



Causal Thinking in Implicit Personality Theories and Person Representations: An Explanation of Egocentric Pattern Projection

by Clayton Roger Critcher

This thesis/dissertation document has been electronically approved by the following individuals:

Dunning, David Alan (Chairperson)

Gilovich, Thomas Dashiff (Minor Member)

Ferguson, Melissa J. (Minor Member)

Cutting, James Eric (Additional Member)

CAUSAL THINKING IN IMPLICIT PERSONALITY THEORIES AND PERSON
REPRESENTATIONS:
AN EXPLANATION OF EGOCENTRIC PATTERN PROJECTION

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
In Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Clayton Roger Critcher
August 2010

© 2010 Clayton Roger Critcher

CAUSAL THINKING IN IMPLICIT PERSONALITY THEORIES AND PERSON
REPRESENTATIONS:
AN EXPLANATION OF EGOCENTRIC PATTERN PROJECTION

Clayton Roger Critcher, Ph. D.

Cornell University 2010

The self is relied on disproportionately as a source of information about the social world. Recent research has found that people's implicit personality theories—beliefs about how personality traits tend to relate in other people—are rooted in perceptions of the self. That is, people engage in *egocentric pattern projection*: they expect traits to be patterned in other people in the same way that they are patterned in the self. Five studies test why pattern projection emerges and why it is egocentric. It is proposed that individuals look to their own personalities and develop causal theories to explain why personality traits relate as they do in the self. These *causal trait narratives* then guide the causal theories that people use to explain why personality traits correlate more generally, their implicit personality theories. This account makes three predictions that the present studies test and support. First, this suggests that implicit personality theories are not represented merely as assumed correlations, but that causal theories (i.e., Trait A influences Trait B) underlie these assumed relationships. Consistent with this possibility, Study 1 found that implicit personality theories are directional—Trait A implies standing on Trait B differently than Trait B implies standing on Trait A. Second, this account predicts that people will be most likely to pattern project a given trait relationship once they have developed a causal theory for the patterning in the self (Study 2). If causal trait narratives do indeed underlie pattern projection, then pattern projection may be egocentric

because people have more comprehensive causal trait narratives to explain the self than to explain other people. People indeed reported having more causal theories in their causal trait narratives of themselves than of someone else (Studies 3 and 4), and such causal trait self-narratives were more accessible (Study 4). But when participants were asked to generate a causal trait narrative to explain a novel social target, their implicit personality theories assimilated toward the patternings observed in the target. (Study 5) This confirms a theoretically-derived limit on pattern projection's egocentrism and supports a key causal claim: causal trait narratives influence implicit personality theories.

BIOGRAPHICAL SKETCH

Clayton Roger Critcher was born in Shreveport, Louisiana, the son of two professors who told him he could be anything in life except for a preacher or a professor. They are more than pleased he chose the latter. Clayton graduated *summa cum laude* from Yale University in 2005, with an A.B. in Psychology. After five years completing his Ph. D. in social and personality psychology at Cornell University, Clayton will join the faculty of the University of California, Berkeley, as an Assistant Professor of Marketing in the Haas School of Business.

ACKNOWLEDGMENTS

I thank my parents for encouraging me to question intelligently. I thank fate and Elliot Aronson for my accidentally happening upon *The Social Animal* at the Barnes & Noble in Shreveport, LA, in December 1999. I read the first chapter of Aronson's book in the café, was mesmerized, and my life has been more interesting ever since. I thank David Armor for his entirely selfless efforts as my advisor at Yale. It was in my second meeting with David, in the summer of 2002, that it was first suggested to me that David Dunning might be a perfect advisor for me. I thank David Dunning for accepting me as his graduate student, for encouraging me when I felt discouraged, and for offering me some of my most intellectually exciting interactions over the last five years. I always leave his office more excited about this whole enterprise. I thank Tom Gilovich for teaching me that some of the most interesting theoretical questions can be studied by examining phenomena that capture anyone's interest. I thank Melissa Ferguson for her patience, intellect, sense of humor, and talent at humbling me. I thank David Pizarro for always being willing to talk at the highest level about the commonest things. I thank Vivian Zayas for our fun meetings at Gimme Coffee that would run past closing time.

I thank Jane Risen and Karlene Hanko for the genuineness of their friendships; Yoel Inbar for his good nature; Nora Williams for her companionship; Shanette Porter for our seemingly endless dissections of the most contentious issues; and Erik Helzer, Jun Fukukura, Emily Rosenzweig, and Chelsea Helion for helping to keep this place alive even after my first family had left. Thanks to all of the social psychology graduate students for making lab meetings and proseminar the high points of my work week.

I thank the National Science Foundation for the support provided by a Graduate Research Fellowship.

I thank my mother for always being willing to talk me through anything.

TABLE OF CONTENTS

Biographical Sketch	iii
Acknowledgments	iv
Table of Contents	v
List of Figures	vi
List of Tables	vii
Chapter 1: Introduction	1
Chapter 2: The Studies	14
Study 1	15
Study 2	24
Study 3	33
Study 4	39
Study 5	42
Chapter 3: General Discussion	53
References	65

LIST OF FIGURES

Figure 1: The impact of implicit personality theories on target judgment as a function of target type and congruence in the directionality of the target judgment and corresponding implicit personality theory (Chapter 2, Study 1)	22
Figure 2: The degree of pattern projection for trait pairs that are or are not bound by causal theories or third-variable theories in causal trait narratives of the self. (Chapter 2, Study 2)	31
Figure 3: Example causal trait narrative (Chapter 2, Study 3)	37
Figure 4: Pattern projection from the novel social targets as a function of processing task condition and the difficulty participants had with the processing task. (Chapter 2, Study 5)	51

LIST OF TABLES

Table 1: The 13 implicit personality theories that differed by directionality (Chapter 2, Study 1)	20
Table 2: Attributes of causal trait narratives describing the self or an other (Chapter 2, Study 3)	37
Table 3: Both versions of the three social targets presented in Study 5 (Chapter 2, Study 5)	45
Table 4: Directional trait theories most overincluded and underincluded in causal trait narratives (Chapter 3)	59
Table 5: Directional trait theories most included in self (versus other) and other (versus self) causal trait narratives (Chapter 3)	60

Chapter 1: Introduction

A person's perspective on his or her social world is fundamentally self-centered. Whether in trying to take another's perspective, trying to predict the opinions of others, or trying to evaluate another person, the self's own perspective, characteristics, and standing influence such judgments. Such self-centeredness, or egocentrism, leads people to inadvertently activate the self when making judgments about others (Dunning & Hayes, 1996). People define traits in self-serving ways when making judgments of others (Dunning, Meyerowitz, & Holzberg, 1989). They use egocentric standards in evaluating the quality of others' performance (Dunning & Cohen, 1992; Dunning & Hayes, 1996) and take an egocentric perspective when considering the viewpoint of others (Epley, Keysar, Van Boven, & Gilovich, 2004; Epley, Morewedge, & Keysar, 2004).

Most demonstrations of the egocentric nature of social judgment can be accounted for by one of two explanations. First, the self is a readily accessible exemplar, and thus an egocentric viewpoint serves as a natural anchor when thinking about the social world. It is effortful to adjust from this egocentric perspective, leaving many social perspectives too egocentrically bound. Second, the self may take an egocentric perspective in judging others in order to manage its own motivational needs. For example, college professors who achieve tenure then adopt more stringent criteria for what constitutes *high productivity* (Eidelman & Biernat, 2007). By egocentrically using one's own levels of achievement to define success, one can make sure that the self is always viewed in the most flattering light as possible.

One of the most well-documented influences of the self on social judgment, *attributional projection*, can be accounted for by both of these mechanisms. People assume that others hold the same traits, attitudes, and behavioral tendencies as does the self—an assumption that has been shown to predict whether or not people see others as honest (Katz & Allport, 1931), happy (Goldings, 1954), sharing their political views (Judd,

Kenny, & Krosnick, 1983), willing to wear an “Eat at Joe’s” sandwich board around campus (Ross, Greene, & House, 1977), and characterized by many different personality traits (Holmes, 1981; Krueger & Stanke, 2001). People will project because the self is a readily accessible guide that may offer reasonable insight as to what others think (Dawes, 1990; Dawes & Mulford, 1996). Also, projection provides a motivationally satisfying sense of social consensus and validation for one’s own beliefs, leading people to believe that even God—perhaps the ultimate arbiter of truth—endorses one’s own belief system (Epley, Converse, Delbosc, Monteleone, & Cacioppo 2009).

But these accounts do not provide a good explanation for a recently identified phenomenon by which self-views impact social judgment: *egocentric pattern projection* (Critcher & Dunning, 2009a). Although egocentric pattern projection is a projection phenomenon—people expect to see features of the self in other people—it is qualitatively distinct from all previously documented projection phenomena. Prior demonstrations of projection reflected variations on a single theme: An attribute is observed in the self, and people expect to or do see that same attribute in other people. In contrast, pattern projection emerges when people observe two attributes in the self, infer a correlation between the attributes based on the way they are patterned in the self, and then expect to see the same correlation in other people. For example, if Jens sees himself as egalitarian and emotional, attributive projection predicts that Jens expects others to be more egalitarian and more emotional than do cold, inequalitarian types. Pattern projection predicts Jens expects egalitarian people to be emotional and prejudicial people to be less so. That is, Jens expects that the patterning of traits in others will recapitulate the patterning of traits in the self. In analysis of variance terms, attributive projection predicts two main effects in social judgment, whereas pattern projection predicts an interaction.

Critcher and Dunning’s (2009a) demonstration of pattern projection challenged a conclusion that had lain dormant in social psychology for more than three decades—that

there were no systematically detectable individual differences in people's implicit personality theories (Schneider, Hastorf, & Ellsworth, 1979). In making inferences about others' personalities, people lean on implicit personality theories—beliefs about how personality traits tend to be configured in people in general (Bruner & Taiguri, 1954). Past research had found that such theories reflected culturally shared beliefs (Wishner, 1960) that were guided in large part by judgments of the degree of semantic overlap between traits (D'Anrade, 1974). These influences would be expected to produce a fair degree of consensus in people's theories. Egocentric pattern projection identifies one reason why there exists variation in people's theories. Individual differences in how traits relate in the self predict individual differences in people's implicit personality theories for such traits (Critcher & Dunning, 2009a).

Critcher and Dunning (2009a) demonstrated that people's implicit personality theories show evidence of pattern projection, that such pattern projection is egocentric (influenced by the way traits are patterned in the self, but not the way traits are patterned in a well-known other), and that the self plays a causal role in bringing about implicit personality theories. In Study 1, participants rated themselves and their freshman year roommates on a number of personality traits, and provided estimates of certain population parameters from which implicit personality theories could be inferred. People pattern projected from themselves, but not from their roommates. That is, beliefs about how personality traits were correlated recapitulated patternings in the self, but not those observed in a roommate. Studies 2 and 3 showed that these egocentric implicit personality theories actually guided people's initial impressions of novel targets and helped organize people's impressions of well-known others. In Studies 4 and 5, Critcher and Dunning (2009a) provided fictitious feedback to participants about their standing on two fictitious personality dimensions: front/back-brainedness and V/Z-dominance. Participants shifted their implicit personality theories to reflect that they assumed the two novel dimensions

were correlated in a way that would produce the patterning observed in the self. For example, individuals who were both V-dominant and front-brained believed that front-brained individuals were more likely to be V-dominant, and that back-brained individuals were more likely to be Z-dominant. No similar shift occurred when participants learned another participant's test feedback. Although Critcher and Dunning (2009a) focused on the *what* of the pattern projection phenomenon, they did not focus on the *why* of it. The present research sought to answer both why pattern projection emerges and why it is egocentric.

An Explanation of Egocentric Pattern Projection

Why would pattern projection emerge, and why might it be egocentric in origin? I hypothesize that pattern projection emerges as a byproduct of trying to understand the self. I suggest that over time people engage in self-reflection that attempts to create a causally coherent narrative explaining why traits relate as they do in the self. For example, creative extroverts may come to decide that exposure to others gives them new ideas and experiences that fuel their creativity, whereas creative introverts may come to believe that time away from others allows them time for solitary reflection that aids in their creative pursuits. Note that both of these causal relationships may be true, but pattern projection emerges when people start to apply these person-specific theories to understand people in general. To the extent that people tend to generate such causal narratives to explain the self more than they do to explain other people, egocentric pattern projection may emerge.

I propose that egocentric pattern projection can be understood to emerge from the convergence of three principles or phenomena. First, and consistent with the account sketched above, I propose that implicit personality theories are not represented merely as correlations, but actually reflect causal theories. That is, people do not merely assume that creativity and extroversion are correlated, but people believe traits are associated because people have a theory for why one trait causes the other. Second, I propose that such

theories are not first developed to explain patternings observed in the world, but are first developed to explain patternings in a sample of one (e.g., the self). This reflects a heretofore unidentified type of person representation, what I call a *causal trait narrative*. Causal trait narratives go beyond piecemeal trait judgments and reflect causal theories that explain why traits are patterned in a person. Such narratives, first used to explain a single person, provide the foundation for people's more general implicit personality theories. The exportation of these theories *is* pattern projection. Third, if this account of pattern projection is true, then the phenomenon would be egocentric because people have more comprehensive (and accessible) narratives to explain the self than someone else. This of course suggests a possible limit to pattern projection's egocentrism. When people develop a comprehensive causal trait narrative about a peer, then people should begin to non-egocentrically pattern projection from that peer. This account is consistent with, makes novel connections between, and extends on theorizing about perceived correlations, inductive reasoning, implicit personality theories, narrative conceptions of the self, and asymmetries in representations of the self and others. I will proceed to outline previous research that supports the tenability of my explanation of egocentric pattern projection.

