's treatment of his main hypothesis and two competing theories. For each explanation, I outline its central tenets, summarize Schultz's reasoning for accepting or dismissing it, and then evaluate each hypothesis (and the evidence provided) using the RAR framework. The new framework yields insights that either diverge from or add crucial nuance to the author's original conclusions, thus opening the door to new and interesting insights about crisis behavior.[23]

6.1. *Main Hypothesis: Affirming the "Confirmatory Effect"*

Schultz's main hypothesis is that French decision-makers backed down once they realized that the British government had support from both the majority and opposition parties in parliament. This hypothesis is an application of a more general proposition known

[22] To provide some background, the Fashoda crisis occurred in 1898 when France and Britain entered into a territorial dispute over control of north-eastern Africa. The dispute quickly escalated, yet France ultimately backed down prior to the dispute escalating to all-out war.

[23] See Appendix B in the supplementary material for an additional example in which I apply the RAR framework to Nina Tannenwald's (2007) *The Nuclear Taboo.*

as the confirmatory effect: "the idea that a signal sent by two actors with opposing interests is more informative than a signal sent by one actor with known incentives to misrepresent its preferences" (2001, 162). Schultz validates his theory by first providing evidence that passes a hoop test, then providing strong confirmatory evidence in favor of this explanation, which in effect satisfies a smoking-gun test.

H1: Support from the opposition party in Britain is a requirement for the confirmatory effect.

> To satisfy this hoop test, Schultz provides primary source evidence indicating the Liberal Party's resolve to back the governing party. The most convincing piece comes from a speech made by the Liberal Party leader in which he states, "Behind the policy of the government is the united strength of the nation...The nation will make any sacrifice and go any length to sustain them" (188).

H2: France backed down due to unified British resolve.

> On this point, Schultz finds evidence that nearly constitutes a smoking-gun for his hypothesis: a letter between key French decision-makers noting that "the Liberals [in London] have come out as much if not more intransigent than the partisans in the government" (190). Schultz concludes from this exchange that "the support of British opposition groups was not lost on French diplomats" (190).

Schultz is able to supply the sort of confirmatory evidence of which most researchers only dream. Yet his treatment of competing hypotheses could be strengthened given a more nuanced conceptualization of their relationship to his main. Reevaluating Schultz's evidence in light of the new framework reveals key oversights regarding his alternative explanations.

6.2. *Neorealism: Establishing Coincident Rivals*

Turning to competing explanations, Schultz first explores the neorealist balance-of-power theory. Since Britain had a decisive military advantage, neorealists would argue that

"France backed down because Britain was stronger" (Schultz 2001, 177). Despite the valid prediction on the outcome, Schultz dismisses balance-of-power theory on the grounds that it fails to explain the onset of the crisis in the first place, and why the conflict escalated to near-war (2001, 177).

However, taking Schultz's evidence through the RAR framework reveals that while power asymmetries alone do not account for *all* aspects of the crisis, this explanation should not be entirely ruled out. I first examine the relationship between Schultz's confirmatory effect hypothesis and balance-of-power theory. Drawing on insights from the relationship, I demonstrate that evidence provided to rule out balance-of-power theory was inadequate.

Q1: Do balance-of-power and the confirmatory effect predict divergent outcomes?

> No—neither would have predicted an escalation to all-out war.

Q2: Do the theories require divergent evidence?

> No—the confirmatory effect hypothesis requires evidence about threat perception and the credibility of threats from unified governments. Balance-of-power theory requires evidence of a power asymmetry and actors' knowledge of the power distribution.

Q3: Does evidence in favor of one have any effect on the other?

> No, the two theories operate via distinct, but non-contradictory pathways. Balance-of-power does not acknowledge sub-state behavior, and the confirmatory effect is not affected at all by the actual power distribution.

It follows that confirmatory effect and balance-of-power are coincident hypotheses: they operate through distinct pathways to make similar predictions about the outcome at Fashoda. In the broader context of the conflict, it is likely that a *combination* of British resolve and military advantage worked together to convince France to back down. After all, if France had evidence of Britain's resolve to go to war, yet the power distribution favored the French, retreat would have been less likely. Since the rival explanations are coincident, ruling out balance-of-power theory requires finding explicit evidence that Britain's military advantage did not factor into French decision-making. The following summarizes Schultz's evaluation of the neorealist explanation.

H1: Preexisting knowledge of the power distribution is a requirement for avoiding war.

> To address this test, Schultz first turns to a statement made by the British Admiralty claiming that "in the event of war, France would not have 'a ghost of a chance'" (177). He argues that "the French had to know from the outset" that they could not withstand a confrontation, though he does not provide explicit evidence of their knowledge.

