Research Report

Psychological Science
2016, Vol. 27(10) 1398–1406

DOI: 10.1177/0956797616656800
pss.sagepub.com

Beliefs About the Causal Structure of the Self-Concept Determine Which Changes Disrupt Personal Identity

Stephanie Y. Chen, Oleg Urminsky, and Daniel M. Bartels
Booth School of Business, University of Chicago

Abstract

Personal identity is an important determinant of behavior, yet how people mentally represent their self-concepts and their concepts of other people is not well understood. In the current studies, we examined the age-old question of what makes people who they are. We propose a novel approach to identity that suggests that the answer lies in people's beliefs about how the features of identity (e.g., memories, moral qualities, personality traits) are causally related to each other. We examined the impact of the causal centrality of a feature, a key determinant of the extent to which a feature defines a concept, on judgments of identity continuity. We found support for this approach in three experiments using both measured and manipulated causal centrality. For judgments both of one's self and of others, we found that some features are perceived to be more causally central than others and that changes in such causally central features are believed to be more disruptive to identity.

Keywords

causal reasoning, personal identity, self-concept, open data, open materials

Received 12/31/15; Revision accepted 6/6/16

Personal identity provides people with norms to follow, scripts for behaviors, and ways to interpret their actions (Akerlof & Kranton, 2000, 2010; Markus & Wurf, 1987; Turner, 1985) and affects a wide range of in-lab and real-world decisions (e.g., Bertrand, Kamenica, & Pan, 2015; Cohn, Fehr, & Maréchal, 2014). In particular, a sense of continuity in one's identity (e.g., connectedness, Bartels & Rips, 2010) provides motivation for making far-sighted choices (Bartels & Urminsky, 2011, 2015) and a sense of disruption in another's identity is related to relationship deterioration (Strohminger & Nichols, 2015).

In this article, we discuss how people represent the self and others, and the features of identity that people believe must be retained for continuity of identity. Prior researchers have debated whether social categories, memories, tastes, personality traits, or moral qualities are most defining of identity (e.g., Blok, Newman, & Rips, 2005; Haslam, Bastian, & Bissett, 2004; Strohminger & Nichols, 2014). To reconcile these approaches to identity, we propose that beliefs about causal relationships among features influence which features are perceived to be most defining of identity. In our approach, unlike those of previous researchers, we do not assume that a single type of feature is most important.

Building on the concepts literature in cognitive psychology, we propose that people's representations of identity incorporate the causal relationships among the features of identity (Sloman, Love, & Ahn, 1998). In general, features that are more causally central (i.e., linked to many other features of a concept or network; Bonacich & Lu, 2012; Pennington & Hastie, 1988; Rehder & Hastie, 2001) are more defining of a concept. We propose that people reason about their self-concepts and concepts of other people in much the same way that they reason about concepts in general. Accordingly, we predicted that

Corresponding Author:
Stephanie Y. Chen, Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637
E-mail: stephanie.chen@chicagobooth.edu

people would believe that causally central features are more defining of identity. For example, the importance of memories, traits, or preferences for the self-concept depends on how these features (and other features of identity) are causally related. Although ideas about causal centrality are extremely influential in the concepts literature, these explorations have mainly tested these ideas in artificial or common everyday categories (but see Kim & Ahn, 2002). We tested these ideas in real-world, highly individualized concepts and incorporated causal centrality into a theory of personal identity for the first time.

In the first two experiments, we measured beliefs about the causal relations between features of identity and the extent to which changes in these features disrupt identity, for the self or for other people. If more causally central features are more defining to identity, as hypothesized, the number of causal connections that a feature has should be positively correlated with perceived disruption to identity resulting from a change in the feature.

In a third experiment, we manipulated the centrality of features to further test the hypothesis that changes to features with more causal connections are perceived as more disruptive to identity. We also examined an alternative approach to centrality, the dependency model (Sloman et al., 1998). This model suggests that centrality depends on a feature's causal depth—a measure of all the feature's direct and indirect downstream effects. We tested whether this alternative approach to centrality explained how causal beliefs influence identity judgments.

Experiment 1

Method

The power analysis from a pilot experiment (for details, see Appendix S1 in the Supplemental Material available online) suggested a sample size of 80 per cell. Two hundred fifty Amazon Mechanical Turk respondents from the United States were randomly assigned to one of three conditions (self, close other, or generic other). Five participants were excluded because of a scripting error, 4 because they failed an attention check, and 2 because they gave identical answers to every question, for a total of 11 exclusions. This resulted in a final sample of 239. Results were similar when we included all participants who provided usable data (for details, see Appendix S1).

All participants completed both a causal-relationships task and an identity-disruption survey. To measure centrality, we asked participants in the self condition to report the causal connections among the features of their own identity. To measure perceived disruption to identity, we asked these participants to rate the extent to which a change in each identity feature would disrupt their own identity. Participants in the close-other condition did the same for a nonromantic close other whom they specified. Participants in the generic-other condition completed the tasks for a generic other person. The order of the tasks was counterbalanced across participants within each condition.

In the causal-relationships task, participants reported the causal relationships that they perceived among 16 features of personal identity (see Table 1). Twelve of the 16 features were intended to be of high importance and were chosen from four categories of personal identity that had been identified as important in previous research: autobiographical memories, personality, morality, and preferences and desires (e.g., Strohminger & Nichols, 2014). The remaining 4 were intended to be of low importance. Two were found in previous research to be less important for identity (instances of semantic memories; Strohminger & Nichols, 2014), and 2 (fillers) were found to be unimportant to identity in a pretest.

After practicing the causal-relationships task with an unrelated concept and receiving feedback, participants completed 16 randomized trials. In each trial, a different feature was the target. Participants indicated which of the other 15 features, if any, was caused by the target feature (see Fig. 1). Then, for each feature selected as a direct effect, participants rated the strength of its relationship (1 = *weak*, 2 = *moderate*, 3 = *strong*) with the target feature.

In the identity-disruption survey, participants rated the extent to which change in each feature would

Table 1. Features Presented to Participants in Experiment 1

Autobiographical memories
Cherished memories of time with parents/family
Important childhood memories
Memories of important life milestones
Morality
Level of wholesomeness
Level of honesty
Level of loyalty
Personality
Intelligence level
Degree of shyness
Reliability
Preferences and desires
Goals for personal life
Favorite hobbies/activities
Aesthetic preferences
Semantic memories
Knowledge of math
Knowledge of music
Fillers
Height
Level of hunger