Premise 1: Causal Theories Underlie Assumed Correlations (IPTs)

First, our explanation relies on cognitive psychology findings that perceived correlations are not represented as matrices of assumed correlation coefficients, but as explanatory theories (Chapman & Chapman, 1967; Kunda, Miller, & Claire, 1990; McNorgan, Kotack, Meehan, & McRae, 2007; Murphy & Wisniewsky, 1989). Illustrating the crucial role of causal theories in correlation detection, Ahn, Marsh, Luhman, and Lee (2002) tried to account for past findings showing that people are much better at detecting certain real-world correlations than they are others (Malt & Smith, 1984). For example, most people recognize that among birds, there is a relationship between a tendency to live close to the ocean and to have fish be part of the bird's diet. In contrast, people are less

likely to notice that among shirts, there is a correlation between a shirt having buttons on it and its being long-sleeved. The former correlation is more intuitive because people can generate a causal narrative to explain it (e.g., “If a bird wants to eat fish, it behooves it to live near the ocean.”) In contrast, people are less likely to have an explanatory narrative that can account for the correlation between buttons and sleeves among shirts.

As Ahn et al. (2002) explicitly noted, their correlational study left open the question of whether “people explicitly notice correlations because they can explain them, or [whether] people impose explanations after they explicitly notice correlations” (p. 115). Note that my explanation of egocentric pattern projection is essentially a mix of these two ideas. I suggest that people first notice a “correlation” in a sample of one (i.e., the self). They then generate a theory to explain that patterning, and this theory serves as the foundation for the expectation of this patterning in others.

Several, independent programs of research suggest that people may readily generalize the theories they have developed about a single exemplar (e.g., the self) to understand all members of a category (e.g., people in general). Examining instances of inductive reasoning, Lassaline (1996) found that people induced general relations from observed feature co-occurrences to the extent that the features were bound by perceived causal relations as opposed to when they were merely temporally contiguous. Showing a similar effect in a more social domain, Risen, Gilovich, and Dunning (2007) found that once participants generated explanations for why a specific minority group member exhibited certain behaviors, people readily generalized this new social knowledge to new minority group exemplars. Both examples show that theories developed to explain single examples can readily be extended to describe members of a more general category. Of course, when Risen et al. (2007) presented participants with odd behaviors (e.g., sleeping in a tent in one’s bedroom) by members of unusual groups (e.g., Jehovah’s Witnesses), then the theories participants would create (e.g., I guess one of the tenants of their religion

is this unusual sleeping practice) should quite reasonably be extended to new group exemplars. Thus, although there is certainly precedent for the generalization of causal theories, it remained less certain whether I would observe that a theory connecting two non-central personality traits would prompt similar generalization.

Although there is clearly precedent for causal theories to underlie perceived correlations, a key question in the present research is whether such causal theories underlie the specific class of correlations that make up implicit personality theories. Historically, there have been 3 conceptions of the structure of implicit personality theories. First, according to an associationistic view, IPTs reflect assumed correlations. As Anderson and Sedikides (1991) explained, an associationistic IPT would assume that “Trait A [would] describe a person who has Traits B, C, and D to the extent that A is assumed to correlate positively with B, C, and D” (p. 203). From this perspective, IPTs are represented as (some version of) a Pearson product-moment correlation. Second, the dimensional view is analogous, except it assumes that people are located on different factors or dimensions in a multi-dimensional trait space (Kim & Rosenberg, 1980; Rosenberg, 1976). Because traits that load on the same factor can be assumed to be correlated, the dimensional view differs from the associationistic view in more neatly delineating traits that are intercorrelated (i.e., those that load on the same factor) from those that are uncorrelated (i.e., those on different factors). Third, Anderson and Sedikides (1991) provide support that implicit personality theories may also be typological. From this perspective, different traits cluster into “person types.” These trait clusters guide social perception above and beyond the assumed correlations between or dimensional structure of traits.

I am proposing a fourth possibility, that implicit personality theories are just that—*theories*. They are not merely Pearson correlation coefficients, factor loadings, or traits clusters. And in fact, Sedikides and Anderson (1994) presaged this perspective by arguing

that causal theories are what underlie the person types they had earlier identified (Anderson & Sedikides, 1991). Sedikides and Anderson (1994) found that traits within a cluster or a person type were seen to be more causally related than traits not in the same cluster. As the title of their paper clearly describes, causal theories may be “The Glue That Holds Person Types Together.” I propose that causal theories do not merely underlie people’s understanding of the trait clusters that underlie person types, but may underlie people’s thinking about trait relationships more generally. Thus, I do not doubt that people may hold particularly elaborated narratives or causal theories to account for certain person types, but I propose that such causal theories may account for people’s implicit personality theories more generally.

Premise 2: Person Representations (Sometimes) Include Causal Trait Narratives

I have argued that beliefs about trait correlations (IPTs) may be rooted in causal theories, and that if people have developed such theories to explain the self, these may color one’s IPTs. But this is premised on the possibility that person representations may include such explanatory causal theories—what I call *causal trait narratives*. Causal thinking is known to be pervasive in many domains of thought. Prior research has noted the central or organizing role of causal thought in people’s understanding of concepts and categories (Murphy & Medin, 1985), in people’s ability to comprehend stories (Black, Galambos, & Read, 1984; Hastie & Pennington, 1991), in attempts to understand social relationships and events (Miller & Read, 1991; Miller, Turnbull, & McFarland, 1989), and in constructing explanations to make sense of life experiences (Baumeister & Newman, 1994). I am arguing that causal thinking may also be important to the way people think about how traits are configured in a person.

People are especially curious about questions of causality in social domains. As Kunda, Miller, and Claire (1990) suggested, most people find it of little interest why inanimate objects have the features they do (e.g., color), but are fascinated with why

people develop into the people they are. And at first glance, the person perception literature, with its historical emphasis on trait ascription (Anderson, 1981; Asch, 1946), might seem to have overlooked the type of causal thinking that may underlie person perception. Given that traits are the most commonly used descriptors used when describing others (Park, 1986), this apparent neglect may seem perfectly reasonable. But if one looks carefully at what social descriptions have been credited as trait-based, one sees that even trait-based descriptions are filled with causal logic. For example, one of Park's (1986) participants provided the following description of another person: "She is wealthy and egotistical, which makes for great fashion sense and good looks." If one merely coded such a sentence for the presentation of piecemeal facts, one would have overlooked the important causal structure that underlies this narrative [(Wealthy + Egotistical) → (Fashionable + Attractive)]. Park (1986) was referring to traits when she said that, "The dominant category of social information appears to be that which is based on inferences as opposed to observed facts" (p. 911). I entirely agree, but suggest that such inferred conclusions do not stop with traits, but also include the causal structure that underlies our conception of traits and the relation between them. That is, people develop and possess causal trait narratives.

Beyond this limited evidence that person conceptions include causal trait narratives, there is reason to believe that people should find causal trait narratives appealing, and perhaps even more appealing than a list of traits. In general, beliefs that can explain other notions are seen to be more valuable than beliefs that can be explained (Preston & Epley, 2005). This suggests that people find explanations that go to deeper levels of analysis (e.g., a causal explanatory trait narrative) more appealing than the features that such beliefs can explain (e.g., traits). And given that people have little difficulty generating ad hoc theories on demand (McInorgan, Kotack, Meegan, & McRae, 2007), even about seemingly contradictory evidence (Asch & Zukier, 1984), people will

likely be successful in generating comprehensive narratives that explain why a person (e.g., the self) is who he is. Given that causal trait narratives would seem to be epistemically satisfying, it is plausible to expect that people develop them.

The idea that people do not think about people merely as lists of traits has been advanced most strongly by Daniel McAdams and his colleagues in their work on life narratives (McAdams, 1985, 2001). McAdams (1995) says trait descriptions offer “the psychology of the stranger” (p. 365). When we know someone better, we think of them in terms of their personal strivings, motivations, hopes, and values. At the most sophisticated level, people strive for a coherent narrative that expresses unity and purpose by stringing together the events of one’s past, present, and anticipated future. Such narratives move beyond mere lists of behavioral episodes (Adler & McAdams, 2007; Pals, 2006), in part by offering coherent, causal explanations that provide structure and order to a life story (Habermans & Bluck, 2000).

I agree with McAdams’s point that our impressions of others are not mere lists of features about them, whether those are traits, strivings, or other intimate details that McAdams (1995) has called “personal concerns.” I believe that McAdams’ three levels of personality have skillfully laid out the raw materials that may go into conceptions of people. Also, I believe that what defines more sophisticated person representations is not merely the content of those representations (e.g., traits vs. behavioral episodes), but their structure as fractured or piecemeal as opposed to (causally) coherent and unified. In a sense, I propose that people have organizing narratives not only to make sense of their life episodes, but to make sense of the structure of their personality (i.e., their traits). These narratives offer a causally coherent structure that explain how different traits influence each other, providing a coherent explanation of how so many individual features can co-occur and give rise to a person.

Bruner's (1990) distinction between paradigmatic and narrative thought is helpful in understanding the basic structure of a causal trait narrative. Paradigmatic thought is characterized by logic and rational analysis, whereas narrative thought is characterized by storytelling. Causal trait narratives reflect a mix of these two types of thought. People have narrative explanations that have their foundation in logic. Thus, when Joe says that his competitive nature leads him to conscientiously pour over the rules and regulations of every contest he enters, he expresses the causal or paradigmatic thought "Competitive → Conscientious" in a more content-rich, narrative form. I suggest that these causal narratives that are meant to characterize a single person (as argued below, typically, the self) become readily generalized to characterize people in general (e.g., Read & Cesa, 1991). In short, the story we tell to understand one person becomes the story we use to understand people in general.

Premise 3: Causal Trait Narratives Are Developed For The Self Especially

I have put forth two main ideas that, in combination, would produce pattern projection. If 1) implicit personality theories are causal theories and 2) people construct causal trait narratives to explain why traits relate in specific people, then it is quite possible that the theories people have to explain specific others may serve as a foundation for their understanding of more general trait correlations. This result is pattern projection. But one mystery would remain. Critcher and Dunning (2009a) found that pattern projection is egocentric. That is, the trait patternings perceived in the self better predicted participants' implicit personality theories than did the trait patternings perceived in one's roommate or in a fellow participant. According to the present account, this result would have emerged only if people are more likely to generate comprehensive causal trait narratives to understand the self as opposed to someone else.

Although much of the original research on self-other differences accounted for ways in which people explained the behavior of the self and others in different ways

(Watson, 1982), more recent research has suggested there are basic representational differences in people's conceptions of themselves versus other people. These differences lend plausibility to the possibility that people are more likely to generate causal trait narratives of the self than of someone else. First, people are more sensitive to the internal dynamics of the self compared to those of other people. The self is represented more in terms of internal, covert reactions, whereas others are understood in terms of their overt social behavior (McGuire & McGuire, 1986). We take a shallower "behavioral" perspective on others in part because we see others' external behaviors as particularly informative. People think that by observing the behaviors of others over time, one will observe more diagnostic information about their "true selves" than would observers of one's own behaviors (Anderson & Ross, 1984). People believe that in understanding the self, external behavior does not provide the whole story (Kruger & Gilovich, 2004). If people feel that descriptions based on external behaviors offer satisfying portraits of other people more so than for the self, they might be more inclined to search for deeper, causal theories to understand the self.

Although there is no direct evidence that people are especially likely to generate causal theories to make sense of the self, there is evidence that representations of others involve mere description, whereas representations of the self reflect more explanation and extrapolation. Prentice (1990) had participants describe the self and others in specific contexts. Participants were also instructed to write more general, decontextualized descriptions of the self and of the same other people. Participants' situation-specific and decontextualized descriptions were similar when describing other people, but showed a divergence when describing the self. That is, our understandings of others involve averaging and aggregating of information in specific contexts. When considering the self, the whole is different than the sum of its parts. Given that causal trait narratives help

transform individual pieces of information into more global, coherent impressions, it may be that people have engaged in more causal narration for the self than for other people.

Of course, it would be naïve to predict that people never engage in narrative construction about others, but I contend that such narratives may be narrower in scope and more simplistic in structure. A number of previous findings support this possibility. First, although people compose “person models” to explain other people, these models tend to be structured around a central trait, with additional information linked to this core concept (Park, DeKray, & Kraus, 1994). Even if causal theories provided the glue for such person models, there would be few traits pairs for such theories to explain. And even when representations of others include many traits, the simplistic structure of those person models suggests there are not elaborate causal trait narratives underlying them. Borkenau and Liebler (1994) found that the factor structure of trait ratings for the self was more complex than the structure of traits ratings for someone else. Beer and Watson (2008) reported that the Big 5 dimensions are less intercorrelated in self ratings than in peer ratings. The more simplistic structure of representations of other people suggests there is little need for casual trait narratives, as much information in such conceptions is redundant. For example, Hampson (1998) found that evaluative inconsistencies are particularly uncommon in representations of others, hinting that people’s understanding of others may reflect a simple evaluative structure (e.g., “She’s a jerk”) that does not necessitate a rich causal narrative structure. In total, self-knowledge may be more nuanced, comprehensive, and complex in structure. Causal trait narratives may provide the glue that unifies this disparate self-knowledge.