H2: Collusion with third parties to tip the balance is a requirement for the weaker state to act aggressively.

> On this point, Schultz provides evidence that Russia was neither willing nor able to assist France in the crisis, and that French leaders did not harbor optimistic expectations of Russian assistance (2001, 180).

Schultz provides convincing evidence that balance-of-power theory cannot explain French aggression or the duration of the crisis. Insights from RAR, however, reveal a critical gap between the evidence needed and the evidence used to eliminate the neorealist explanation. Merely demonstrating that a theory provides an incomplete explanation is an insufficient basis to eliminate it from consideration. The evaluation of balance-of-power theory could have been stronger by acknowledging its partial role in explaining the outcome at Fashoda.

6.3. *Institutional Constraints: Insights from Congruent Rivals*

Schultz's second alternative explanation derives from the institutional democratic peace proposition (hereafter, IDPP). He draws on the underlying mechanism to derive a set of monadic expectations about the behavior of British leaders during the crisis.[24] The

[24] Although the democratic peace proposition largely focuses on explaining a dyadic outcome, Schultz derives a set of monadic predictions and implications about "the relationship between democracy and

central tenet of this theory is that democratic institutions—specifically, a legitimated opposition party and the public's ability to sanction leaders via elections—exert constraints on democratic leaders' behavior. In short, since democratic leaders usually face multiple and varied sources of opposition to war, it is difficult for them to take unilateral military action, thus impeding hawkish behavior (2001, pp. 13, 18, 182, 183). In light of Britain's steadfast belligerence, Schultz quickly rules out the IDPP as "problematic" (2001, 182).

Since the institutional constraints argument typically explains why conflict is *less* likely, this explanation indeed appears inappropriate. To evaluate the validity of his conclusion, I reanalyze his evidence through the RAR framework: first by considering the relationship between the IDPP and the confirmatory effect, then by retesting the IDPP in light of the insights revealed. I demonstrate that a more systematic evaluation of evidence in line with the prescriptions of RAR would have (1) generated the *opposite* conclusion regarding the role of institutional effects in explaining the Fashoda crisis, and (2) revealed an interesting and counterintuitive prediction regarding conflict behavior.

Q1: Do the theories predict divergent outcomes?

> No—largely because they make predictions about different aspects of conflict. Neither would predict war, and their respective logics stem from different places.

Q2: Do the theories require divergent evidence?

> The evidence required to test the confirmatory effect is (1) a unified opposition, (2) a threat, (3) the target observing that the government is unified behind the threat.
>
> The evidence required to test whether institutional constraints shape leader behavior is (1) a specification of the variety of opinions across the polity (the governing party, the opposition party, and the public), (2) an indication that the leader's behavior was a function of those opinions.

Since the IDPP does not deal with threats, the last two conditions for the confirmatory effect (threat existence and perception of unanimity) can be ruled out as possible triggers

war" from the institutional mechanism underlying the proposition (2001, 12-13).

of mutual exclusivity. We are then left asking how evidence of a unified government would bear on the evidentiary requirements of IDPP. Since one requirement of demonstrating institutional constraints is to specify the range of opinions across the polity, evidence of a unified government represents one possible configuration of opinions out of many. Unification does not controvert IDPP; rather, it describes a rare condition in which the governing party and opposition party agree. As such, we can rule out evidentiary mutual exclusivity as well.

Beyond ruling out mutual exclusivity, the rarity of a unified government raises an important question: what would the institutional constraints argument predict under this condition? This question highlights a crucial and testable assumption on which IDPP predictions are based. IDPP assumes that the more people to whom a leader is accountable, the more likely she is to encounter a wide variety of opinions pulling in opposing directions. Schultz and others predict restraint and pacifism because navigating *countervailing* opinions leaves the leader no choice but to be balanced and circumspect. However, when all of the constraining opinions are exerting force in the same direction, the theory of institutional constraints is also *supported* if a leader is consequently forced to move excessively far in an undesirable policy direction (e.g. going to war). Testing that assumption reveals a novel insight.

A1: The presence of multiple factions in government leads to a variety of countervailing opinions about going to war.

> As Schultz previously stated, the nation was "unified," and the Liberals "have come out as much if not more intransigent than the partisans in the government" (2001, 190). Thus, the assumption on which IDPP predictions are based does not hold.