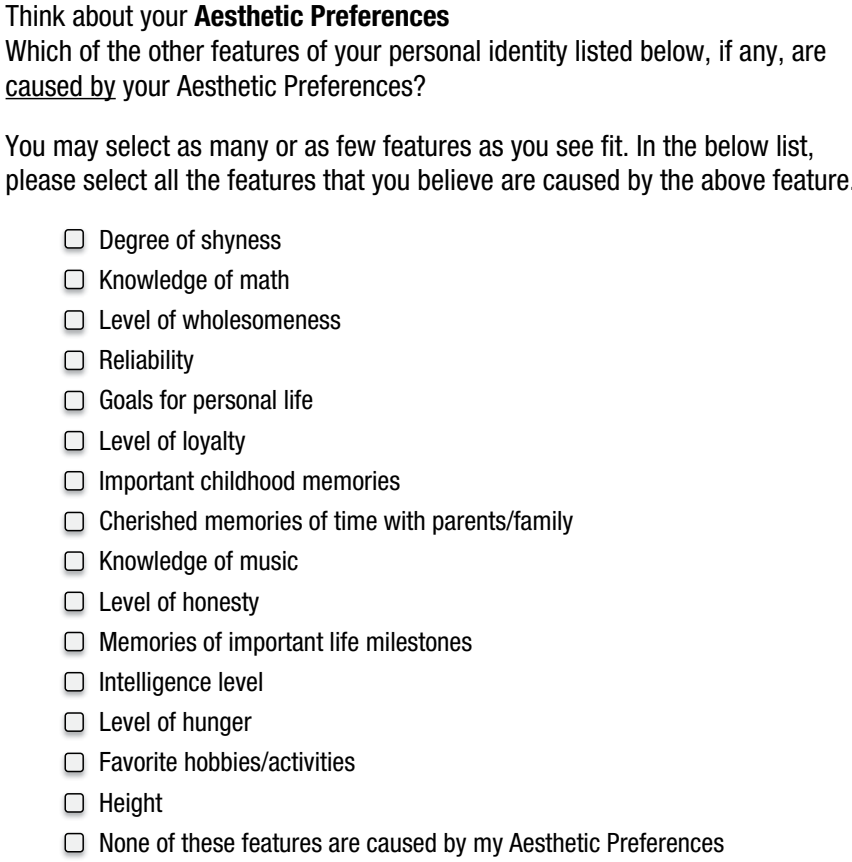

Fig. 1. Example of a question from the causal-relationships task used in Experiment 1.

disrupt the identity of the person that corresponded to their condition (i.e., self, close other, or generic other). They rated disruption on a scale of 0 to 100; larger numbers indicated greater disruption. (For the wording of the question, see Appendix S1 in the Supplemental Material.)

Results

Our analyses used the number of causal connections (i.e., the number of other features to which a target feature was directly linked, either as a cause or as an effect) as the measure of causal centrality. More links indicated greater centrality. Our findings were similar when we used an alternative approach, causal depth (the dependency model; Sloman et al., 1998), as the measure of causal centrality (for details of analysis and results, see Appendix S2 in the Supplemental Material).

On average, participants reported 37.9 causal links among the 16 features of identity. The number of links did not differ by condition (self condition: $M = 35.6$; close-other condition: $M = 39.1$; generic-other condition: $M = 38.9$), $F(2, 236) = 0.69$, $p > .250$, which suggests that participants perceived that other people's personal identities were as complex as their own.

An analysis of variance[1] with condition as a between-subjects factor and feature as a repeated measures factor revealed that the number of causal connections differed across features, $F(17, 2113) = 156.34$, $p < .001$. There was also a significant Condition × Feature interaction, $F(17, 2113) = 1.62$, $p = .049$, suggesting that differences in causal connections across the features varied by condition (for results by condition, see Appendix S1 in the Supplemental Material). As expected, features selected as being of low importance were less central than those identified by prior research as being of high importance (low importance: $M = 2.0$; high importance: $M = 5.6$), $F(1, 42) = 64.48$, $p < .001$, and this difference did not vary by condition, $F(1, 42) = 0.04$, $p > .250$.

Overall, changes in features with more causal connections were rated as more disruptive to identity in the self and close-other conditions and marginally more disruptive in the generic-other condition (self condition: $r_s = .60$, $p = .015$; close-other condition: $r_s = .62$, $p = .013$; generic-other condition: $r_s = .44$, $p = .093$).[2] There were no significant differences between conditions regarding

the perceived relationship between a feature's causal centrality and the extent to which a change to it would be disruptive (self condition vs. close-other condition: p = .842; self condition vs. generic-other condition: p = .171; close-other condition vs. generic-other condition: p = .116).

Likewise, the average individual-level correlations (within individual participants, across all items) were positive in all conditions (see Table 2). The majority of participants in all conditions rated changes in the features with more causal connections as being more disruptive (r_s was positive for 77%, 84%, and 74% of participants in the self, close-other, and generic-other conditions, respectively). Similar results were found when we analyzed only the 12 high-importance features (see Table 2). The relationship between causal connections and disruptiveness of change in the self condition was replicated in another experiment using a task in which participants drew the causal connections among the same 16 features of identity (for method and results of this experiment, see Appendix S1 in the Supplemental Material).

The strength of the individual-level correlations between causal connections and disruptiveness of change across features did not differ by condition, $F(2, 236) = 0.81$, $p >$ .250 (see Table 2). These results suggest that the perceived strength of the relationship between a feature's causal centrality and the extent to which that feature defines identity is similar for the self and others.

Experiment 2

In Experiment 1, using prespecified features from prior literature, we found that changes in more causally central features were seen as more disruptive to identity. In Experiment 2, we tested whether these findings would generalize to important features of identity generated by participants.

Method

A power analysis based on the results of Experiment 1 suggested sample sizes of 95 per cell. Two hundred two Amazon Mechanical Turk respondents in the United States were randomly assigned to one of two conditions (self or close other). Excluding 13 participants (5 failed an attention check and 8 gave the same answers to all questions) yielded a sample size of 189. Similar results were found when we included all participants who provided usable data (for details, see Appendix S1 in the Supplemental Material).

Participants' first task was to generate 16 important features for either their own identity or for a close other's identity. Participants listed the 3 most important features in each of the following categories: memories, goals and desires, preferences, moral qualities, and the four most important personality traits. To keep the level of specificity similar for all feature types, we asked participants to describe how each moral quality or personality trait was expressed (e.g., for the humor trait, a participant stated that he or she "jokes all the time"). These specific descriptions were used as the features. As in the self and close-other conditions in Experiment 1, participants then performed the causal-relationships task and the identity-disruption survey using the 16 self-generated features. The order of these two tasks was counterbalanced across participants. In Experiment 2, however, we also instructed participants that "cause" meant that a feature shaped or influenced another feature.

Results

On average, participants reported 35.3 causal links among the 16 features of identity. The number of links was lower in the self condition (M = 31.8) than in the close-other condition (M = 38.6), $t(187) = 2.39$, p = .018,

Table 2. Results From Experiment 1: Correlations Between a Feature's Number of Causal Connections and Ratings of the Extent to Which Change in That Feature Would Disrupt Identity

Condition	Aggregate Spearman correlation: all features	Mean individual-level Spearman correlation	
		All features	High-importance features[a]
Self	r_s = .60, p = .015	r_s = .34, 95% CI = [.25, .44], $t(78) = 7.29$, $p < .001$	r_s = .15, 95% CI = [.05, .25], $t(77) = 3.04$, p = .003
Close other	r_s = .62, p = .013	r_s = .38, 95% CI = [.29, .46], $t(78) = 9.08$, $p < .001$	r_s = .17, 95% CI = [.09, .25], $t(77) = 4.38$, $p < .001$
Generic other	r_s = .44, p = .094	r_s = .30, 95% CI = [.20, .39], $t(80) = 6.28$, $p < .001$	r_s = .10, 95% CI = [.00, .20], $t(79) = 1.94$, p = .056

Note: The t values are from one-sample t tests of the mean r_s (with Fisher transformation) and 0. CI = confidence interval.