Chapter 2: The Studies

I propose that egocentric pattern projection emerges as a byproduct of frequent self-reflection that strives to create a causally coherent narrative about the self. These causal trait narratives created to explain the self (or, at times, someone else) contain causal theories that explain why traits are linked in that person. If implicit personality theories—beliefs about how personality traits tend to be related in the population at large—have causal theories at their core, then theories developed to explain a single person may become the theories used to explain people in general. If people develop richer causal trait narratives for the self than for other people, this would explain why pattern projection is egocentric.

This perspective makes a number of predictions that the five studies presented here were designed to test. First, this suggests that implicit personality theories may not be represented in mental matrices of intertrait correlations, but as causal theories of one trait causing or influencing another (Study 1). Second, people will be more likely to pattern project from a given trait pair in the self when that trait pair is included in people's causal trait narrative of the self (Study 2). That is, if people have developed a causal theory to account for the cooccurrence of two traits in the self, then they should be likely to export that theory such that it is reflected in their implicit personality theories. Third, to explain why pattern projection has been found to be egocentric, people should have more elaborated causal theories to explain the self than to explain someone else (Studies 3 and 4). Fourth, unifying across the second and third predictions above, people should begin to pattern project from someone else once they generate a causal theory that accounts for why traits relate in that other person (Study 5). This would establish that causal trait narratives causally impact implicit personality theories and simultaneously demonstrate that pattern projection is not inherently egocentric.

Study 1

I begin by testing the nature of implicit personality theories. If it is the case that implicit personality theories reflect causal theories (of the form trait A influences trait B), and are not merely represented by correlations, dimensions, or trait clusters, then there should be a certain directionality to people's theories (Ahn et al., 2002). That is, people may have distinct theories for how much trait A predicts trait B, as opposed to trait B predicting trait A. If an implicit personality theory simply reflects the assumed correlation between A and B (or A and B's proximity in multidimensional "trait space"), then A should imply B in the same way that B implies A. For example, it feels more intuitive that a father's height predicts his son's height than the reverse because of the causal theory that underlies this correlation (Tversky & Kahneman, 1980). Of course, logically, these two judgments should be identical (see Kunda & Nisbett, 1986). I examined whether implicit personality theories have directionality in two ways.

First, I tested whether participants' implicit personality theories differed depending on the direction in which they were assessed. Participants expressed their implicit personality theories about fifty-five trait pairs. Between participants, I varied the directionality of these measures. That is, half of participants answered what percentage of people who were above-average on trait A would also be above-average on trait B, whereas the other half made judgments of the opposite form, predicting what percentage of people who were above-average on trait B would also be above-average on trait A. Although the hypothesis that IPTs have directionality does not necessitate that the mean of the two directional theories (a particular directional theory and its converse) differ at the group level, it does follow that if the two judgments

differ there must be directionality to the theories.¹

Second, I went beyond this most direct test of whether directionality matters by looking at the role of directionality in how implicit personality theories are applied. A week before participants expressed their implicit personality, participants completed a supposedly unrelated web-based pretest. They judged 20 targets about whom they had limited information. For each target, participants learned one piece of trait-based information (i.e., that the target was characterized by a trait or that trait's opposite) before estimating what the person was like on a second trait. I varied, for each target, which of 2 traits participants learned the target's standing on. In this way, for each target, there was a match or a mismatch in the directionality of the IPT and the target judgment. That is, some participants judged a target on trait A after learning the target's standing on trait B, whereas others learned the target's standing on trait B before judging them on trait A. To assess whether implicit personality theories are actually directional theories (and do not merely reflect assumed correlations), I tested whether implicit personality theories better predicted trait judgments when they were of matching directionality. If implicit personality theories are merely represented as assumed correlations, then directionality should not matter. For example, imagine rating how happy-go-lucky someone is if all you know about him is that he is very wordy. If there is directionality to IPTs, then this judgment can be predicted better from the IPT $p(\text{happy-go-lucky} \mid \text{wordy})$ than from the IPT $p(\text{wordy} \mid \text{happy-go-lucky})$. If IPTs merely reflection assumed correlations, both IPTs should be equally predictive of the target judgment.

¹ It should be noted that the importance of directionality may disprove that implicit personality theories are merely represented as correlations, but it would not necessarily establish that causal theories underlie IPTs. Instead, such a finding would provide evidence consistent with this possibility. The remaining studies provide more conclusive positive evidence that causal theories bridge person representation and IPTs (to produce pattern projection).

Method

Participants and Design

Two hundred nineteen undergraduates at Cornell University participated in exchange for extra course credit. Participants were instructed to complete a Target Rating Task on a web-based pretest approximately a week before coming to the lab. Two hundred three of the 219 participants (93%) actually did so. All participants completed an Implicit Personality Theory measure in the lab. Participants completed one of two versions of each of these measures.

Procedure

A week before coming to the lab for the main portion of the study, participants completed a web-based Target Rating Task. All participants were informed that they were taking part in the second half of the study. Supposedly, in the first part, all of the students on the floor of a campus residence hall had rated each other on a variety of traits. The second part of the experiment was said to concern whether the participants could form impressions of others based on minimal information.

I created the fictitious targets using eleven personality traits that had been used in previous pattern projection research (Critcher & Dunning, 2009a): *bashful*, *considerate*, *dependent*, *happy-go-lucky*, *idealistic*, *persistent*, *prideful*, *reserved*, *resigned*, *skeptical*, and *wordy*. Of the 55 total trait pairs for which implicit personality theories would be measured in the lab, I randomly selected 20 of these trait pairs in order to form social targets. For ten of the trait pairs, a target was said to very much have one of the traits. For the other ten trait pairs, a target was said to not have one of the traits. I clarified that not having the trait meant that the target was seen to be the opposite of the trait. Thus, if a target was not at all (the opposite of) *bashful*, then this was meant to communicate that the target was particularly outgoing. I created two versions of each target by varying which trait was the one participants received target

information about, and which was the trait participants judged. For example, participants either learned that “Person A was described as not at all (the opposite of) WORDY” and had to estimate how skeptical Person A was, or they learned that the target was not at all skeptical and had to estimate how wordy Person A was.

Once in the lab, participants expressed their implicit personality theories for all 55 possible trait pairs, which included the 20 trait pairs that corresponded to the 20 targets participants had judged a week before. I used a simpler implicit personality theory measure than that used by Critcher and Dunning (2009a). Instead of having participants provide both conditional (e.g., p (skeptical | wordy)) and marginal (e.g., p (skeptical) and p (wordy)) probability judgments, I had participants answer a single question that combined both of these elements. Participants indicated what percentage of people who were more [trait 1] than the average person were also more [trait 2] than the average person. This simplification parallels the measurement strategy from Critcher and Dunning (2009a), but the marginal probabilities are implicit in the question wording. That is, given that 50% of the population is “above-average” (a premise we also made explicit in the instructions to participants), participants made estimates about traits that were assumed to be had by 50% of the population. In this way, higher [lower] responses to the single conditional probability question reflected a higher [lower] perceived correlation between the traits. The exact question participants saw was, “If you know that someone is more [trait X] than the average person, how likely is it that they are also more [trait Y] than the average person?” Half of participants expressed the implicit personality theory in the reverse order, in which trait X and trait Y were switched.

The materials were carefully constructed such that half of the IPTs were assessed directionally to match the corresponding target judgment, whereas half of the IPTs were measured in the reverse direction. In this way, I could see if the

directionality of the implicit personality theory mattered in predicting target judgments, or whether implicit personality theories are actually just represented as correlations (i.e., without direction). Directionality should be observed if causal theories underlie IPTs, but not if IPTs are merely represented as assumed correlations or as trait clusters (both of which are “directionless”).

Results

First, I tested whether the implicit personality theories participants expressed varied depending on the direction of measurement. Participants completed one of two versions of the implicit personality theory measure. The measures differed only in the directionality of the theories assessed. I therefore submitted all 55 implicit personality theories to a one-way MANOVA to test whether the expressed theories differed between the two versions. Indicating that directionality did matter, there was a significant effect of the measure version (i.e., directionality), $F(55, 163) = 3.09, p < .001$. Thirteen of the 55 trait pairs differed at the $p < .05$ level. These theories, and the mean IPTs by directionality, are listed in Table 1. This analysis in itself suggests that implicit personality theories are directional and do not merely reflect assumed correlations, but these results do not confirm the stronger claim that implicit personality theories are actually applied in a directional manner.

Second, I tested whether the social judgments that participants had made a week in advance could be better predicted by implicit personality theories that were matching in directionality than by those whose directionality mismatched. It is first necessary to define several factors that I included in a hierarchical linear model. First, for every trait pair, participants expressed an Implicit Personality Theory, a probability judgment from 0 to 100. Second, each target that participants judged was said to be high or low on (i.e., the opposite of) a trait, its Target Level. Evidence that people were leaning on their implicit personality theories in judging targets would come from

Table 1: The 13 Implicit Personality Theories that Differed by Directionality (Study 1)

Trait A	Trait B	p (Trait A Trait B)	p (Trait B Trait A)
Skeptical	Resigned	64.6	48.1
Skeptical	Bashful	49.5	44.1
Bashful	Considerate	48.9	54.9
Considerate	Dependent	52.1	45.3
Considerate	Wordy	45.6	40.9
Bashful	Prideful	24.6	31.2
Dependent	Prideful	29.7	35.9
Wordy	Prideful	64.2	58.5
Skeptical	Reserved	63.7	50.7
Bashful	Reserved	70.8	77.7
Considerate	Reserved	61.9	56.4
Reserved	Persistent	26.3	34.3
Wordy	Idealistic	58.5	54.8
Reserved	Idealistic	37.1	44.8

Note. All theories are stated as percentages.

an Implicit Personality Theory (IPT) X Target Level interaction. There should be a positive effect of IPTs in judging targets high on a trait, but a negative effect of IPTs in judging targets low on a trait. But the key prediction is that the application of these theories would depend on directionality. Of the 20 trait pairs for which participants

provided both IPTs and Target Judgments, ten matched and ten mismatched. (The precise ten pairs that matched or mismatched varied between participants.) The third factor distinguished between IPTs measured in a direction that matched versus mismatched the direction in which the target judgment was made (Directionality). I expected that the ability of IPTs to account for social judgment (the IPT X Target Level interaction) would be strongest when there was matching directionality. This should produce an IPT X Target Level X Directionality interaction.

I tested a random-slope, random-intercept multi-level model that predicted trait judgments from implicit personality theories, which were nested in trait pairs. I included two additional Level 1 variables: Target Level (-1 = target has the opposite of a trait, +1 = target has a trait) and Directionality (-1 = mismatching IPT/Target judgment directionality, +1 = matching IPT/Target judgment directionality). The three predictor variables were fully crossed to test for interaction effects.

Demonstrating that people did lean on their implicit personality theories in judging social targets, the IPT X Target Level interaction was significant, $t(19.07) = 7.55, p < .001$. But as predicted, this predictive power depended on the matching directionality, $t(4149.23) = 2.39, p = .02$.

To understand the nature of this interaction, I compared the effect of implicit personality theories on social judgments for matching and mismatching directionality for targets that were said to either have or have the opposite of a particular trait. When the implicit personality theory and the target judgment were of matching directionality, there was a positive effect of IPTs on social judgment when the target was said to have a trait, $t = 7.58, p < .001$, and a negative effect of IPTs on social judgment when target was said to have the opposite of a trait, $t = -3.62, p = .001$. When the directionality was mismatched, there was a smaller effect of IPTs on social judgment when the target had a trait, $t = 6.68, p < .001$, and a non-significant effect

when the target had the opposite of a trait, $t = -1.78$, $p = .09$. The unstandardized betas from the Hierarchical Linear Model are depicted in Figure 1.

To determine whether implicit personality theories more generally show evidence of directionality, or whether the observed effects may have been driven by the handful of trait pairs that showed group-level directionality effects, I conducted a follow-up analysis to my test of whether directional IPTs better predict social




Figure 1: The impact of implicit personality theories on target judgment as a function of target type and congruence in the directionality of the target judgment and corresponding implicit personality theory . Pattern projection is reflected by positive betas for targets that have a trait, and negative betas for targets that do not have (have the opposite of) a trait. (Study 1).

judgments than do IPTs of the reverse (mismatched) directionality. I divided the twenty relevant trait pairs into two equally-sized groups—those for which there was a larger average difference between the two possible IPTs (i.e., $p(A | B)$ and $p(B | A)$), and those for which the two IPTs were nearly equivalent. I then tested the same mixed-level model, but tested whether the influence of directionality was greater for those trait pairs that showed group-level mean differences. This 4-way interaction was non-significant, $t < 1$. This increases our confidence that IPTs, more generally, possess directionality. That is, even though many IPTs and their converse elicited similar judgments aggregated across people, the two judgments had distinct predictive power at the level of the individual.

Discussion

Study 1 establishes in two ways that IPTs are indeed directional theories and are not merely assumed correlations. Across 55 measured IPTs, the reported theory depended on the direction in which the theory was assessed. As one example, participants who considered people who were more resigned than average thought there was a 65% chance that they would also be more skeptical than average. But when participants considered people who were more skeptical than average, they thought there was only a 47% chance that they would be more resigned than average. Logically, these two beliefs are not compatible, but they do reflect the directionality of theories that underlie people's social judgments. It is easy to imagine why a person who is resigned may have grown skeptical about the world, but there is a less clear theory of why skeptical people have grown resigned (versus further emboldened). Furthermore, the directionality of theories mattered in predicting social judgments.