At this point, Schultz overlooks an interesting and unprecedented outcome of the IDPP, which in turn leads to an inconsistency in his argument. The falseness of the assumption does not undermine the theory, rather, it changes the prediction. Since all actors were pushing for war, the IDPP would, ironically, predict that Salisbury was constrained into adopting a hawkish stance, which is precisely what Schultz's evidence shows. Thus, the evidence Schultz provides to *counter* the IDPP actually corroborates it. Consequently, the confirmatory effect and the institutional constraints argument are congruent: they make similar predictions and are corroborated by similar evidence.

In accordance with the RAR framework, an analysis of the relationship between IDPP and the confirmatory effect hypothesis reveals not only their congruence, but also a critical assumption on which the institutional constraints argument relies. Rather than undermining IDPP, Schultz's corroborating evidence for the confirmatory effect—unified government opinion—unearths a rare political situation that fundamentally alters the predictions of the theory. This insight would have allowed Schultz to exploit an interesting opportunity in IR theory: predictions about conflict-proneness change as a function of government unification. Schultz has the evidence to demonstrate that when all constraining actors happen to agree on a policy issue, democratic leaders may be forced *away* from pacifism and towards belligerence. This insight suggests that some of the IR literature broadly conflates institutional *constraint* and institutional *restraint*.

7. DISCUSSION

Although process tracing is widely touted as a powerful tool for adjudicating among competing explanations, the literature lacks comprehensive guidelines in service of this goal. The pitfalls of assuming exclusivity are not entirely lost on process-tracing scholars, and savvy practitioners do acknowledge the partial roles of competing explanations. However, the continued bracketing of non-exclusivity for the sake of advancing other aspects of the method has caused process-tracing scholarship to get ahead of itself on a critical dimension. This article removes those brackets and develops a framework that serves as both a sound methodological foundation for further process-tracing scholarship and a crisp practical guide for substantive research.

The first major contribution of the RAR framework is a typology of the four possible relationships among competing hypotheses. Mutual exclusivity describes the rare situation when two explanations cannot simultaneously be true in a single case. Coincident hypotheses are akin to statistical independence: the truth or existence of one in no way affects that of the other. Congruent hypotheses operate via some similar mechanisms so that evidence in favor of one explanation also corroborates the other. Finally, as a special case of congruence, inclusive hypotheses occur when one explanation represents a theoretical extension of another, and thus can be subsumed under the same heading.

One of the foremost advantages of the RAR framework is that it guides researchers through a systematic consideration of alternative explanations without forcing researchers into one or another process-tracing approach. Crucially, however, RAR does have important implications for how each approach is conducted—especially when researchers are tasked with adjudicating among non-exclusive alternatives. The most extensive consequences of integrating RAR into existing approaches are found in the Bayesian approach. While scholars in this tradition have acknowledged the added complexity of estimation when mutual exclusivity breaks down, this article is the first to derive the full expansion of Bayes' rule in order to guide practitioners using this approach.

With the RAR framework in hand, I demonstrate that even the best examples of process tracing can benefit from this new tool. For both Schultz's work and Tannenwald's work, employing RAR procedures yields novel inferences that represent important avenues for new research in the authors' respective fields. The RAR framework represents a much-needed advancement in process tracing by bringing the procedures of the method in line with its goals. With the advent of this approach, scholars can proceed both more systematically and more self-consciously through process-tracing research.

REFERENCES

Beach, Derek & Rasmus Brun Pedersen. 2012. *Process-Tracing Methods: Foundations and Guidelines*. Ann Arbor: The University of Michigan Press.

Bennett, Andrew. 2009. Process Tracing: A Bayesian Perspective. In *The Oxford Handbook of Political Methodology*, ed. Janet Box-Stefensmeier, Henry Brady & David Collier. Oxford University Press pp. 702–722.

Bennett, Andrew. 2010. Process Tracing and Causal Inference. In *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, ed. Henry E. Brady & David Collier. Second ed. Lanham, MD: Rowman and Littlefield.

Bennett, Andrew. 2014. "Process Tracing with Bayes: Moving Beyond the Criteria of Necessity and Sufficiency." *Qualitative and Multimethod Research* 12(1):46–51.

Bennett, Andrew. 2015. Systematizing Process Tracing with Bayesian Analysis. In *Process Tracing: From Metaphor to Analytic Tool*. Cambridge: Cambridge University Press.

Bennett, Andrew & Jeffrey T. Checkel, eds. 2015. *Process Tracing: From Metaphor to Analytic Tool*. Cambridge: Cambridge University Press.

Blatter, Joachim & Markus Haverland. 2012. *Designing Case Studies: Explanatory Approaches in Small-N Research*. New York: Palgrave Macmillan.