[a]Data from 1 participant from each condition could not be included in the analyses of high-importance features because he or she gave the same disruption-to-identity ratings for all high-importance features.

Table 3. Results From Experiment 2: Correlations Between a Feature's Number of Causal Connections and Ratings of the Extent to Which Change in That Feature Would Disrupt Identity

Condition	Mean individual-level Spearman correlation
Self	r_s = .24, 95% CI = [.17, .31], $t(91) = 6.54$, $p < .001$
Close other	r_s = .21, 95% CI = [.13, .30], $t(96) = 5.00$, $p < .001$

Note: The t values are from one-sample t tests of the mean r_s (with Fisher transformation) and 0. CI = confidence interval.

95% confidence interval (CI) for the difference between conditions = [−12.54, −1.19]. This finding suggests that, on average, participants perceived greater complexity among the features they generated for a close others' personal identity than among the features they generated for their own identity.

As in Experiment 1, changes in features with more causal connections were rated as more disruptive to identity. Because participants generated different features, aggregate correlations were not calculated. For the majority of participants in all conditions (76% and 60% in the self and close-other conditions, respectively), there were positive individual-level correlations between features' connections and the rated disruptiveness of change. On average, individual-level correlations were positive in all conditions (see Table 3). There was no significant difference between the two conditions' average individual-level correlations, $t(187) = 0.43$, $p > .250$. This again suggests that the relationship between causal connections and identity disruption is perceived to be similar in magnitude for the self and for others. As in Experiment 1, similar results were found when we used an alternative approach, causal depth (the dependency model; Sloman et al., 1998), as the measure of causal centrality (for details of analysis and results, see Appendix S2 in the Supplemental Material).

Experiment 3

The previous studies found strong correlational evidence that causal centrality determines the extent to which a feature defines identity. In Experiment 3, we manipulated the centrality of features in vignettes to test whether making a feature more causally central affected how defining that feature was for identity. We also tested whether the causal-connections approach or the causal-depth approach better captured how causal beliefs influence identity judgments. The causal-depth approach suggests that features that have more direct and indirect effects (i.e., that are deeper in the causal chain) are more causally central. Thus, this approach, unlike the causal-connections approach, suggests that causes are more important than their effects (Sloman et al., 1998; for details, see Appendix S2 in the Supplemental Material).

Method

In prior research (Ahn, Kim, Lassaline, & Dennis, 2000), manipulated centrality had a large effect on the extent to which features influenced categorization judgments (d = 0.8). Power analysis suggested a sample size of 22 per cell, so we set a target of approximately 30 per set of vignettes. Sixty Amazon Mechanical Turk participants in the United States were randomly assigned to read one of two sets. Removing 4 participants for failing either a comprehension check or an attention check left 56 participants for analysis. Similar results were found when we included all participants in the analysis (for results with all participants, see Appendix S1 in the Supplemental Material).

The first aim of Experiment 3 was to manipulate the centrality of features to test whether making a feature more causally central affected the extent to which the feature was perceived to define identity. We constructed vignettes that described the causal relationships among four salient features of a person in a common-cause structure. For example, one vignette described four of Jack's features as being related to one another via a single cause—Jack's memories of being a lonely child caused his shyness, his preference for solitary activities, and his awkward demeanor (Fig. 2, Version A).

To manipulate whether a given feature was causally central or peripheral, we created two versions of each vignette. In the other version of the vignette, the position of two target features (shyness and memories) were flipped so that Jack's shyness caused his memories, preferences, and demeanor (Fig. 2, Version B). Thus, the same features were counterbalanced to be either the causally central cause feature (memories in Version A and shyness in Version B) or the causally peripheral effect feature. This was done to control for any idiosyncratic influences of specific features.

The focal task involved selecting which individual—one missing the effect feature (e.g., shyness in Version A) and one missing a cause feature (e.g., memories in Version A)—was more likely to be the character in the vignette. Given that the cause feature in these vignettes is involved in more causal connections and comes earlier in the causal chains, both approaches to causal centrality make the same prediction: Retaining the cause feature should be more important for continuity of identity. Thus, we predicted that participants would choose the individual who was missing the effect feature (and retained the cause feature) as the one more likely to be the character in the vignette.

The second aim of Experiment 3 was to understand which approach to causal centrality better described how

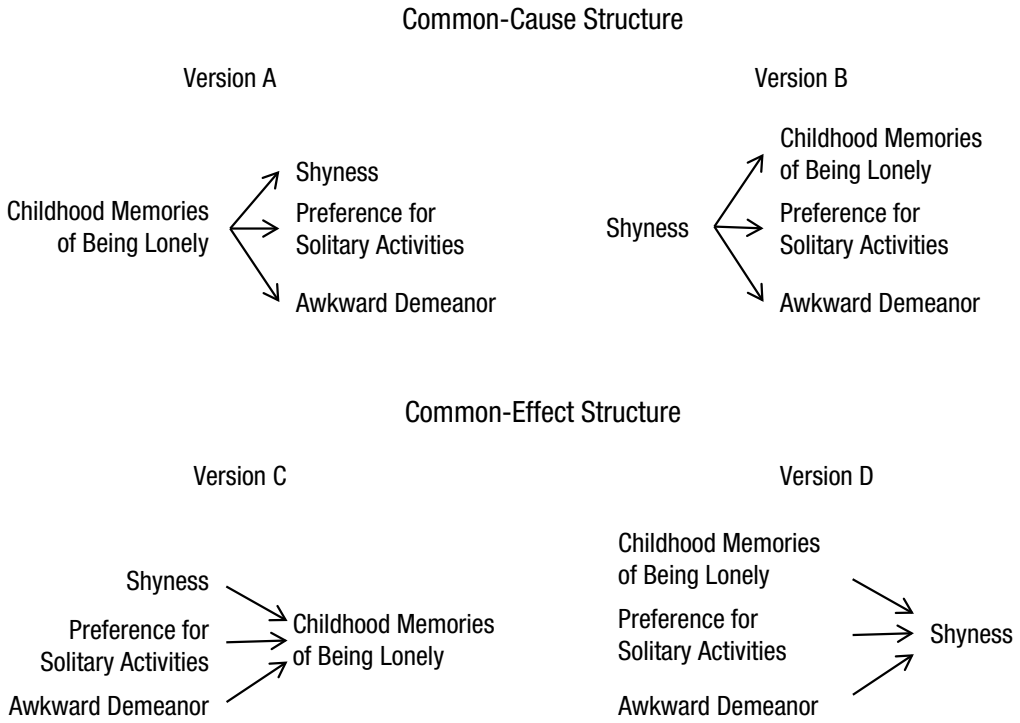

Fig. 2. Illustration of the structure of the vignettes used in Experiment 3. There were four versions of each vignette, two versions for each causal structure. The two versions for the common-cause structure referred to the same four features but differed in which feature was described as the cause of the other three (in the example shown here, the common cause is childhood memories of being lonely in Version A and shyness in Version B). The two versions for the common-effect structure also referred to the same four features but differed in which feature was described as an effect of the other three (in the example shown here, the common effect is childhood memories of being lonely in Version C and shyness in Version D).

causal beliefs influenced identity judgments. To do this, we created common-effect versions of all the vignettes (one effect with three causes). For example, the common-effect version of Jack's vignette presented childhood memories as an effect of his other three features (including shyness) rather than as the cause (Fig. 2, Version C). As with the common-cause vignettes, we created two versions of each common-effect vignette to counterbalance the position of two target features in the causal structure. The other common-effect version of Jack's vignette presented shyness as an effect of his other three features, including his memories (Fig. 2, Version D).