These findings support one condition that is necessary for my explanation of pattern projection, and confirm a prediction made by Sedikides and Anderson (1994) more than 15 years ago: “For some person types, causal connections between traits are

likely to be directional” (p. 301). Sedikides and Anderson were referring to theories that may hold together certain person types or clusters of related traits. I have established that people hold directional theories more generally: to understand the way that personality traits in general are interconnected. I have proposed that in projecting patterns, people are not projecting correlations, Instead, they are projecting causal theories. One feature of causal theories is that they have directionality; correlations do not. According to my account, people will generate theories to explain a single person (typically, the self). These egocentric theories are then exported to explain people in general. Study 1 merely establishes that IPTs reflect directional theories. Our subsequent studies examine directly whether people do indeed construct causal trait theories to explain individuals, and whether such theories underlie IPTs and thus explain pattern projection.

Study 2

When people are egocentrically pattern projecting, are they merely projecting the patterns they see in themselves onto other people? Or instead, are they actually projecting the *theories* they have developed to describe themselves? Study 1 found that implicit personality theories take the form of directional theories and are not merely “directionless” correlations that describe an association between two traits. In other words, people do not simply have a theory that trait A correlates with trait B; they have a theory of how much trait A leads to trait B, which may be distinct from how much trait B leads to trait A. If such theories are rooted in theories to explain the self, this may reflect the crucial bridge between person representations and implicit personality theories that account for pattern projection.

In Study 2, I test whether egocentric pattern projection emerges when people develop directional theories to explain the self. This explicit theory may then be projected onto people in general. For example, Christine may decide that her own

idealism about the world is what is driving her conscientiousness. Even though Christine created this explanation to explain herself, the relationship it implies may be exported to explain people in general. In this way, egocentric pattern projection may emerge from the causal trait narratives people spin to explain the self.

According to this account, it is not enough for two traits to be connected in one's causal trait narratives of the self. For example, people may explain the co-occurrence of two traits in the self not by appealing to a causal theory (A causes B), but to a third variable theory (A and B share a directionless association because both stem from C). For example, perhaps Sarah sees her own dishonesty and conscientiousness as caused by the same underlying aspect of her personality—her desire to be president of the United States. But given work showing that causal theories ($A \rightarrow B$) underlie assumed correlations (Ahn et al., 2002), one wouldn't necessarily (but could in theory) expect a theory in one's causal trait narrative that C causes both A and B (a third-variable theory) to lead to a more general assumed correlation between A and B (pattern projection). Also, given that A and B's co-occurrence is premised on something else about a person (the presence of C), it is not clear that a general A-B correlation would be projected.

In this study, I asked participants to state whether they had theories to explain why their traits related as they did. For some participants, they answered whether as part of their causal trait self-narrative, they had a causal theory for why how much they had one trait may be causing how much they had a second trait (causal theory). For other participants, they answered whether as part of their self-narratives, they had a theory that a single underlying aspect of their personality (e.g., another trait, a goal) gave rise to or caused how much they had each of the two traits (third-variable theory). I had two central predictions. First, if generating causal self-narratives gives rise to pattern projection, then I should find that for any given trait pair, those

participants who have a causal explanation for why those traits are connected in the self should be those who are pattern projecting that theory onto others. Second, I expected that this would be especially (or perhaps only) true when participants had a theory for why one trait was causing the other trait (causal theory). Given that Study 1 showed that the directionality of the theory is important, I always assessed IPTs and causal trait theories in the same direction.

Method

Participants and Design

Three hundred forty-five undergraduates at Cornell University participated in exchange for extra course credit. Participants were randomly assigned to one of two *theory type* conditions: causal or third-variable theories. The order in which participants indicated their implicit personality theories and the extent to which they thought they had each of eleven traits was counterbalanced. Participants always completed questions about their causal trait narratives last.

Self Judgments

Participants rated themselves on the same eleven personality traits used in Study 1: *bashful*, *considerate*, *dependent*, *happy-go-lucky*, *idealistic*, *persistent*, *prideful*, *reserved*, *resigned*, *skeptical*, *wordy*. For each trait, participants expressed how much that trait characterized their own personality on an eleven-point scale anchored at 1 (not at all) and 11 (extremely). As in Critcher and Dunning (2009a), I calculated absolute value difference scores (i.e., $|\text{trait}_x - \text{trait}_y|$) for each of the 55 trait pairs as a measure of how “correlated” each trait pair was in the self. A lower absolute value difference score indicates the traits occur similarly in the self, and a higher absolute value difference reflects that the traits occur in opposition in the self. Pattern projection emerges when such participants assume the traits are similarly correlated in the population-at-large.

Implicit Personality Theories

I again used single items, of the form, “If you knew that someone was more [trait X] than the average person, how likely would it be that the person was also [more Y] than the average person?” Again, I made explicit that “50% of people are above-average on any trait.” Unlike in Study 1, I also noted that if knowing someone’s standing on one trait provided no information about the person’s likely standing on a second trait, then the answer to the above question would be 50%.

Causal Trait Narratives

Instructions explained to participants the notion of causal trait narratives. Participants were told that they would provide information about the content of their own causal trait narratives. For each of fifty-five trait pairs, I assessed whether the causal trait narratives participants had of themselves included a theory that explained why two traits related the way that they did in the self. Participants answered a dichotomous-response (yes / no) questions for each trait pair. Depending on participants’ theory type condition, these questions were in one of two forms:

Causal Theories. Those in the causal theory condition were asked whether in their self-narrative, they had a causal theory explaining why how much they had one trait was causing how much they had a second trait. For example, one item read, “Does how CONSIDERATE you are cause how WORDY you are?” Participants were instructed to press ‘Y’ if “Yes, I can think of a reason why there would be a causal relationship in that direction.” Participants pressed ‘N’ if they could not think of such a reason.

Third-variable Theories. Those in the third-variable theory condition were asked whether in their causal trait narrative, they had a theory explaining why another underlying aspect of their personality (e.g., another trait, a goal) was influencing how much they had two traits. For example, one of the items read, “Does a single

underlying aspect (for example, a goal you have or another trait you possess) of your personality help explain how much you are CONSIDERATE and how much you are WORDY?” Participants pressed ‘Y’ to indicate that “Yes, I can think of an underlying aspect that would help explain why the two traits exist the way they do.” Otherwise participants were to press ‘N.’

Results

I used hierarchical linear modeling in order to assess whether pattern projection emerged more strongly for trait pairs that were causally linked in participants’ own causal trait narratives. Critcher and Dunning (2009a) demonstrated that pattern projection emerges when the participant is the unit of analysis or when the trait pair is the unit of analysis. That is, variation in how different trait pairs are patterned within a person predicts variation in that person’s corresponding implicit personality theories. Also, variation across people in how a specific pair of traits are patterned in their personalities predicts variation between people in their IPTs for that trait pair. The present analyses capitalize on the strengths of the within-subjects and within-trait-pairs analyses simultaneously. The trait difference scores were nested within the trait pair, meaning that pattern projection was defined at the level of the trait pair. But I also standardized (i.e., z-scored), at the level of the individual participant, each input into our model.

The primary impetus for this (potentially conservative) standardization was that there was great variability between participants in how many trait pairs they indicated were in their own causal trait narratives. Standardization eliminates this source of variation. Such variance in participant responses could be driven by two sources. For one, it could reflect real differences between participants in how much they tended to construct theories to explain the self (a meaningful source of variation). Instead, it could reflect differences between participants in the subjective threshold

they used in responding to the theory measures (e.g., “I’ll say YES as long as I *feel like* there is a reason even though I cannot completely articulate it” vs. “I’ll say YES only if I have spelled out in my mind a good argument that would convince the experimenter”) or other artifactual sources of variation. The former source of variance, according to our predictions, should produce real differences between people in how much they are pattern projecting. By controlling for person-level variance, I conservatively eliminate this source of variance that could be legitimately contributing to our effect. The latter (artifactual) source of variance merely adds noise that hampers our ability to detect a real effect. To the extent that the clarity gained by eliminating this noise is greater than the true “signal variance” lost, this data analytic strategy is useful.

First, to test whether there was indeed evidence of pattern projection, I tested whether the trait difference scores (nested within trait pairs) predicted variation in implicit personality theories. Replicating Critcher and Dunning (2009a), pattern projection emerged, $B = 1.62$, $t(50.89) = 7.72$, $p < .001$ ². In subsequent analyses, I will refer to the impact of the trait difference scores variable as *pattern projection*.

Second, I wanted to test whether the tendency for a participant to pattern project on a given trait pair depended on whether the participant had a causal theory in their causal trait self-narratives of why one trait caused the other. I extended on our model by adding in two additional Level 1 variables: a dichotomous *theory present* variable (z-scored at the level of the individual) indicating whether participants indicated they had a theory for that specific trait pair, and a pattern projection X theory present interaction term. As expected, the tendency to pattern project any given trait pair depended on whether participants reported having a theory explaining the co-

² Pattern projection is actually reflected by negative betas, but for ease of interpretation, all such betas, in this and all studies, have been reversed so that positive values reflect pattern projection

occurrence of the traits in the self, $B = 0.56$, $t(15,347.14) = 3.93$, $p < .001$. Because I standardized the theory present variable at the level of the participant, there is not a straightforward way to test for the degree of pattern projection observed for trait pairs for which a theory exists compared to those for which no theory exists. For illustrative purposes, I present analyses from a similar model that did not perform that standardization. From that model, I find that there is much greater pattern projection from those trait pairs for which there is a self-theory, $B = 2.19$, $t(111.22) = 8.71$, $p < .001$, compared to those trait pairs for which there is not a self-theory, $B = .92$, $t(103.05) = 3.73$, $p < .001$.

Third, I did not predict that just any causal theory linking two traits in one's self-narrative would be sufficient to produce pattern projection. Although causal theories developed to explain a specific instance can be easily generalized to become the foundation of a more general perceived correlation (Ahn et al., 2002), I did not necessarily expect to see a similar pattern for third-variable theories. I thus extended the model further by adding a new Level-1 predictor, a between-condition variable *theory type* ($-1 = \text{third-variable theory}$, $+1 = \text{causal theory}$). I also added the new interaction terms.

The predicted three-way interaction—pattern projection \times theory presence \times theory type—was significant, $B = 0.44$, $t(17,021.64) = 3.12$, $p = .002$. Figure 2 depicts this three-way interaction, showing the degree of pattern projection when a theory of each type (causal or third-variable) is present or absent. I then conducted follow-up tests to understand how the impact of theories in one's causal trait narrative on one's IPTs differed depending on the nature of the theory. As can be seen in Figure 2, participants pattern projected a trait patterning more when the pair was bound by a causal theory in one's causal trait narrative of the self, $B = 0.88$, $t(18,619.25) = 6.20$, $p < .001$. In contrast, when traits were merely linked by third-variable theories, there




Figure 2. The degree of pattern projection for trait pairs that are or are not bound by causal theories or third-variable theories in causal trait narratives of the self. (Study 2)

was no greater tendency to project those patternings, $B = 0.17$, $t(18,614.27) = 1.20$, $p > .22$.

Discussion

I have suggested that pattern projection may emerge as a byproduct of coming to understand the self. That is, people will generate causal trait narratives that explain why one aspect of their personality causes another aspect of their personality (e.g., “My being considerate leads me to be so wordy: I am always trying to word everything in the kindest light possible, even if that causes me to be long-winded.”). Because people’s causal trait narratives do not include causal theories explaining

every possible pairwise relationship between all of their traits, some trait pairs should be more likely sources for pattern projection than are others.

Confirming this hypothesis, Study 2 found that for any trait pair, those participants who pattern projected a trait patterning most strongly were those who had causally linked the two traits in their own causal trait narratives. But more specifically, it was not any type of theory that predicted increased pattern projection. Instead, causal self-theories (trait A causes trait B) predicted increased pattern projection, whereas third-variable self-theories (both trait A and trait B are caused by trait C) did not. If theories to explain the self are extended to explain people in general, it fits that theories that take the same form as implicit personality theories—directional causal theories (Study 1)—are those most likely to influence people’s IPTs.

Critcher and Dunning (2009a) highlighted that it is most important to measure pattern projection at the level of the trait pair. When pattern projection only emerges at the level of the individual, it is most likely that differences between trait pairs are responsible for artifactually producing pattern projection. It was particularly important for the present analyses that pattern projection emerged while implicit personality theories were nested within trait pair, for it helps to rule out an obvious alternative explanation for the present results. There were certain trait pairs that, across all participants, tended to be included in or excluded from people’s trait narratives. For example, 81% of participants had an explanation for why how reserved they were caused how bashful they were, whereas only 26% of participants had a theory why how dependent they were was causing how wordy they were. Had our measures been nested within participant (instead of trait pair), a plausible reverse-causality story could have artifactually produced our results. That is, it may be easier for everyone to explain why certain trait pairs are related (e.g., due to the semantic overlap of reserved and bashful), and these trait pairs co-occur similarly in everyone. These traits pairs

would then be more likely to be in participants' self-theories, and pattern projection would emerge more clearly for these pairs. This and similar alternative causal stories are made less plausible by the analytic approach employed.

If people automatically generalize the narratives they create to explain a specific person and use such narratives to explain people in general, this suggests a reason that pattern projection may be egocentric. People may tend to generate causal theories to explain the self more than they do for other people. This string of logic makes two predictions that the final three studies test. First, people should have more elaborate, accessible causal theories to explain why traits relate in the self than why traits relate in a well-known other (Studies 3 and 4). Second, pattern projection need not be egocentric. If one were to find an other about whom one had generated a causal narrative, then people may naturally pattern project from that person. Similarly, and more convincing in establishing causality, if people were prompted to generate a causal narrative to explain a social target, then people should begin to pattern project from that person (Study 5). Our account prioritizes the influence of the self only insofar as people's causal trait narratives of themselves tend to be more elaborate and comprehensive than their causal narratives of others.