Brady, Henry E. 2006. "Toward a Pluralistic Vision of Methodology." *Political Analysis* 14(3):353–368.

Brady, Henry E. 2010. Data-Set Observations versus Causal-Process Observations: The 2000 U.S. Presidential Election. In *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, ed. Henry E. Brady & David Collier. Second ed. New York, NY: Rowman and Littlefield.

Collier, David. 2011. "Understanding Process Tracing." *PS: Political Science and Politics* 44(4).

Collier, David, Henry E. Brady & Jason Seawright. 2010. *Introduction to the Second Edition*. Second ed. New York: Rowman and Littlefield chapter 1.

Collier, Paul & Anke Hoeffler. 2004. "Greed and Grievance in Civil War." *Oxford Economic Papers* 56(4):563–595.

Fairfield, Tasha & Andrew Charman. 2015. "Bayesian Probability: The Logic of (Political) Science." University of California, Berkeley. Prepared for the annual meeting of the American Political Science Association: San Francisco, CA.

Freedman, David A. 2010. On Types of Scientific Inquiry: The Role of Qualitative Reasoning. In *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, ed. Henry E. Brady & David Collier. Second ed. Rowman and Littlefield pp. 1–21.

George, Alexander L. 1979. Case Studies and Theory Development: The Method of Structured, Focused Comparison. In *Diplomacy: New Approaches in History, Theory and Policy*, ed. Paul Gordon Lauren. The Free Press pp. 43–68.

George, Alexander L. & Andrew Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. BCSIA Studies in International Security Cambridge, MA: MIT Press.

George, Alexander L. & Timothy J. McKeown. 1985. *Case Studies and Theory Development in the Social Sciences*. Cambridge, MA: MIT Press.

Goertz, Gary & James Mahoney. 2012. *A Tale of Two Cultures: Qualitative and Quantitative Research in the Social Sciences*. Princeton, NJ: Princeton University Press.

Hall, Peter A. 2006. "Systematic Process Analysis: When and How to Use It." *European Management Review* 3(1):24–31.

Humphreys, Macartan & Alan Jacobs. 2015. "Mixing Methods: A Bayesian Approach." *American Political Science Review* 109(4):653–673.

Kay, Adrian & Phillip Baker. 2015. "What Can Causal Process Tracing Offer to Policy Studies? A Review of the Literature." *Policy Studies Journal* 43(1):1–21.

Mahoney, James. 2008. "The Logic of Historical Explanation in the Social Sciences." *Comparative Political Studies* 42(1):114–146.

Mahoney, James. 2012. "The Logic of Process Tracing Tests in the Social Sciences." *Sociological Methods & Research* 41(4):570–597.

Mahoney, James & Gary Goertz. 2006. "A Tale of Two Cultures: Contrasting Quantitative and Qualitative Research." *Political Analysis* 14(3):227–249.

Maoz, Zeev & Bruce Russett. 1993. "Normative and Structural Causes of Democratic Peace." *American Political Science Review* 87(3):624–638.

Platt, John R. 1964. "Strong Inference." *Science* 146(3642):347–353.

Rohlfing, Ingo. 2012. *Case Studies and Causal Inference: An Integrative Framework*. New York, NY: Palgrave Macmillan.

Rohlfing, Ingo. 2014. "Comparative Hypothesis Testing Via Process Tracing." *Sociological Methods & Research* 43(4):602–642.

Schultz, Kenneth. 2001. *Democracy and Coercive Diplomacy*. Cambridge: Cambridge University Press.

Van Evera, Stephen. 1997. *Guide to Methods for Students of Political Science*. Ithica, NY: Cornell University Press.

Waldner, David. 2014. What Makes Process Tracing Good? Causal Mechanisms, Causal Inference, and the Completeness Standard in Comparative Politics. In *Process Tracing: From Metaphor to Analytic Tool*, ed. Andrew Bennett & Jeffrey T. Checkel. Cambridge University Press.

Zaks, Sherry. 2012. Relationships Among Rivals: Contending Hypotheses and the Logic of Process Tracing. Paper Prepared for the Annual Conference of the American Political Science Association New Orleans, LA: .

Notes

[1]The scope of this article is limited to process tracing as it is used for theory-testing purposes. For discussions of other uses of process tracing see Kay & Baker (2015).

[2]See also (George & Bennett 2005, Brady 2006, Bennett 2010, Collier 2011).

[3][Link to online appendix.]

[4]I use "analytic narratives" to refer to the classical approach to process tracing in the vein of George (1979), Van Evera (1997), and George & Bennett (2005), among others. This term should not be confused with the usage of "analytic narrative" in Bates et al. 1998.