The common-effect structure allowed us to distinguish between the two approaches to causal centrality, which make different predictions for these versions. The cause features are deeper in the causal chain than the effect features. Thus, according to the causal-depth approach, when participants select which individual is the character in the story, they should prefer the individual missing the effect (by this definition, the more peripheral feature) to the individual missing the cause. In contrast, according to the causal-connections approach, participants should pick the person missing the cause feature because the effect feature is linked to all three cause features, whereas each cause is linked to only one other feature, the effect feature.

We constructed six vignettes that described the causal relationships among four different features of a person's identity. Each vignette had four versions (two common cause, two common effect) with the same four features, thereby counterbalancing which of the two focal features was a cause and which was an effect (see Fig. 2).

The vignettes were split into two sets (Sets 1 and 2). Each set contained the two common-cause versions for three vignettes and the two common-effect versions for the other three vignettes (e.g., for the Jack vignette, Set 1 contained Versions A and B, Set 2 contained Versions C and D; see Fig. 2). Participants were randomly assigned to Set 1 or Set 2 and then, for each of the six vignettes, to one of the versions included in that set. That is, participants read only one version of each of the six vignettes. Diagrams like those in Figure 2 accompanied the vignettes.

After reading each vignette, participants completed a comprehension check to confirm that they understood the causal structure. Participants then selected which of two people—one missing the focal cause feature and one missing the focal effect feature—they believed was most

likely to be the character in the vignette. Participants then rated the plausibility of the vignette on a scale from 0 (*not at all plausible*) to 100 (*extremely plausible*). We wanted to ensure that participants made a careful choice; thus, they were shown the two people again, presented in a different spatial layout, and were asked to report the person who they had previously selected.

Results

We excluded trials in which participants failed the comprehension check (12% of trials) or provided inconsistent answers regarding which person was the vignette character (5% of trials). Results were similar when we performed the analysis with no trials excluded (for results, see Appendix S1 in the Supplemental Material).

The dependent measure was the average of the individual-level proportion of trials in which the participant selected the person missing the effect feature. In the common-cause trials, the cause feature should be more central than the effect feature (i.e., the cause feature is connected to more features and is deeper in the causal chain) according to both approaches. Therefore, participants should pick the person missing the effect feature, and they did (M = .70, SD = .31, 95% CI = [.62, .78]), $t(55)$ = 4.76, $p < .001$. This result was replicated in another experiment using a different number of features (for details on this experiment, see Appendix S1 in the Supplemental Material).

However, the two approaches to centrality yield different predictions in the common-effect trials. According to the causal-connections approach, a missing effect should disrupt identity more than a missing cause because the effect feature has more causal connections than the cause feature does. In contrast, according to the causal-depth approach, a missing cause should disrupt identity more than a missing effect because order in the causal chain is what matters.

Our results are more consistent with the predictions of the causal-connections approach. In the common-effect condition, participants selected the person missing the effect feature—the feature that had more connections but was less deep in the causal chain than the cause feature—at lower than chance levels (M = .39, SD = .32, 95% CI = [.31, .47]), $t(55) = 2.60$, p = .012. On average, participants selected the person missing the effect feature significantly more in the common-cause condition (M = .70) than in the common-effect condition (M = .39), $t(55)$ = 4.90, $p < .001$, which is consistent with predictions of the causal-connections approach.

The results of both conditions replicate the prior findings using an experimental manipulation of causal centrality. When a feature had more causal connections, changes in that feature were perceived as being more inconsistent with continuity of identity. The findings were further moderated by participants' perceptions of the plausibility of each vignette; this finding is consistent with research that suggests our concepts, in general, are influenced by our prior knowledge (Murphy, 2002; Murphy & Medin, 1985). Ratings of vignette plausibility were correlated with the average proportion of trials in which participants selected the person missing the causally peripheral feature with fewer connections (r = .64, 95% CI = [.32, .83], $p < .001$). The average proportion of trials in which participants selected the person missing the causally peripheral feature was significantly higher among the 12 most plausible vignettes (M = .74, SD = .16) than among the 12 least plausible vignettes (M = .56, SD = .20), $t(22) = 2.51$, p = .020, 95% CI for the difference between the most and least plausible vignettes = [.03, .34], which suggests that participants' use of the causal information from the vignettes was moderated by how it fit their beliefs about which causal relationships are relatively more likely to occur.

Discussion

People perceived more causally central features as being more necessary for continuity of identity, both for the self and for others (Experiments 1 and 2). Furthermore, when we experimentally increased a feature's causal centrality, perceptions of the extent to which that feature defined identity also increased (Experiment 3).

Prior research has focused on comparing the individual importance of different types of features. These approaches seem to have missed the critical aspect of people's beliefs about the causal relationships among features; such relationships influence the extent to which a feature is perceived to define identity. The incorporation of causal beliefs into a theory of personal identity is consistent with people's general drive to explain the world using causal relations (Gopnik, 1998; Keil, 2006) and with narrative-based views of identity (McAdams, 2001, 2013).

Experiment 3 found that the causal-connections approach better described how causal beliefs influenced identity judgments. Likewise, in Experiment 1 and the pilot experiment, although both the number of causal connections and causal depth related to the extent to which a feature was perceived to define identity, only causal connections remained significant in a multiple regression (for multiple regression results of Experiment 1, see Appendix S2 in the Supplemental Material; for results of the pilot experiment, see Appendix S1 in the Supplemental Material). This suggests that features that cause many other features or are caused by the combination of many other features (or both) will be most defining of identity. Thus, changes to or the addition of features may be less disruptive when people can causally connect these new aspects to existing identity features. In fact,

prior research has found that students whose personal narratives included more causal descriptions of experienced changes had greater emotional stability (Lodi-Smith, Geise, Robins, & Roberts, 2009).

Differences in beliefs about the causal structure of identity may have important implications for identity-based motivations for behavior. If people who anticipate disruptions to more causally central features are less connected to their future selves, they may make more short-sighted decisions (Bartels & Urminsky, 2011). The effectiveness of interventions that appeal to identity features (e.g., Bryan, Walton, Rogers, & Dweck, 2011) may depend on the causal centrality of the targeted feature.