Study 3

In Study 3, I asked participants to make a causal trait narrative map, either of themselves or their freshman year roommate. I chose freshman year roommates as a comparison target for two reasons. First, roommates have been used as a "familiar other" in prior research (Prentice, 1990), and thus could serve as a similar comparison target. Second, and more important, Critcher and Dunning (2009a) repeatedly established that college students pattern project from themselves but not from their freshman year roommates. Thus, if our account of pattern projection is true, people

should have more causally elaborated conceptions of the self than of their freshman year roommates.

In Study 3, participants were given a set of note cards. On each was written a different personality trait. Participants were to create a visual representation of their causal trait narrative of themselves or of their freshman year roommate. Participants drew arrows to indicate causal relationships in their visual narratives. My primary prediction was that participants' causal narratives of themselves would be more comprehensive and causally rich than their narratives of their roommates. Given that pattern projection is a particularly robust phenomenon and given Study 2's finding that pattern projection arises from causal (but not third-variable) theories, I had a secondary prediction that narratives would be characterized much more by causal theories than by third-variable theories.

Method

Participants and Design

Two hundred eight undergraduates at Cornell University participated in exchange for \$5 or extra course credit. Participants were randomly assigned to draw a visual map of their own causal narrative (*self* condition) or their narrative of their freshman year roommate (*other* condition).

Procedure

All participants began by rating themselves or their freshman year roommates on sixteen personality traits: *bashful*, *considerate*, *cunning*, *dependent*, *extravagant*, *generous*, *happy-go-lucky*, *idealistic*, *opportunistic*, *persistent*, *prideful*, *prudent*, *reserved*, *resigned*, *skeptical*, *wordy*. I was not directly interested in the ratings themselves. Instead, I wanted all participants to have formulated a clear picture of their own and their roommate's personality before reporting on their causal trait

narratives. I did not want narrative differences to emerge only because trait knowledge about the self was more accessible.

In the next step, participants were given 16 index cards. On each card was written one of the sixteen traits for which each participant had provided ratings for themselves or their roommates. At this point, participants were told that people sometimes construct narratives to explain people. The instructions explained that such narratives link together different aspects of one's personality in a causal story. To facilitate thinking about causal trait narratives, participants were first asked to look through the cards and lump together cards that would cluster together as part of the same sub-narrative. That is, it was noted that as part of a causal trait narrative, people may have one or more subnarratives that explain why different clusters of traits relate as they do in a person. As an example, it was noted that someone might lump together the cards (which no participant actually had) *carefree*, *creative*, and *extroverted* if they had a narrative causally combining the three traits:

“As an example, imagine that you are very creative, not at all extroverted, and not at all carefree. You may see that this creativity is in large part caused by all of the time you have to devote to creative projects given your lack of extroversion. Furthermore, a desire to achieve and be successful may lead you to continue to develop your creative abilities, and the seriousness with which you undertake this endeavor makes you not all that carefree.”

Those in the roommate condition saw similar instructions, though they were written to describe one's freshman year roommate. The instructions stressed that each narrative had to have at least two traits in it. It was stressed to participants that they need not use all the cards. Participants recorded which cards they placed in each cluster or sub-narrative. Although there are several interesting, measurable features of these maps (e.g., how many traits are part of narratives, how many clusters or “*sub-narratives*”

participants create), this step was largely a prelude to the next stage, in which I had people draw out more complete causal narrative maps. Nonetheless, given that prior research has found that the factor structure of self perceptions is more complex than the structure of social perceptions (Borkenau & Liebler, 1994), I expected that there would be more distinct clusters or sub-narratives for the self than for someone else.

Finally, participants were told that they would draw out their more complete causal trait narratives. Specifically, participants were told they should indicate the ways in which a specific personality trait influenced another trait, or how two traits were influenced by the same underlying aspect of personality. These of course reflect the two types of theories measured in Study 2—causal and third-variable theories, respectively. Two examples were offered to illustrate the difference between these types of theories. I illustrated a causal link with, “In me, the lack of extroversion leads to an increase in creativity.” Participants represented a directional causal link by drawing a directional arrow from one trait to another. A third-variable link was reflected by, “My desire to grow up to be a successful artist leads me to further develop my creative abilities and to spend a lot of time on solitary activities that are not very extroverted.” Participants represented a third-variable link by connecting two traits with a line, and then drawing an arrow that pointed at the line (see Figure 3). It was stressed to participants that they should not feel bound by the sub-narratives they had created in the first portion. I counted the number of directional and third-variable causal links depicted.

Results

By every metric, the trait narratives of the self were more comprehensive and causally rich than the trait narratives of the roommates. The descriptive statistics for these narratives appear in Table 2. When describing the self, people listed more clusters or sub-narratives than when describing an other, $t(202) = 2.22, p = .03$.




Figure 3. Example causal trait narrative. This narrative expresses 5 causal theories and 2 third-variable theories. Note that two traits can be bound by both types of theory (e.g., Traits 1 and 7).

Table 2: Attributes of Causal Trait Narratives Describing the Self or an Other (Study 3)

Attribute	Self	Other
Clusters / Sub-narratives	4.65 (0.99)	4.31 (1.21)
Traits used	13.40 (1.89)	12.52 (2.55)
Causal Theories	10.89 (5.03)	9.53 (3.88)
Third-Variable Theories	1.27 (1.58)	0.81 (1.19)

Note. Each mean is followed by the corresponding standard deviation. The means in each row are significantly different from one another.

Furthermore, these sub-narratives describing the self were more comprehensive: They included more of the 16 traits than those describing someone else, $t(182.26) = 2.81, p = .01$). Given Study 2 linked the presence of causal theories to pattern projection, it would be especially informative if people have more causal theories in the trait narratives of the self than in their narratives of others. And indeed, participants saw more causal relationships in the self than they did in someone, $t(206) = 2.17, p = .03$. Even though Study 2 suggested third-variable theories are not related to pattern projection, such theories were more numerous in narratives about the self than in narratives about someone else, $t(195.17) = 2.38, p = .02$. Given there is evidence that causal but not third-variable theories underlie pattern projection, it is notable that causal theories were much more numerous in the narratives of both the self, $t(105) = 20.39, p < .001$, and of the other, $t(101) = 22.32, p < .001$.

Discussion

As a first step in explaining the egocentric nature of pattern projection, I tested whether people do indeed have richer causal narratives for the self than for someone else. And indeed, narratives about the self were more comprehensive and more causally thorough. I imagine that these results provide an extremely conservative estimate of the differences in people's narratives of themselves versus other people. Even if participants did not already have well-formed narratives about their freshman year roommates, the demands of the experimental task likely led them to try and create one. It is thus notable that differences emerged nonetheless.

Even though Study 2 found that causal theories, and not third-variable theories, were a plausible prerequisite for pattern projection to emerge, it is helpful that both types of theorizing were assessed. That is, it was particularly informative not only that the causal maps of the self were more comprehensive and causally rich, but that they were composed primarily of causal and not third-variable theories. Given that self-

narratives are dominated by the types of theories that seemingly produce pattern projection, it may not be surprising that pattern projection is a robust phenomenon.

The question remains of whether participants drew concept maps that reflected pre-existing representations, or whether those causal maps were created in the moment. Study 4 tried to address this question of whether causal trait narratives of the self are not only more numerous but more likely to be preexisting (and therefore accessible) by including reaction time measures.

Study 4

Much like in the causal theory condition in Study 2, participants in Study 4 were asked fifty-five yes/no questions about whether two traits were linked causally in their causal trait narratives of themselves. Unlike in Study 2, they also answered the same questions concerning their causal trait narratives of their freshman year roommate. Also unlike in Study 2, I recorded how long it took participants to respond to each question. To the extent that a representation is already well-formed, people should be faster to indicate the presence or absence of information in that representation (Park, 1986; Prentice, 1990). I had two central predictions. First, I expected that even with this more structured and constrained measurement technique, I would once again find that people report more causal theories to explain the self than to explain someone else. Second, I expected that these causal narratives for the self would be more accessible than the narratives for the roommate. In other words, people should be faster to indicate whether they have a theory to explain the self than a theory to explain the roommate.

Method

Participants and Design

Participants were 41 undergraduates from Cornell University. In exchange for their participation, participants received \$5 or course credit. Participants answered

questions about the content of the causal trait narratives of themselves and of their freshman year roommates.

Procedure

Participants began by rating themselves and their freshman year roommates on the same eleven personality traits used in Study 2. These were completed in a counterbalanced order. These ratings were completed for a similar reason that they were in Study 3: to give participants a chance to settle on their self- and roommate-trait ascriptions before I assessed whether participants had accessible causal trait narratives for each target. I did not want the baseline accessibility of trait knowledge about the self to be able to account for differences in the accessibility of causal trait narratives.

Next, participants answered 55 questions about their self-narratives and 55 questions about their causal trait narratives of their freshman year roommates. The concept of causal trait narratives was explained in the same way as in Study 2. To assess the presence of particular theories in participants' causal trait narratives, participants answered questions of the form, "Does how SKEPTICAL you are [your roommate is] cause how PRIDEFUL you are [your roommate is]?" Responses were coded dichotomously (yes = 1; no = 0). Unlike in Study 2, participants were told that they would answer these causal theory questions about both themselves and their freshman year roommate. The order in which these were completed was counterbalanced across participants. Also, the computer recorded not only participants' responses, but also how long it took them to make each dichotomous response.

Results

I first tested whether participants' own causal trait narratives were more causally comprehensive than their narratives of their roommates. I submitted the total

number of “yes” responses to a mixed-model ANOVA, with the counterbalancing order variable as a between-subjects factor and target (self or other) as a within-subjects factor. As predicted, participants’ self-narratives were more comprehensive ($M = 28.26$, $SE = 1.42$) than their causal narratives about their roommates ($M = 24.47$, $SE = 1.44$), $F(1, 39) = 13.93$, $p < .001$. Also, participants were faster to answer questions about themselves ($M = 4.02s$, $SE = .24s$) than about their freshman year roommates, ($M = 4.36s$, $SE = .29s$), $F(1, 39) = 6.43$, $p = .02$. This accessibility difference was equally strong regardless of whether participants were indicating that a particular trait pair was or was not in their causal narrative for their roommate or for themselves, $F < 1$.

Discussion

These findings conceptually replicate and extend on those of Study 3. Not only do people possess more causally rich causal trait narratives to represent the self compared to someone else, but these theories are more accessible. It is notable that participants were faster to indicate both the presence and the absence of a theory for the self. This most likely indicates that participants had a more extensive pre-existing causal trait narrative of the self, which included (or permitted easy assessment of) both what was theorized about and that which was not. Furthermore, given that prior research has found that reaction times in this paradigm only differentiate knowledge that is core to a person representation as opposed to knowledge that is more peripheral (Park, 1986), the present findings indicate that such causal theories form a core part of people’s self-knowledge.

The accessibility findings are particularly important, because they help to rule out an alternative explanation for the results of Study 3. In Study 3, participants had little guidance or constraint as they drew their causal concept maps. The experimenter did not tell participants how long was a reasonable length of time to spend on the task.

Although this open-ended technique has the benefit of assessing spontaneously produced representations of the self versus someone else (see also Chen, 2003), the unconstrained nature of the task leaves the results open to a criticism. Participants may have been more motivated to draw a more nuanced narrative of themselves, perhaps because they enjoyed indulging in such self-thought and self-analysis. Of course, to the extent that people in their daily lives enjoy self-analysis more than other-analysis, this would not reflect a threat to internal validity as much as it would reflect an origin of these representational differences. Regardless, this alternative cannot account for the findings of Study 4. Participants took less time to report on their own self-concept. If participants walked into the experiment with similar baseline representations of the self and of their roommates, but then were more willing to take the time to construct theories to explain the self than to explain the roommate, I would have seen that it took participants longer to answer the self questions than the roommate questions.

Study 5

The present study was designed to tie together multiple strands in the present studies that have suggested that: a) the causal richness of the self-narrative explains why pattern projection emerges (Study 2), and b) that the special richness of the self-narrative compared to a narrative of someone else (Study 3 and 4) may explain the egocentric nature of pattern projection (Critcher & Dunning, 2009a). If the causal richness of the self-narrative explains why pattern projection is typically egocentric, it follows that people will pattern project from another person (and thereby reduce the egocentric nature of the effect) once they have developed a causal narrative to explain the other person's personality. This would more conclusively establish the mechanism behind pattern projection that was supported by correlational evidence in Study 2 and made plausible by the self-other differences observed in Studies 3 and 4.

In Study 5, I presented participants with trait information about three novel targets. Some participants constructed a causal narrative to explain why the traits related as they did in those targets. Other participants processed the information about these targets in a more piecemeal fashion that did not involve narrative building. Although past research has found that people have little trouble writing out descriptions of others' personalities (e.g., Park et al., 1994), no prior work has requested that people generate causal trait narratives. Nonetheless, given that people have little trouble generating ad hoc theories on demand (McNorgan, Kotack, Meegan, & McRae, 2007), I expected that participants would be able to generate causal trait narratives to explain other people.