[5]Bennett's (2010) table draws on the criteria of Van Evera's (1997) four ideal-type tests based on the same dimensions. Uniqueness refers to evidence that is so specific to one theory that finding it is sufficient to confirm the hypothesis. Certainty refers to evidence that must be found for the theory to be true.

[6]The sample space, Ω, describes the set of all possible outcomes, and the corresponding probability space is the assignment of probabilities to each event ω_i. For example, the roll of a single fair die can be represented by a sample space with six disjoint outcomes $(\omega_1, \omega_2, \dots, \omega_6)$ of equal probability, $p(\omega_i) = \frac{1}{6}$.

[7]Mutual exclusivity is defined as $P(A \cap B) = 0$ (i.e. the two events can never co-occur), independence describes is defined as $P(A \cap B) = P(A) * P(B))$ (i.e. the two events may co-occur, but neither affects the other), and dependence is defined as $P(A \cap B) = P(A) * P(B|A))$ (i.e. the likelihood of one event depends upon the occurrence of the other).

[8]The respective predictions of Newtonian and quantum mechanics illustrate this case well. The kinematic and dynamic predictions of Newtonian mechanics are accurate for a wide range of phenomena, but the laws break down at the atomic level. For objects smaller than 10^{-9}m, quantum mechanics makes different, and more accurate predictions about how particles behave. Thus, at the atomic level, Newtonian and quantum mechanics exhibit mutual exclusivity; yet, the predictions of both theories converge for larger phenomena, which suggests that their mutual exclusivity is conditional on size.

[9]In the Bayesian approach, for example, this problem would manifest as an artificially low prior on a given hypothesis.

[10]The distinction between expecting *different* and *divergent* evidence is potentially the source of many errors in process tracing. Two congruent theories may nonetheless need different types of evidence from one another to be verified. Too often, however, are researchers inclined to rule out one hypothesis because they found the unique kind of corroborating evidence in favor of the other.

[11]For an example of congruence in practice, see the Tannenwald analysis in the online appendix (Appendix B).

[12]It is worth noting that Bayesian scholars argue that the construction of analytic narratives is merely an implicit use of Bayesian inference. The purpose here is not to create sharp (let alone false) divides among approaches, but rather to examine the implications for practitioners.

[13]Mahoney (2012), for example, provides a detailed description of identifying joint sufficiency, and then separately poses the question, *"How are rival hypothesis eliminated?"*. This separate treatment suggests that scholars lack a conceptual framework for integrating alternative explanations that may work in conjunction with the main hypothesis.

[14]This notation draws on Humphrey and Jacobs (2015), which represents the successful search of a piece of evidence as $K = 1$.

[15]Methodologists are engaged in a debate about the benefits and drawbacks of adopting this explicit mathematical approach. The debate itself is beyond the scope of this article. For a comprehensive introduction to the Bayesian approach and the debates surrounding it see, Bennett (2009, 2014, 2015), Beach & Pedersen (2012), Humphreys & Jacobs (2015), Fairfield & Charman (2015).

[16]When estimating the coincidence of two statistically independent events, $P(A \cap B) = P(A) \times P(B)$.

[17]And this definition does not include the additional complication of overlap among *other* alternatives.

[18]Though Kay & Baker (2015) encourage researchers to begin research in a similar way, their recommendation omits a consideration of relationships among competing hypotheses.

[19]In Appendix A (online), I demonstrate that process tracing conclusions (especially—though not limited to—conclusions in Bayesian process tracing) are particularly susceptible to bias when a piece of evidence is expected under two different hypotheses in the rival space, but is rare overall. Having a visual construction of the rival space populated by the competing explanations and the evidence expected under each can help reveal where these problems may arise. The appendix also suggests a modified test to overcome this issue.

[20]Thus, one should not just write $A \rightarrow Z$, but instead, she should specify $A \rightarrow B, B \rightarrow C$, etc.

[21]I am grateful to one of my reviewers for pointing out that every piece of evidence need not be evaluated using the same approach.

[22]To provide some background, the Fashoda crisis occurred in 1898 when France and Britain entered into a territorial dispute over control of north-eastern Africa. The dispute quickly escalated, yet France ultimately backed down prior to the dispute escalating to all-out war.

[23]See Appendix B in the supplementary material for an additional example in which I apply the RAR framework to Nina Tannenwald's (2007) *The Nuclear Taboo*.

[24]Although the democratic peace proposition largely focuses on explaining a dyadic outcome, Schultz derives a set of monadic predictions and implications about "the relationship between democracy and war" from the institutional mechanism underlying the proposition (2001, 12-13).