People's representations of themselves and others are not simply a list of features or social categories. These representations incorporate beliefs about the causal relations among aspects of identity. The answer to the riddle of who people are lies at the nexus of causal connections among their features of identity.

Action Editor

Marc J. Buehner served as action editor for this article.

Author Contributions

S. Y. Chen and D. M. Bartels developed the study concept. All the authors contributed to the study design. S. Y. Chen conducted the data collection and data analysis. S. Y. Chen and O. Urminsky wrote the manuscript, and D. M. Bartels provided critical revisions. O. Urminsky and D. M. Bartels contributed equally to this work.

Acknowledgments

We thank Douglas Markant and Sarah Molouki for helpful comments and Halley Bayer for help with data collection.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported by John Templeton Foundation Grant No. 40357.

Supplemental Material

Additional supporting information can be found at http://pss.sagepub.com/content/by/supplemental-data

Open Practices

All data have been made publicly available via the Open Science Framework and can be accessed at https://osf.io/u9bsv/. The complete Open Practices Disclosure for this article can be found at http://pss.sagepub.com/content/by/supplemental-data. This article has received badges for Open Data and Open Materials. More information about the Open Practices badges can be found at https://osf.io/tvyxz/wiki/1.%20View%20the%20Badges/ and http://pss.sagepub.com/content/25/1/3.full.

Notes

1. The results reported were Huynh-Feldt corrected when sphericity could not be assumed.
2. The correlations reported for Experiments 1 and 2 are Spearman's rank-order correlations. Fisher transformations were performed before *t* tests.

References

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*, 361–416.

Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, *115*, 715–753.

Akerlof, G. A., & Kranton, R. E. (2010). *Identity economics: How our identities shape our work, wages, and well-being*. Princeton, NJ: Princeton University Press.

Bartels, D. M., & Rips, L. J. (2010). Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, *139*, 49–69.

Bartels, D. M., & Urminsky, O. (2011). On intertemporal selfishness: How the perceived instability of identity underlies impatient consumption. *Journal of Consumer Research*, *38*, 182–198.

Bartels, D. M., & Urminsky, O. (2015). To know and to care: How awareness and valuation of the future jointly shape consumer savings and spending. *Journal of Consumer Research*, *41*, 1469–1485.

Bertrand, M., Kamenica, E., & Pan, J. (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics*, *130*, 571–614.

Blok, S., Newman, G., & Rips, L. (2005). Individuals and their concepts. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin* (pp. 127–149). Washington, DC: American Psychological Association.

Bonacich, P., & Lu, P. (2012). *Introduction to mathematical sociology*. Princeton, NJ: Princeton University Press.

Bryan, C. J., Walton, G. M., Rogers, T., & Dweck, C. S. (2011). Motivating voter turnout by invoking the self. *Proceedings of the National Academy of Sciences, USA*, *108*, 12653–12656.

Cohn, A., Fehr, E., & Maréchal, M. A. (2014). Business culture and dishonesty in the banking industry. *Nature*, *516*, 86–89.

Gopnik, A. (1998). Explanation as orgasm. *Minds and Machines*, *8*, 101–118.

Haslam, N., Bastian, B., & Bissett, M. (2004). Essentialist beliefs about personality and their implications. *Personality and Social Psychology Bulletin*, *30*, 1661–1673.

Keil, F. C. (2006). Explanation and understanding. *Annual Review of Psychology*, *57*, 227–254.

Kim, N. S., & Ahn, W. (2002). Clinical psychologists' theory-based representations of mental disorders predict their diagnostic reasoning and memory. *Journal of Experimental Psychology: General*, *131*, 451–476.

Lodi-Smith, J., Geise, A., Robins, R. W., & Roberts, B. W. (2009). Narrating personality change. *Journal of Personality and Social Psychology*, *96*, 679–689.

Markus, H., & Wurf, E. (1987). The dynamic self-concept: A social psychological perspective. *Annual Review of Psychology*, *38*, 299–337.

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

McAdams, D. P. (2013). The psychological self as actor, agent, and author. *Perspectives on Psychological Science*, *8*, 272–295.

Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

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

Pennington, N., & Hastie, R. (1988). Explanation-based decision making: Effects of memory structure on judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 521–533.

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, *130*, 323–360.

Sloman, S., Love, B., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, *22*, 189–228.

Strohminger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, *131*, 159–171.

Strohminger, N., & Nichols, S. (2015). Neurodegeneration and identity. *Psychological Science*, *26*, 1468–1479.

Turner, J. C. (1985). Social categorization and the self-concept: A social cognitive theory of group behavior. In E. J. Lawler (Ed.), *Advances in group processes: Theory and research* (Vol. 2, pp. 77–122). Greenwich, CT: JAI Press.

Appendix S1

Causal Depth Analyses

To calculate causal centrality based on the causal depth of a feature, we used the dependency model of causal centrality (Sloman et al., 1998). According to this iterative model, c_i, the centrality of feature *i*, is determined (at each time step) by summing across the centrality of the concept's other features (at time, *t*), $c_{j,t}$, multiplied by how dependent each feature, *j*, is on feature *i*, d_{ij}:

$$c_{i,t+1} = \sum_j d_{ij} c_{j,t}. \quad (1)$$

The implementation of the model is a repeated matrix multiplication that converges on a stable ranking within a small number of iterations (Sloman et al., 1998; Kim & Park, 2009). The ConceptBuilder software (Kim & Park, 2009) used in the pilot experiment (see Appendix S2) performs 15 iterations and the initial centrality of all features (at time 0) is set to 0.5. All causal depth analyses follow this convention.

As d_{ij} is a positive value when feature *i* causes feature *j*, according to this model, the deeper a feature is in the causal chain, the more defining it is to the concept. In the concept map task and the listing causal relationships task, the dependence, d_{ij}, is the value (1 = weak, 2 = moderate, 3 = strong) that participants assigned to the strength of each causal relationship they drew or reported.

Causal Depth Analysis – Experiment 1

The results of the causal depth correlational analyses revealed a similar pattern to the causal connections analyses. The overall correlation between causal depth and disruption to identity was significant in the close-other condition, and marginally significant in the self and generic-

other conditions (see Table S1). The results of the individual-level analysis revealed that the Spearman correlation for causal depth was positive for all conditions (M_{self} = .26, $M_{close\text{-}other}$ = .32, $M_{generic\text{-}other}$ = .22), $t(78) = 5.14$, $t(78) = 6.61$, $t(80) = 4.78$, ps < .001[1]. The majority of participants in all conditions had a positive individual-level correlation between features' causal depth and rated disruptiveness of change (72%, 80%, and 72% in the self, close-other, and generic-other conditions, respectively). A one-way ANOVA revealed that the mean Spearman correlation did not differ by condition, $F(2, 238) = 1.14$, p >.250, suggesting that the relationship between casual depth and disruption to identity was similar across conditions (see Table S1).