After processing the target information in one of the two ways, participants stated their implicit personality theories. I expected that participants who generated causal narratives to explain a specific other would begin to pattern project from that person. Of course, my account does not predict that it is enough to *try* to create a causal trait narrative about someone else, but that one should actually be successful in doing so before such non-egocentric pattern projection will emerge. This predicts a more nuanced hypothesis that *only* to the extent that a participant reports success in generating a causal narrative about the target will he or she pattern project from the target. At the end of the study, participants completed a surprise memory test for the targets they had seen. This allowed me to test whether any differences between the conditions in the degree of pattern projection could be explained by differences in memory for the targets.

Method

Participants and Design

Participants were 405 undergraduates at Cornell University who completed the experiment in exchange for extra course credit. Participants were randomly assigned

to one of four conditions in a 2 (processing task: causal narrative or control) X 2 (target version) full-factorial design.

Targets

I chose twelve traits that I had used in prior studies. These traits were randomly grouped into three groups of four traits. The four traits in each group would form the basis for a description of a novel target. I then constructed two versions of each person by randomly determining, for each target, whether the target was described as “very much” having the trait, or “not at all (being the opposite of)” the trait. To minimize the likelihood that two contrasting traits would be nonsensically paired within the same person (e.g., very bashful but not at all reserved), I added the constraint that the four traits used to describe a person had to be fairly uncorrelated. (I used trait ratings from past studies to confirm that the absolute value of the correlations was less than .20.)

The two versions of the three targets are listed in Table 3. What is less important than the level of each trait in each target is how each pair of traits relates in each target. For example, even though Target 3’s skepticism and prudence differ by version, the two traits are “negatively correlated” in each version and thus do not constitute a trait pair of interest. Across the three targets, 11 of the 18 observed trait relationships differ between the two targets. If participants are pattern projecting from a target, then they should infer more of a positive correlation between two traits when the traits relate similarly in the target (“very much” – “very much” or “not at all” – “not at all”). They should infer more of a negative correlation when the traits exist dissimilarly in the target (“very much” – “not at all” or “not at all” – “very much”).

Procedure

Participants were first informed that they would receive information about three different people. This information would comprise four sentences, each conveying trait-relevant information. As seen in Table 3, each sentence was of the

Table 3: Both Versions of the Three Social Targets Presented in Study 5

Target 1 (Version A; Version B)

Person 1 is (very much NOT; very) generous.

Person 1 is very much NOT cunning.

Person 1 is very resigned.

Person 1 is very dependent.

Target 2 (Version A; Version B)

Person 2 is (very much NOT; very) happy-go-lucky.

Person 2 is very bashful.

Person 2 very much NOT prideful.

Person 2 is (very much NOT; very) idealistic.

Target 3 (Version A; Version B)

Person 3 is (very; very much NOT) skeptical.

Person 3 is (very much NOT; very) prudent.

Person 3 is very opportunistic.

Person 3 is very wordy.

Note. Trait information is listed in the order in which it was presented.

form, “Person X is [very; very much NOT] trait Y.” After each sentence appeared on the screen, 45 seconds elapsed before the next sentence would appear. Even after the next sentence would appear, the prior sentences were still visible. During the 45-sec

period, participants were to engage in one of two processing tasks, depending on their condition:

Causal Narrative. The instructions in the causal narrative condition prompted participants to generate a causal narrative explaining how the traits all influenced each other to give rise to a single, coherent individual. “Your task will be to incorporate each new piece of information you learn into a coherent picture of the person. You want to try and link together individual traits to understand how they influence or affect each other, why they fit together as they do in the same person. For example, after you learned that a target was ‘very much NOT extroverted’ but ‘very creative’, you might type that ‘Time to himself probably helps him enrich his creative inner-life. Also, since he probably isn’t that social that probably lets him devote time to artistic or whatever creative activities he is into,’” It was emphasized to participants that they should try to analyze and type for the full forty-five seconds.

Control. As a control task, it was important that participants still focus on information about the target, but I did not want them to generate causal theories about why the traits were configured in the targets in the manner stated. Accordingly, these participants were to elaborate on what it meant for the target to possess each of his or her traits. Thus, when each new sentence appeared, instead of spending 45 seconds trying to generate theories to connect the newly presented trait to the other traits, the participant spent 45 seconds elaborating on what the trait meant. “For example, if you learned a person was ‘very much NOT extroverted,’ you might type that the person is ‘a sociable, affable kind of person, interested in socializing, not at all aloof or shy, warm, gregarious...’” It was emphasized that participants were to generate these descriptions only about the most recently presented piece of information. This was stressed so participants would not think their task was to synthesize across the traits and describe what the person as a whole was like.

Procedure

Participants were told that they were taking part in the second half of a two-part experiment. In the first part, past participants had supposedly generated four traits that a person “very much was” or “very much was not (was the opposite of)”. In the present study, participants would see those descriptions. Depending on their condition, participants received one of the two sets of processing instructions described above. Then, participants were exposed to the three targets while either generating causal trait narratives to explain the targets, or by elaborating on what each of their traits meant.

After participants completed the full three-minute processing task for each target, participants were asked how difficult it was to successfully complete the processing task for that target by pressing 1 (not at all difficult), 2 (a little difficult), 3 (somewhat difficult), or 4 (very difficult). I was primarily interested in this measure in the causal narrative condition, for it provided an indication of whether participants were able to generate a causal narrative with no difficulty, or whether they found it very, or perhaps too difficult.

After participants had been exposed to all three targets, participants stated their implicit personality theories for all eighteen possible trait pairs, even though only eleven of these trait pairs would allow us to assess whether participants were pattern projecting from the targets. I measured implicit personality theories using Critcher and Dunning’s (2009a) three-judgment method—assessing $p(\text{trait 1})$, $p(\text{trait 2})$, and $p(\text{trait 1} \mid \text{trait 2})$. The IPT was derived from the following linear expression: $p(\text{trait 1} \mid \text{trait 2}) * p(\text{trait 2}) - p(\text{trait 1}) * p(\text{trait 2})$.

At the end of the experiment, participants encountered a surprise recognition task. Participants were presented with the twelve traits that had been associated with the three targets. They had to indicate whether the target in question “very much” or “very much did NOT” have the trait. In this way, I could assess whether any tendency

to pattern project differently by condition could actually be attributable to a superior explicit memory for the information about the target.

All participants had indicated their own standing on each of the traits on a web-based pretest, completed at least 24 hours before coming to the lab. These ratings were made on eleven-point scales anchored at 1(not at all me) and 11 (completely me).

Results

First, I wanted to test whether those assigned to generate a causal trait narrative of someone else would begin to pattern project from that person. Recall that there were eleven key trait pairs that differentiated the two targets. Therefore, I was interested in whether generating a causal narrative to explain a target would change participants' implicit personality theories for those key 11 trait pairs. More specifically, I expected that those who wrote a causal narrative to explain the person would state implicit personality theories that recapitulated the intertrait patterns that resided in the targets.

To test for this pattern, I first created a new variable called *Patterning*. This dichotomous variable was used to differentiate those trait pairs that implied the same direction of correlation (positive or negative) within a version. That is, I chose one of the versions, coded each trait pair implying a positive correlation as +1, and each trait pair implying a negative correlation as -1. These codings were used to describe both versions. In this way, the Patterning X Version interaction tests whether participants are pattern projecting from the targets. In describing results, I will refer to the Patterning X Version interaction as "Pattern Projection." In this way, the Patterning X Version X Processing Task interaction tests whether the degree of pattern projection from the target depends on the processing condition. Given that the past studies have indicated that variability in the way traits relate in the self accounts for variability in people's implicit personality theories, I used participants' pretest ratings of themselves

to create absolute value difference scores for all relevant trait pairs (i.e., | Self rating on trait i – Self rating on trait j |). These self-ratings were included as a covariate in all analyses.

A multi-level model was constructed to assess these hypotheses. Patterning was nested within the subject. There were four additional predictors: Version, Processing Task, Difficulty, and Self-ratings. All higher-order interaction terms were included in the full model. Replicating the standard pattern projection effect, I found a significant effect of Self-ratings, indicating that there was significant pattern projection from the self, $B = .08$, $t(4,278.00) = 5.38$, $p < .001$. But also, for the first time, I found that participants pattern projected from the targets as well, $B = .05$, $t(396.83) = 3.25$, $p = .001$. Of course, I did not expect that all participants would pattern project from the targets, and as predicted, people pattern projected from a target more once they generated a causal trait narrative of the target, $B = .03$ $t(396.83) = 1.95$, $p = .05$. Confirming the central prediction, participants who wrote causal trait narratives of the targets began to pattern project from them, $B = .08$, $t(400.51) = 3.61$, $p < .001$. But those who merely processed information about the target in a piecemeal fashion showed no hint of pattern projection from the target, $B = .02$, $t < 1$.

Of course, I am not predicting that merely trying to generate a causal trait narrative will produce pattern projection. Instead, participants would need to be reasonably successful in doing so. Thus, a Patterning X Version X Processing Task X Difficulty³ interaction would provide more detailed support for our proposed mechanism by testing whether only those who successfully generated causal narratives about a target were those who were pattern projecting from the targets. And indeed,

³ People found it marginally more difficult to complete the control task ($M = 2.70$, $SE = .037$) than the causal narrative task ($M = 2.60$, $SE = .036$), $F(1, 399) = 3.23$, $p = .07$. As such, I standardized the difficulty ratings separately by processing task condition before running the model.

this qualification by difficulty was also significant, $B = .05$, $t(1,792.36) = 3.18$, $p = .001$. I completed additional analyses to further understand the Pattern Projection X Processing Task X Difficulty interaction. I ran four additional versions of the model to assess the level of pattern projection for those who had an easy or a difficult time with each of the two processing tasks. The degree of pattern projection from the target for each difficulty / processing task combination is depicted in Figure 4. When participants completed the control processing task (and merely elaborated on what each trait of the target's in isolation), they did not pattern project, regardless of whether participants found this elaboration task difficult, $F < 1$, or simple, $F(1, 5507) = 1.12$, $p > .28$. This conceptually replicates Critcher and Dunning's (2009a) consistent failure to find pattern projection from an other. But for those who generated causal narratives about the target, they pattern projected from the target when they found it relatively easy to generate this narrative, $F(1, 5507) = 24.23$, $p < .001$, but not when they found it difficult to do, $F < 1$. No interaction terms with the self-ratings approached significance.

These findings still leave open the question of whether writing a causal trait narrative about a person influenced the way people thought about people in general (i.e., their implicit personality theories) for an artifactual reason—enhanced memory for details of the target. Past research has found that people have enhanced memory for information about people when they are trying to form an impression of them than when they are merely trying to memorize the information about them (Hamilton, Katz, & Leirer, 1980). And indeed, participants in the causal narrative conditions did have a better memory for the targets' standing on the 12 traits ($M = 10.19$, $SD = 2.06$) than those in the control condition ($M = 9.20$, $SD = 2.40$), $t(386.53) = 4.46$, $p < .001$. But there was no evidence that superior memory for the trait information was the mediator responsible for the impact of causal trait narrative generation on pattern projection.




Figure 4. Pattern projection from the novel social targets as a function of processing task condition and the difficulty participants had with the processing task. High and Low difficulty is predicted at ± 1 standard deviation from the mean level of self-reported difficulty in the relevant processing task condition (Study 5).

When memory was added to the model, it was not the case that accurately recalling the way that traits were positioned in a target increased the chance that participants would pattern project from that target, $B = -.02$, $t(2,772.94) = 1.30$, $p > .19$.

Discussion

Study 5 makes two central contributions. First, it demonstrates that generating a causal narrative about someone *causes* one to pattern project from that person. In

this sense, the story we tell about a person, not merely the self, becomes a story about human nature. Study 2 hinted at this causal connection by showing that individual differences in the presence of these self-narratives predicted the degree of pattern projection from the self. In the present study, I *manipulated* whether people generated causal, explanatory narratives, and this led people to pattern project from the person about whom those narratives were generated.

Second, the study finds a theoretically predicted limit to the egocentric nature of pattern projection. Studies 3 and 4 began to answer the question of why pattern projection is egocentric—people had more comprehensive and more accessible narratives about themselves than they did about someone else. Even though people may tend to have fewer narratives about other people, the present study shows that when participants were led to construct a narrative about someone else, they began to pattern project from that person. In this way, by decomposing the origin of pattern projection into processes that could characterize people's thinking about anyone, we begin to see the circumstances in which pattern projection will emerge from others as well. Note this clarifies that there is a natural limit to the egocentric nature of pattern projection. When people generate causal trait narratives about others, pattern projection should occur not only from the self but from these others as well. Future research could ask *who* is likely to generate causal narratives about *what kinds of people* and *in what circumstances*. In answering this, researchers would have a better idea of what social targets besides the self tend to inform people's basic beliefs about human nature.

Chapter 3: General Discussion

Critcher and Dunning (2009a) provided evidence of pattern projection—a qualitatively novel way in which the self influences social judgment—but conceded that it was unclear exactly why pattern projection occurs. The present studies answered the question left open by past research, and in so doing demonstrated heretofore unidentified representational bases of implicit personality theories and of the self. First, I found that implicit personality theories reflect directional theories; they do not simply reflect assumed correlations (Study 1). This suggests that when people pattern project, they are not projecting a correlation from a sample of one onto people in general. Instead, they are projecting a causal theory that explains a sample of one, thereby applying that theory to explain people in general. Consistent with this, and second, I found that when people have an explicit theory to explain why two traits co-occur in the self (i.e., a causal trait narrative), they are then more likely to project that patterning onto people in general (Study 2). That is, the narrative used to explain the self becomes the narrative used to explain people in general. These two findings suggest an answer to why pattern projection is egocentric: People have more elaborated (Studies 3 and 4) and more accessible (Study 4) causal trait narratives to explain the self than to explain a familiar other. But once people are prompted to generate a narrative to explain someone else, they begin to pattern project from this person as well (Study 5). Thus, although pattern projection may tend to be an egocentric phenomenon, people will begin to pattern project from an other once they have developed a causal trait narrative to understand this other person. In this way, coming to understand a person—whether that be the self or someone else—changes the way one thinks about people in general.