Table S1

Correlations Between a Feature's Causal Depth and Ratings of How Disruptive Change in that Feature Would Be

	Causal Depth Approach	
Condition	Aggregate Spearman Correlation	Individual Spearman Correlation
Self	r_s = .49, p = .05	mean r_s = .26, t(78) = 5.15, p < .001, 95% CI = [.16, .36]
Close-other	r_s = .65, p = .01	mean r_s = .32, t(78) = 6.61, p < .001, 95% CI = [.22, .42]
Generic-other	r_s = .42, p = .11	mean r_s = .22, t(80) = 4.79 p < .001, 95% CI = [.13, .31]

Note: The t values are from one-sample t tests of the mean r_s (with Fisher transformation) and 0. CI = confidence interval.

Causal Depth vs. Causal Connections Approaches – Experiment 1

To examine the relative impact of the two forms of causal centrality on disruption to identity scores, for each subject, we regressed identity disruption ratings on both measures of causal centrality (all measures z-scored within subject). We performed a 3 (condition: self, close-other,

[1] Fisher transformations were performed prior to t-tests.

generic-other) × 2 (causal centrality approach: causal connections vs causal depth) ANOVA on the resulting betas. Neither the main effect of condition, $F(2, 236) = 1.08$, $p > .250$, nor the condition × causal centrality approach interaction, $F(2, 236) = .71$, $p > .250$, were significant, suggesting that condition did not influence the predictive value of these two causal centrality measures.

We found a main effect of causal centrality approach, $F(1, 236) = 36.13$, $p < .001$. For all conditions the mean regression coefficient for the causal connections term ($M_{self} = .26$, $M_{close\text{-}other} = .32$, $M_{generic\text{-}other} = .31$) was significantly positive, ts > 5.45, ps < .001, and greater than the mean coefficient for the causal depth term ($M_{self} = .04$, $M_{close\text{-}other} = .03$, $M_{generic\text{-}other} = -.05$), ts > 2.8, ps < .01 (see Table S2). The mean coefficient for the causal depth term was not significantly positive for any condition, ts < 1.03, ps > .250. The regression analysis suggests that the causal depth approach does not significantly add to the predictive power of a model that includes the number of causal connections.

Table S2

Summary of Experiment 1 Regression Results

	Mean Coefficient		
Condition	Causal Connections Term, *M (SD)*	Causal Depth Term, *M (SD)*	t-tests Comparing Terms
Self	.26 (.42) 95% CI=[.16, .35]	.04 (.37) 95% CI=[-.05, .11]	$t(78) = 2.89$, $p = .005$
Close-other	.32 (.42) 95% CI=[.23, .42]	.03 (.45) 95% CI=[-.07, .13]	$t(78) = 3.43$, $p = .001$
Generic-other	.31 (.45) 95% CI=[.21, .41]	-.05 (.46) 95% CI=[-.16, .05]	$t(80) = 4.03$, $p < .001$

Causal Depth Analysis – Experiment 2

Using causal depth as an alternative measure of causal centrality, we again found that causal information influenced identity judgments. The Spearman rank correlation coefficient was

positive for the majority of participants in the self and close-other conditions (76% and 70%, respectively). The average correlation coefficient was positive between causal centrality and disruption to identity for both conditions (M_{self} = .25, 95% CI = [.18 .32]; $M_{close\text{-}other}$ = .24, 95% CI = [.16 .32]), $t(91) = 6.8$, $t(96) = 5.78$, ps < .001.

An independent-samples t-test revealed that the mean Spearman correlation did not differ by condition, $t(187) = .25$, p >.250, suggesting that the relationship between casual depth and disruption to identity was similar across conditions.

Causal Depth vs. Causal Connections Approaches – Experiment 2

To examine the relative impact of the two forms of causal centrality on disruption to identity, for each subject, we regressed identity disruption ratings on both measures of causal centrality (all measures z-scored within subject). We performed a 2 (condition: self, close-other) × 2 (causal centrality approach: causal connections vs causal depth) ANOVA on the resulting betas. The main effect of condition was not significant, $F(1, 187) = .66$, $p > .250$, nor was the condition × causal centrality approach interaction, $F(1, 187) = .57$, $p > .250$, suggesting that condition did not influence the predictive value of these two causal centrality measures.

There was no main effect of causal centrality approach, $F(1, 187) = 1.64$, $p = .202$, suggesting the predictive value of the two causal centrality approaches did not differ. Overall, the mean coefficient for the causal connections term (M = .15, 95% CI = [.08 .21]) and the causal depth term (M = .07, 95% CI = [.01 .13]) were both significantly positive, $t(188) = 4.29$, p<.001, $t(188) = 2.25$, p<.025. (See Table S3 for detailed results.)

Table S3

Summary of Experiment 2 Regression Results

	Mean coefficient		
Condition	Causal Connections Term, *M (SD)*	Causal Depth Term, *M (SD)*	t-tests Comparing Terms
Self	.17 (.50) 95% CI=[.07, .28]	.06 (.43) 95% CI=[-.02, .15]	$t(92) = 1.22$, $p = .225$
Close-other	.12 (.44) 95% CI=[.03, .21]	.08 (.43) 95% CI=[-.01, .16]	$t(96) = .60$, $p > .250$

Level of Description Analysis – Experiment 2

The features of identity reported in Experiment 2 varied on how abstract of concrete they are as well as whether they were plural or singular. To ensure that these factors did not influence our results, we had an independent coder (blind to the hypotheses) indicate whether each feature participants listed was best characterized as singular or plural, as well as how specific/concrete vs. generic/abstract it was[2]. We performed a partial correlation for each subject to determine the relationship between causal connections and disruption to identity while controlling for level of abstraction and plurality. The results revealed that when controlling for these variables the Spearman correlation coefficients were significantly positive for both approaches to causal centrality for both experiments (see Table S4).

[2] For coding, a scale of 1 to 5 was used where 1 meant that the listed feature was "not at all specific and could apply to anyone/is not a tangible feature" and 5 meant that the listed feature was "extremely specific to the participant/is a concrete feature."

Table S4

Summary of Experiment 2 Partial Correlation Results

Condition	Causal Connections	Causal Depth
Self	mean r_s = .22, t(91) = 5.39, p < .001, 95% CI = [.14, .30]	mean r_s = .24, t(91) = 6.04, p < .001, 95% CI = [.16, .32]
Close-other	mean r_s = .19, t(96) = 4.68, p < .001, 95% CI = [.11, .28]	mean r_s = .17, t(96) = 3.96, p < .001, 95% CI = [.08, .25]

Note. T-tests are one-sample t-tests of the mean Spearman rho (with Fisher transformation) against 0.

Plausibility Analysis – Experiment 3

Because different causal structures may also differ in how natural they appear to be (Ahn, 1999), we also examined the plausibility of the two different types of causal structures. The common cause and common effect vignettes were rated as equally plausible ($M_{commoncause}$ = 74.2, $M_{commoneffect}$ = 72.5, $t(11) = .47$, $p > .250$). So, the observed difference in selections between the two conditions cannot be explained by a difference in the believability of the two causal structures.

References

Ahn, W. (1999). Effect of causal structure on category construction. *Memory & Cognition*, *27*, 1008-1023.