Why are causal trait narratives more elaborated for the self than for others?

Critcher and Dunning (2009a) demonstrated egocentric pattern projection, but did not demonstrate why pattern projection emerged or why it was egocentric. The current research answers these questions, but similarly raises new ones. The present results show that pattern projection emerges as a byproduct of generating causal trait narratives to explain the structure of someone's personality, and accounts for the egocentric nature of pattern projection by showing that such causal trait narration is egocentric. But it does not answer why people generate more elaborate trait narratives for the self than for other people. One intuitively appealing answer, but one that I find ultimately incomplete, is that the self is the object of its own thoughts more than are other people. In all of that egocentric thought, there would simply be more of a chance for people to elaborate causal narratives to understand the self. Although there is no doubt some truth in this statement, it seems unlikely to offer a complete explanation given that Critcher and Dunning (2009a) consistently found that people displayed no hint of pattern projection from their freshman year roommates. It seems that over the course of a year of living in the same room as someone else, people certainly would have spent considerable time thinking about that familiar other. Thus, even if pattern projection were egocentric—driven more by self-views than views of others—one might still expect a (reduced) effect of people's roommates on their implicit personality theories.

A second possibility is that people may be more motivated to construct self-narratives, because such narratives may have a perceived functional benefit of assisting in producing long-term personality change. Given that traits are assumed to be fairly stable labels of personality, causal trait narratives may be of little use in predicting immediate behaviors, suggesting they may not be very useful in forecasting the short-term behavior of other people. In other words, in the short run, causal trait narratives may provide little additional predictive power above and beyond the trait

labels themselves. But if people wish to change unpleasant aspects of their own personality, a causal trait narrative would help to focus people on the deeper causes of their own imperfections. This logic would explain why people would construct narratives to explain negative aspects of one's personality, but given that both the present research and Critcher and Dunning (2009) examined pattern projection from relatively neutral traits (Anderson, 1968), this idea cannot account for the documented examples of pattern projection.

A third possibility, and one that I favor, stems from research showing that trait-based understanding of the self is more nuanced than it is of others (Sander, Goethals, & Radloff, 1988). This nuance stems from people realizing that people do not always behave consistently with their traits, but display some variety across situations (Baxter & Goldberg, 1987). For example, even a person who considers himself extroverted can think of certain contexts in which he is more reserved. And when social expectations or desires are violated, people spontaneously engage in causal narrative building to construct a story to make sense of the discrepancy (Kunda, Miller, & Claire, 1990; Risen, Gilovich, & Dunning, 2007). For example, Charlie and his friends may all consider Charlie to be conscientious. But what Charlie realizes, but few other people know, is that although Charlie is extremely conscientious at work, he is disorganized with his finances. Thus, even though Charlie may consider himself a conscientious person, the exceptions to this label may prompt Charlie to narrate that it is his career aspiration that drive his conscientiousness. Thus, people may make sense of the imperfectness of their own trait labels by appealing to causal trait narratives (see Chen, 2003).

If this explanation is true, it provides insight into when we would expect that people would generate causal narratives to explain other people. Prentice (1990) described that with increasing familiarity about other people, one is likely to encounter

more and more inconsistent information about the person. She speculated that person understanding may then include explanations or theories that bring coherence to a diverse set of information. Given that this notion of Prentice's is compatible with our conception of causal trait narratives, it may be that we have causal trait narratives for individuals with whom we have sufficient familiarity that we are aware of behavioral exceptions to their traits.

But note that while knowledge of others' inconsistencies may prompt *some* causal theorizing about other people, it is unlikely to be as comprehensive as the causal theorizing we generate to understand ourselves. Given that we observe ourselves in a wider variety of contexts than we observe anyone else, we are more likely to know our own inconsistencies. People may believe they know everything about their spouses or their children, but these others may be very different people while at school or work. Also, any narrative we develop to explain someone else may provide more of a sense of explanatory completeness than a similar narrative about the self would, limiting the likely scope of our narratives of others. We look inward and see complex selves, with much hidden beneath the surface, whereas we see any information we receive about others as providing a more complete picture of their personalities (Pronin, Kruger, Savitsky, & Ross, 2001). For example, Kunda et al. (1990) found that when participants learned of a gay construction worker, they developed the causal narrative that the man's homosexuality prompted him to compensate by seeking out a traditionally masculine career. I suspect that those participants saw such a narrative as a fairly complete description of the person, whereas I suspect that the construction worker would find such a description of himself as horribly incomplete. Because pattern projection is typically assessed across a large variety of randomly selected traits, pattern projection from a target would be detected only to the extent that the social perceiver has constructed comprehensive

narratives to explain the target. Incomplete or simplistic narratives (e.g., “Everything about her stems from her being a jerk”) will not give rise to pattern projection.

What makes it into a person’s causal trait narrative?

Given that people’s causal trait narratives will not be entirely comprehensive, one can ask what the nature of people’s own narratives is likely to be. In the present paper, I have focused on the existence of and a consequence of having causal trait narratives. Left unexamined is the pattern to the content of these narratives. Three questions naturally emerge. First, what trait pairs are included in causal trait narratives, and what pairs are not? Second, are there certain trait pairs that are particularly likely to be included in one’s causal trait narratives of the self (compared to narratives of other people)? Third, although Study 3 found that directional theories are more pervasive than third-variable theories, it remains unknown which traits are given causal priority in narratives, and which are most easily explained by other traits and goals.

To examine whether certain trait pairs are especially likely or unlikely to be included in causal trait narratives, I returned to Study 4 to conduct some exploratory analyses on the trait narratives of the self and of one’s freshman year roommate. Of course, some traits are highly correlated, and thus it may be obvious why they are related. For example, the semantic overlap of *bashful* and *reserved* lead the two traits to be highly correlated ($r = .54$ in Study 4), and a majority of people (73%) had a theory for why one trait was causing the other. So what is perhaps most interesting are trait pairs that are especially likely or unlikely to be included in trait narratives, even controlling for the degree to which they are correlated.

In order to identify such pairs, I calculated—for both the self judgments and the roommate judgments—the (absolute value) correlations between the two traits that formed each trait pair and the percentage of participants reporting that they had a

causal theory to explain the way the traits related in the self as they did. Reflecting the way that causal theories underlie the understanding of perceived correlations (Ahn et al., 2002), the more correlated two traits were, the more likely participants were to have a causal theory connecting them ($r = .55$, for the self; $r = .59$, for the roommate). I then regressed the percentage of participants who claimed to have a theory for each trait pair on the correlation of each trait pair, separately for the self and for the roommate. The residuals reflect the degree to which trait pairs are “overincluded” or “underincluded” in narratives.

Given the pattern of the residuals was quite similar between the model predicting self narratives and that predicting roommate narratives ($r = .82$), I summed across these residuals to get a clearer idea of which trait pairs are over- or underincluded in people’s causal trait narratives. Listed in Table 4 are the five most over-included and under-included trait pairs. Instead, one can subtract the residuals from the analyses on roommate narratives from the residuals for the same model about the self. This shows what trait pairs are especially unlikely or especially likely to appear in a self narrative (compared to a roommate narrative). Listed in Table 5 are the five most over-included and under-included trait pairs in self narratives (compared to roommate narratives).

I hesitate to draw firm conclusions on observed patterns, because the ability to detect narratives was of course constrained by the trait pairs I happened to measure. Nonetheless, there are some hints of patterns that I expect may be more general. For example, participants seemed to have a surprising number of theories in which *skeptical* was a causal antecedent and *persistent* was a consequent. *Skepticism* seems to be more of an internal trait, characterizing the way one processes information about the world. *Persistence*, in contrast, refers to the intensity of one’s behavioral responses. Thus, I expect that there may be a more general tendency to create causal

narratives in which internal traits give rise to external traits. It is also notable that the two theories that are most overrepresented in self-narratives explain whether one is

Table 4: Directional Trait Theories Most Overincluded and Underincluded in Causal Trait Narratives (Study 4)

Theories Occurring in Causal Trait Narratives More Than Expected

1. Prideful → Persistent
2. Idealistic → Persistent
3. Skeptical → Persistent
4. Bashful → Wordy
5. Skeptical → Idealistic

Theories Occurring in Causal Trait Narratives Less Than Expected

1. Dependent → Wordy
 2. Dependent → Idealistic
 3. Skeptical → Bashful
 4. Idealistic → Bashful
 5. Considerate → Bashful
-

Note. Expected degree of inclusion is based on the actual sample correlation between the traits. Theories are ranked based on the extremity of the residuals from regressing the proportion of participants reporting having such a theory in their causal trait narratives on the absolute value of the correlation of the two trait ratings. Note that each theory is not that having one trait causes having the second trait, but that how much a person has one trait causes how the person has the second trait.

Table 5: Causal Trait Theories Most Included in Self (versus Other) and Other (versus Self) Causal Trait Narratives (Study 4)

Theories More Likely To Appear in Self Narratives

1. Prideful → Wordy
2. Bashful → Wordy
3. Skeptical → Considerate
4. Dependent → Idealistic
5. Reserved → Prideful

Theories More Likely To Appear in Other Narratives

1. Dependent → Considerate
 2. Resigned → Happy-go-lucky
 3. Reserved → Dependent
 4. Skeptical → Bashful
 5. Idealistic → Wordy
-

Note. Expected degree of inclusion is based on the actual sample correlation between the traits. Theories are ranked based on the extremity of the residuals from regressing the proportion of participants reporting having such a theory in their causal trait narratives on the absolute value of the correlation of the two trait ratings. Note that each theory is not that having one trait causes having the second trait, but that how much a person has one trait causes how the person has the second trait.

wordy. Wordiness may be a trait for which people are particularly attentive to variation in the self, but less so to variation in the behavior of others. If self theorizing

is especially comprehensive because we are especially attentive to inconsistencies in our own behavior, it makes sense that people would be especially likely to generate theories about themselves to explain their own level of wordiness. Future research should be designed to identify more specific patterns in people's causal trait narratives.

Why does social experience not eliminate egocentric implicit personality theories?

Given that people have substantial experience with others, why is it that such experience does not force them to realize that their own egocentric narratives may not be applicable to people in general? One possibility is that people do not carefully attend to others, or to all aspects of others, equally. Explanatory theories focus attention on compatible feature co-occurrences (Murphy & Wisniewski, 1989). Thus, egocentric causal trait narratives may direct attention toward others who are compatible with the narrative, and this selective attention may provide biased support for one's own initial theories

Second, it takes more to directly contradict a projected pattern than a projected attribute. This is because it takes two pieces of disconfirming evidence to violate a pattern. For example, if David sees himself as liberal and warm, as soon as he learns that Thad maxed out on his contribution to John McCain's election campaign, David will no longer see Thad as a liberal, and attributive projection is blocked. In contrast, David can continue to pattern project onto Thad if David decides that Thad must be a cold-hearted conservative. That is, by shifting his view of Thad on the second characteristic (warmth), there remains consonance between the pattern David observed in Thad (conservative and cold) and the pattern David observes in himself (liberal and warm). As such, even as people learn more about social targets, they can revise their impressions so that egocentric implicit personality theories are not violated. Consistent with this possibility, even as the academic year progressed, students living in a

residence hall showed no increase in their consensus in judgments of each others' personalities (Park & Kraus, 1993; cf. Funder, Kolar, & Blackman, 1995), and pattern projection onto such well-individuated targets remains strong (Critcher & Dunning, 2009a). Thus, greater experience with other people does not communicate unambiguous information about those targets' personality.

Third, even when people construct causal trait narratives of other people, these may not water down the egocentric nature of pattern projection because such social narratives may be egocentric themselves. Park et al. (1994) found that participants provided with the exact same information about a target developed very different narratives to understand the target. The authors reported that they could not find any individual difference variable that was able to predict what narrative participants used to organize their information about the novel targets. According to our research, people may have gravitated toward social narratives that were most consistent with their own self narratives. In this way, people may choose to structure their narratives about other people in a way that confirms the universality of patterns observed in the self.

Fourth, causal trait narratives may not only guide our attention to or understanding of other people, but the narratives themselves may serve to hold one's self-views in place (see Critcher & Dunning, 2009b). Although past work has found that there is surprisingly low intertemporal consistency in people's own life narratives (McAdams et al., 2006), the epistemic function of causal trait narratives likely makes them more stable. That is, although the perceived centrality of defining episodes in one's own life may wax and wane, there is little reason to believe that people need to continually revise their explanations for why the different aspects of their personalities relate as they do. Such stability may hold people's trait self-perceptions in place. That is, once people have developed a reason why their resigned attitude gives rise to their

dependent interpersonal style, the theory may cement people's self-perception. This means that self-views may become rigid and thus flawed, but such fixed self-views provide constant validation for one's own theories.

Finally, people do not come to understand themselves in a vacuum. Self-understanding is a social process. Life narratives are socially constructed (Adler & McAdams, 2007; Pasupathi, 2001), and causal trait narratives likely are as well. People likely try out their causal trait narratives on other people, who are known to provide an important validating function for the narratives people construct (Pasupathi & Rich, 2005). People may confuse others' supportive reactions to causal trait narratives of the self for feedback that the narratives are the accurate theories of human behavior.

Interpersonal Trait Narratives, Interpersonal Pattern Projection

Causal trait narratives need not refer only to intrapersonal trait dynamics. They could refer to interpersonal trait dynamics as well. People will include their conceptions of close others in the self (Aron, McLaughlin-Volpe, Mashek, Lewandowski, Wright, & Aron, 2004; Wright, Aron, & Tropp, 2002), and members of collectivist cultures naturally have a more expansive view of the self (Markus & Kitayama, 1991). Thus, people's narratives may expand to include features of others as well. Such narratives may include explanations for why traits in the self relate to, influence, or are influenced by traits in close others.

Interpersonal trait narratives may then produce pattern projection at the dyadic level. For example, romantic couples may generate theories for why a trait in one partner has given rise to a trait in the second partner. The couple may then generalize these theories and use them as bases for expectations about new couples they meet. In this way, interpersonal pattern projection will look similar to standard pattern projection, except the two component traits reside in separate people instead of in the

same person. But interpersonal pattern projection also permits projection of the same traits. For some traits, members of a dyad will display similarity, whereas on other traits they will evince complementarity. Couples may expect to see the patterns in their own dyad recapitulated in other couples. This may be especially true to the extent that the couples have created narratives to explain their own dyadic patterns.

Conclusion

The present research has both advanced and bridged our understanding of the nature of implicit personality theories and self-conceptions. People's beliefs as lay personality psychologists and their beliefs about themselves as people have at their core causal theories. Pattern projection is but one phenomenon that results from the construction of causal trait narratives to understand the dynamics of a person, especially the self. Future research will no doubt uncover additional consequences of this heretofore unidentified type of person representation.

REFERENCES

- Adler, J. M., & McAdams, D. P. (2007). Time, culture, and stories of the self. *Psychological Inquiry, 18*, 97-128.
- Ahn, W., Marsh, J. K., Luhmann, C., & Lee, K. (2002). Effect of theory-based feature correlations on typicality judgments. *Memory & Cognition, 30*, 107-118.
- Andersen, S. M., & Ross, L. (1984). Self-knowledge and social inference: I. The impact of cognitive/affective and behavioral data. *Journal of Personality and Social Psychology, 46*, 280-293.
- Anderson, C. A., & Sedikides, C. (1991). Thinking about people: Contributions of a typological alternative to associationistic and dimensional models of person perception. *Journal of Personality and Social Psychology, 60*, 203-217.
- Anderson, N. H. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology, 9*, 272-279.
- Anderson, N. H. (1981). *Foundation of information integration theory*. New York: Academic Press.
- Aron, A., McLaughlin-Volpe, T., Mashek, D., Lewandowski, G., Wright, S. C., & Aron, E. N. (2004). Including close others in the self. *European Review of Social Psychology, 15*, 101-132.
- Asch, S. E., & Zukier, H. (1984). Thinking about persons. *Journal of Personality and Social Psychology, 46*, 1230-1240.
- Asch, S. B. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology, 41*, 258-290.
- Baumeister, R. F., & Newman, L. S. (1994). How stories make sense of personal experiences: Motives that shape autobiographical narratives. *Personality and Social Psychology Bulletin, 20*, 676-690.

- Baxter, T. L., & Goldberg, L. R. (1987). Perceived behavioral consistency underlying trait attributions to oneself and another: An extension of the actor-observer effect. *Personality and Social Psychology Bulletin*, 13, 437-447.
- Beer, A., & Watson, D. (2008). Asymmetry in judgments of personality: Others are less differentiated than the self. *Journal of Personality*, 76, 535-559.
- Black, J. B., Galambos, J. A., & Read, S. J. (1984). Comprehending stories and social situations. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 3, pp. 45-86). Hillsdale, NJ: Lawrence Erlbaum.
- Borkenau, P., & Liebler, A. (1994). The factor structure of trait ratings depends on the extent of information available to the judges. *European Review of Applied Psychology*, 44, 3-7.
- Bruner, J. (1990). *Actual minds, possible worlds*. Cambridge, MA: Harvard University Press.
- Bruner, J. S., & Taiguri, R. (1954). The perception of people. In G. Lindsay (Ed.), *Handbook of social psychology* (Vol. 2). Cambridge, MA: Addison-Wesley.
- Chapman, L. J., & Chapman, J.P. (1967). Genesis of popular but erroneous diagnostic observations. *Journal of Abnormal Psychology*, 72, 193-204.
- Chen, S. (2003). Psychological-state theories about significant others: Implications for the content and structure of significant-other representations. *Personality and Social Psychology Bulletin*, 29, 1285-1302.
- Critcher, C. R., & Dunning, D. (2009a). Egocentric pattern projection: How implicit personality theories recapitulate the geography of the self. *Journal of Personality and Social Psychology*, 97, 1-16.
- Critcher, C. R., & Dunning, D. (2009b). How chronic self-views influence (and mislead) self-assessments of performance: Self-views shape bottom-up experiences with the task. *Journal of Personality and Social Psychology*, 97,

931-945.

- D'Anrade, R. G. (1974). Memory and the assessment of behavior. In H. M. Blalock (Ed.), *Measurement in the social sciences*. Chicago: Aldine-Atherton.
- Dawes, R. M. (1990). The potential nonfalsity of the false consensus effect. In R. M. Hogarth (Ed.), *Insights in decision making: A tribute to Hillel J. Einhorn* (pp. 179–199). Chicago: University of Chicago Press.
- Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgments or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes*, 65, 201–211.
- Dunning, D., & Cohen, G. L. (1992). Egocentric definitions of traits and abilities in social judgment. *Journal of Personality and Social Psychology*, 63, 341–355.
- Dunning, D., & Hayes, A. (1996). Evidence for egocentric comparison in social judgment. *Journal of Personality and Social Psychology*, 71, 213–229.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 57, 1082–1090.
- Eidelman, S., & Biernat, M. (2007). Getting more from success: Standard raising as esteem maintenance. *Journal of Personality and Social Psychology*, 92, 759–774.
- Epley, N., Converse, B.A., Delbosc, A., Monteleone, G., & Cacioppo, J. (2009). Believers' estimates of God's beliefs are more egocentric than estimates of other people's beliefs. *Proceedings of the National Academy of Sciences*, 106, 21533–21538.

- Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology, 87*, 327-339.
- Epley, N., Morewedge, C., & Keysar, B. (2004). Perspective taking in children and adults: Equivalent egocentrism but differential correction. *Journal of Experimental Social psychology, 40*, 760-768.
- Funder, D. C., Kolar, D. C., & Blackman, M. C. (1995). Agreement among judges of personality: Interpersonal relations, similarity, and acquaintanceship. *Journal of Personality and Social Psychology, 69*, 656-672.
- Goldings. (1954). On the avowal and projection of happiness. *Journal of Personality, 23*, 30-47.
- Habermas, T., & Bluck, S. (2000). Getting a life: The emergence of the life story in adolescence. *Psychological Bulletin, 126*, 748-769.
- Hamilton, D.L., Katz, L.B., & Leirer, V.O. (1980). Cognitive representation of personality impressions: Organizational processes in first impression formation. *Journal of Personality and Social Psychology, 39*, 1050-1063.
- Hampson, S. E. (1998). When is an inconsistency not an inconsistency? Trait reconciliation in personality description and impression formation. *Journal of Personality and Social Psychology, 74*, 102-117.
- Hastie, R., & Pennington, N. (1991). Cognitive and social processes in decision making. In L. Resnick, J. Levine, & S. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 308-327). Washington, DC: American Psychological Association.
- Holmes, D. S. (1981). Existence of classical projection and the stressreducing function of attributive projection: A reply to Sherwood. *Psychological Bulletin, 90*, 460-466.

- Judd, C. M., Kenny, D. A., & Krosnick, J. A. (1983). Judging the positions of political candidates: Models of assimilation and contrast. *Journal of Personality and Social Psychology*, 44, 952–963.
- Katz, D., & Allport, F. (1931). *Students' attitudes*. Syracuse, NY: Craftsman Press.
- Kim, M.P., Rosenberg, S. (1980). Comparison of two structural models of implicit personality theory. *Journal of Personality and Social Psychology*, 38, 375–89.
- Krueger, J., & Stanke, D. (2001). The role of self-reference and otherreferent knowledge in perceptions of group characteristics. *Personality and Social Psychology Bulletin*, 27, 876–888.
- Kruger, J., & Gilovich, T. (2004). Actions, intentions, and trait assessment: The road to self-enhancement is paved with good intentions. *Personality and Social Psychology Bulletin*, 30, 328-339
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, 14, 551-577.
- Kunda, Z., & Nisbett, R. E. (1986). Prediction and the partial understanding of the law of large numbers. *Journal of Experimental Social Psychology*, 22, 339-354.
- Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 22, 754-770.
- Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning & Verbal Behavior*, 23, 250-269.
- Markus, H., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, 98, 224-253.
- McAdams, D. P. (1985). *Power, intimacy, and the life story: Personological inquiries into identity*. New York: Guilford Press. McAdams, D. P. (2001). The psychology of life stories. *Review of General Psychology*, 5, 100–122.
- McAdams, D. P. (1995). What do we know when we know a person? *Journal of*

Personality, 63, 363-396.

McAdams, D. P. (2001). The psychology of life stories. *Review of General Psychology*, 5, 100-122.

McAdams, D. P., Bauer, J. J., Sakaeda, A. R., Anyidoho, N. A., Machado, M. A., Magrino-Failla, K., White, K.W., & Pals, J. L. (2006). Continuity and change in the life story: A longitudinal study of autobiographical memories in emerging adulthood.

McGuire, W. X., & McGuire, C. V. (1986). Differences in conceptualizing self versus other people as manifested in contrasting verb types used in natural speech. *Journal of Personality and Social Psychology*, 51, 1135-1143.

McNorgan, C. M., Kotack, R. A., Meehan, D. C., & McRae, K. (2007). Feature-feature causal relations and statistical co-occurrences in object concepts. *Memory & Cognition*, 35, 418-431.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.

Murphy, G. L., & Wisniewski, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol. 2. Theory and applications* (pp. 23-45). Chichester, U.K.: Ellis Horwood.

Pals, J. L. (2006). Constructing the “springboard effect”: Causal connections, self-making, and growth within the life story. In D. P. McAdams, R. Josselson, and Lieblch (Eds.), *Identity and story: Creating self in narrative* (pp. 175–199). Washington, DC: American Psychological Association Press.

Park, B. (1986). A method for studying the development of impressions of real people. *Journal of Personality and Social Psychology*, 51, 907-917.

Park, B., DeKray, M. L., & Kraus, S. (1994). Aggregating social behavior into person models: Perceiver-induced consistency. *Journal of Personality and Social*

Psychology, 66, 437-459.

Park, B., & Kraus, S. (1993). *Longitudinal trends in impression consensus*.

Unpublished manuscript, University of Colorado

Pasupathi, M. (2001). The social construction of the personal past and its implications for adult development. *Psychological Bulletin*, 127, 651–672.

Pasupathi, M., & Rich, B. (2005). Inattentive listening undermines self-verification in personal storytelling. *Journal of Personality*, 73, 1051–1085.

Prentice, D. A. (1990). Familiarity and differences in self- and other-representations. *Journal of Personality and Social Psychology*, 59, 369-383.

Preston, J., & Epley, N. (2005). Explanations versus applications: The explanatory power of valuable beliefs. *Psychological Science*, 16, 826-832.

Pronin, E., Kruger, J., Savitsky, K., & Ross, L. (2001). You don't know me, but I know you: The illusion of asymmetric insight. *Journal of Personality and Social Psychology*, 81, 639-656.

Read, S. J., & Cesa, I. L. (1991). This reminds me of the time when . . . : Expectation failures in reminding and explaining. *Journal of Experimental Social Psychology*, 27, 1-25.

Risen, J. L., & Gilovich, T., Dunning, D. (2007). One-shot illusory correlations and stereotype formation. *Personality and Social Psychology Bulletin*, 33, 1492-1502.

Rosenberg, S. (1976). New approaches to the analysis of personal constructs in person perception. In A. W. Lanfield (Ed.), *Nebraska symposium on motivation* (pp. 179-242). Lincoln: University of Nebraska Press.

Ross, L., Greene, D., & House, P. (1977). The false consensus phenomenon: An attributional bias in self-perception and social perception processes. *Journal of Experimental Social Psychology*, 13, 279-301.

- Sande, G. N., Goethals, G. R., & Radloff, C. E. (1988). Perceiving one's own traits and others': The multifaceted self. *Journal of Personality and Social Psychology*, 54, 13-20.
- Schneider, D. J., Hastorf, A. H., & Ellsworth, P. C. (1979). *Person perception*. Reading, MA: Addison-Wesley.
- Sedikides, C., & Anderson, C. A. (1994). Causal perceptions of intertrait relations: The glue that holds person types together. *Personality and Social Psychology Bulletin*, 20, 294-302.
- Tversky, A., & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp. 84-98). Hillsdale, NJ: Erlbaum.
- Watson, D. (1982). The actor and the observer: How are their perceptions of causality divergent? *Psychological Bulletin*, 92, 682-700.
- Wishner, J. (1960). Reanalysis of "impressions of personality." *Psychological Review*, 67, 96-112.
- Wright, S. C., Aron, A., & Tropp, L. R. (2002). Including others (and groups) in the self: Self-expansion and intergroup relations. In J. P. Forgas & K. D. Williams (Eds.), *The social self: Cognitive, interpersonal and intergroup perspectives* (pp. 343-363). Philadelphia: Psychology Press.