Kim, N. S., & Park, E. Y. (2009). ConceptBuilder: An open-source software tool for measuring, depicting, and quantifying causal models. *Behavior Research Methods*, *41*, 128-136.

Sloman, S., Love, B., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22,* 189-228.

Appendix S2

Experiment 1

Wording of Survey Questions

Disruption to identity question: *We would now like to understand how a change in each of the features below would change your identity. That is, for each of the features below, imagine that you are completely different on that dimension (e.g., for height, if you are tall, imagine that something changed so that you are now short). Do you think that you would still be the same person you are now, or would you be a different person? Please indicate your answer with each of the sliders below where 0 means, "I would be the exact same person" and 100 means, "I would be a completely different person."*

Expected change question: *We would now like to understand how much you think each of the features below may change in the future. Please indicate how much you think each feature will change in the* ***next 5 years*** *with each of the sliders below where 0 means, "Will not change at all" and 100 means, "Will change completely."*

Details of ANOVA

An ANOVA[1] with condition (self vs. close-other vs. generic-other) as a between-subjects factor and features as repeated measures found an effect of feature, $F(17, 2113)=156.34$, $p<.001$, on causal connections. There was also a significant condition × feature interaction, $F(17, 2113) = 1.62$, $p = .049$, suggesting that differences in centrality across the features varied by condition.

[1] Results reported here were Huynh-Feldt corrected when sphericity could not be assumed.

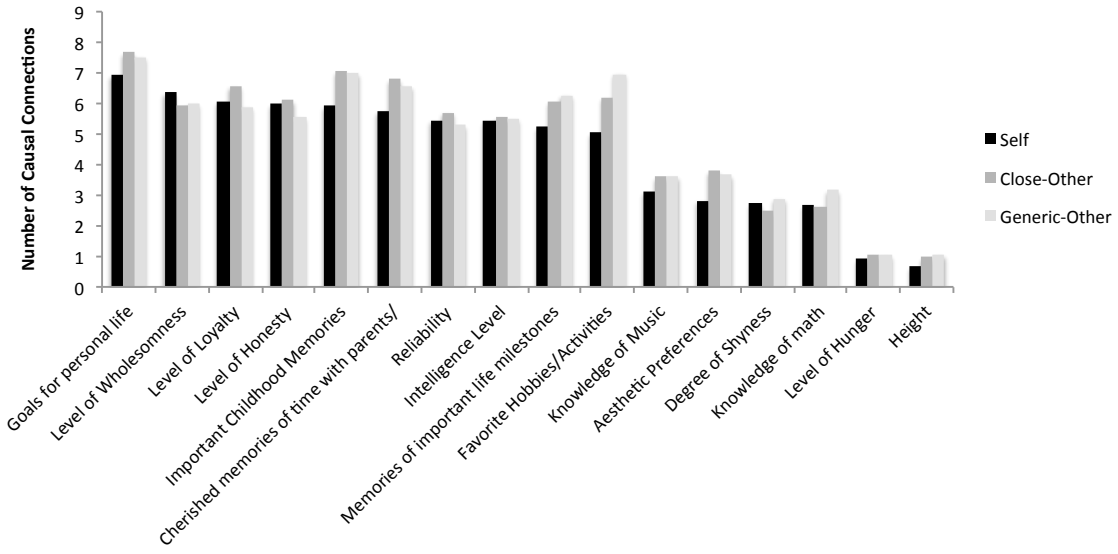

Figure SU1. Results of Experiment 1. Number of causal connections for each feature by condition.

Table SU1

Features with Most Causal Connections in Experiment 1

Condition	Features with most causal connections (# links)	% participants who list feature as most connected	% participants who list feature in top three most connected
Self	Goals for personal life (6.9)	24%	47%
	Level Wholesomeness (6.4)	20%	47%
	Level of Loyalty (6.1)	6%	42%
Close-other	Goals for personal life (7.7)	30%	55%
	Important Childhood Memories (7.1)	26%	45%
	Cherished memories w family (6.8)	10%	37%
Generic-other	Goals for personal life (7.5)	28%	47%
	Important childhood memories (7.0)	19%	48%
	Favorite hobbies/activities (7.0)	20%	46%

Correlations Between Causal Centrality and Disruption to Identity with All Participants

We performed both the aggregate and individual-level Spearman correlational analysis including all participants (i.e. even those who failed the attention checks)[2]. The results are similar to those presented in the main manuscript and are summarized in Table SU2.

Table SU2

Correlations Between a Feature's Number of Causal Connections and Ratings of How Disruptive Change in that Feature Would Be (Experiment 1, Including All Participants who Failed Attention Checks)

	Causal Connections Approach	
Condition	Aggregate Spearman Correlation	Individual Spearman Correlations
Self	r_s = .60, p =.014	mean r_s = .34, $t(79)$ = 7.20, p < .001, 95% CI = [.25, .43]
Close-other	r_s = .62, p = .013	mean r_s = .37, $t(80)$ = 8.88, p < .001, 95% CI = [.28, .45]
Generic-other	r_s = .42, p = .110	mean r_s = .29, $t(81)$ = 6.19, p < .001, 95% CI = [.20, .38]

Note. T-tests in Individual Correlations column are one-sample t-tests of the mean Spearman rho (with Fisher transformation) against 0.

Correlations Between Causal Centrality and Disruption to Identity with only High Importance Features

Causal Connections. Changes to features with more causal connections on average were not rated as more disruptive to identity, in all three conditions (self: r_s = .17, p = .588; close-other: r_s = .18, p = .573; generic-other: r_s = -.17, p = .604). However, the majority of participants in all

[2] All participants who provided usable data are included in this analysis. Correlations for two participants who gave all the same answers could not be calculated so they are excluded from this analysis. A technical error prevented five participants from viewing any features, so they are also excluded from this analysis.

conditions had a positive correlation between features' causal connections and rated disruption (65% of participants in each condition). On average, the individual-level Spearman correlations were significantly positive in the self and close-other conditions and marginally so in the generic-other condition ($\bar{r}_{Self}$=.15, p = .003, 95% CI = [.05, .25]; $\bar{r}_{Close-other}$ = .17, p < .001, 95% CI [.09, .25]; $\bar{r}_{Generic-other}$ = .10, p = .056, 95% CI = [.00, .20]).

Causal Depth. Changes to features with higher causal depth scores (see SOM-R for details on how this was calculated) on average were not rated more disruptive to identity in all three conditions (self: r_s = -.17, p = .588; close-other: r_s = .22, p = .485; generic-other: r_s = -.23, p = .471). The majority of participants in all conditions had a positive correlation between features' causal depth and rated disruption (55%, 63%, 59% of participants in the self, close-other, and generic-other conditions). However, the average individual-level Spearman correlations was significantly positive in only the close-other condition ($\bar{r}_{Self}$ = .08, p = .127, 95% CI = [-.02, .19]; $\bar{r}_{Close-other}$ = .12, p = .015, 95% CI [.03, .22]; $\bar{r}_{Generic-other}$ = .07, p = .144, 95% CI = [-.03, .17]).

Experiment 2

Correlations Between Causal Centrality and Disruption to Identity with All Participants

We performed the individual level Spearman correlational analysis including all participants (i.e. even those who failed the attention checks)[3]. The results are similar to those presented in the main manuscript and are summarized in Table SU3.

[3] All participants who provided usable data are included in this analysis. Correlations for eight participants who gave all the same answers could not be calculated so they are excluded from this analysis.

Table SU3

Correlations Between a Feature's Number of Causal Connections and Ratings of How Disruptive Change in that Feature Would Be (Experiment 2, Including All Participants who Failed Attention Checks)

Condition	**Causal Connections**
Self	mean r_s = .23, t(93) = 6.19, p < .001, 95% CI = [.15, .30]
Close-other	mean r_s = .22, t(99)= 6.26, p < .001, 95% CI = [.14, .30]

Note. T-tests in Individual Correlations column are one-sample t-tests of the mean Spearman rho (with Fisher transformation) against 0.

Experiment 3

Table SU4

*Vignettes used in Experiment 3 and Supplemental Experiment 4**

Common Cause Vignettes	
Version A	Version B
Jack has a few salient characteristics. Among them are that he is very shy, he likes solitary activities, he has an awkward demeanor, and he has very strong memories from his childhood.	Jack has a few salient characteristics. Among them are that he is very shy, he likes solitary activities, he has an awkward demeanor, and he has very strong memories from his childhood.
Jack didn't have many friends as a child so many of his memories from childhood are of him playing alone. As a result of these memories and experiences, he has always been shy, he developed an awkward demeanor from not interacting much with his peers, and his favorite activities are generally solitary ones like building model airplanes. He has always thought his memories of his childhood experiences caused his shyness, awkward demeanor, and his preference for solitary activities.	Jack he has always been shy. As a result of his shyness, he didn't have many friends as a child so many of his memories from childhood are of him playing alone, he developed an awkward demeanor since he was too shy to interact much with his peers, and his favorite activities are generally solitary ones like building model airplanes. He has always thought his shyness caused these memories of his childhood experiences, his awkward demeanor, and his preference for solitary activities.
Mary has a few salient characteristics. Among them are that she is very ambitious, her professional goal is to become a neurosurgeon, she enjoys doing volunteer work, and many of her important memories are of her various academic accomplishments.	Mary has a few salient characteristics. Among them are that she is very ambitious, her professional goal is to become a neurosurgeon, she enjoys doing volunteer work, and many of her important memories are of her various academic accomplishments.
As long as anyone can remember, Mary wanted to be a neurosurgeon. As a result of this goal, Mary became a very good student so many of her important memories are of her various academic accomplishments. Her goal also led her to develop a very high level of ambition and an affinity for volunteering in hospitals to prepare for her future career. She has always thought that her desire to become a neurosurgeon caused these important memories of her accomplishments, her enjoyment of volunteer work, and her ambitious personality.	As long as anyone can remember, Mary has been very ambitious. As a result of her ambition, Mary was a very good student and many of her important memories are of her various academic accomplishments. Her ambition also led her to develop the professional goal of being a neurosurgeon and to enjoy volunteering at hospitals as she really wanted to make a difference. She has always thought that her ambitious personality caused her desire to be a neurosurgeon, her enjoyment of volunteer work, and these memories of her academic experiences.
Henry has a few salient characteristics. Among them are that he is extremely honest, he has many memories of the lessons his parents taught him as a child, he is naive, and he has a lot of close friends.	Henry has a few salient characteristics. Among them are that he is extremely honest, he has many memories of the lessons his parents taught him as a child, he is naive, and he has a lot of close friends.
When he was a child, Henry's parents emphasized the importance of honesty and he has many memories of them praising him for being honest. As a result of these lessons and memories, Henry has always had many close friends, he is a bit naive because he often thinks everyone is honest, and he developed into an extremely honest person; no one can remember an instance of Henry being dishonest. Henry has always thought that his memories of his parents emphasis on and praise of honesty caused his close friendships, his naiveté, and his honesty.	Henry has always been an extremely honest person; no one can remember an instance of Henry being dishonest. As a result of his honesty, Henry is a bit naive because he often thinks everyone is as honest as he is. In addition, his honesty caused him to have many close friendships and many memories of his parents praising him for being honest and emphasizing the importance of honesty. Henry has always thought that his honesty caused his parents to emphasize these life lessons, his naiveté, and his many close friendships.

Version A	Version B
Jane has a few salient characteristics. Among them are that she is very intelligent, she has many cherished childhood memories of time spent with her parents, she is a very curious person, and her favorite activities generally involve learning new things.	Jane has a few salient characteristics. Among them are that she is very intelligent, she has many cherished childhood memories of time spent with her parents, she is a very curious person, and her favorite activities generally involve learning new things.
As long as anyone can remember, Jane has always been very intelligent. Her intelligence caused her parents to teach her about everything from science to music and she has many cherished memories of her parents teaching her all sorts of things about the world. Her intelligence also cause her develop into a very curious person and a love for activities that involve learning—one of her favorite activities is going to museums. Jane has always thought that her intelligence caused her memories of her parents teaching her about the world, her curiosity, and her love of learning-related activities.	As long as Jane can remember, her parents always wanted to teach her about the world. From an early age they taught her about everything from science to music. As a result of these memories, Jane developed a high level of curiosity, a high level of intelligence and a love of activities that involve opportunities to learn new things—one of her favorite activities is going to museums. Jane has always thought that her memories of her parents teaching her about the world caused her intelligence, her curiosity, and her love of learning.
Anne has a few salient characteristics. Among them are that she is very reliable, she has many memories of her parents giving her a lot of responsibility while she was growing up, she is very organized, and she enjoys being involved in group activities.	Anne has a few salient characteristics. Among them are that she is very reliable, she has many memories of her parents giving her a lot of responsibility while she was growing up, she is very organized, and she enjoys being involved in group activities.
As long as anyone can remember, Anne has always been very reliable. As a result, she developed into a very organized person and enjoys organizing group activities-- for example, she was the class president in middle school and captain of her basketball team. In addition, because she was so reliable, Anne's parents gave her lots of responsibilities when she was growing up--she was often asked to take care of her siblings and she often held part time jobs while she was in school. She has always thought her reliability caused her to be organized, her participation in group activities, and her parents to give her so many responsibilities.	As long as she can remember, Anne's parents always gave her a lot of responsibility. She was often asked to take care of her siblings and she often held part time jobs when she was in school. As a result, she developed into a very organized person and a very reliable person. In addition, these experiences of having many responsibilities led to her enjoyment of organizing group activities-- for example, she was the class president in middle school and captain of her basketball team. She has always thought her memories of the responsibilities her parents gave her growing up caused her to be organized, her high level of responsibility and her participation in group activities.
Robert has a few salient characteristics. Among them are that he is very loyal, he has many memories fond memories of childhood, his favorite activity is playing team sports, and he is very dependable.	Robert has a few salient characteristics. Among them are that he is very loyal, he has many memories fond memories of childhood, his favorite activity is playing team sports, and he is very dependable.
As long as anyone can remember, Robert has always been very loyal. He is extremely loyal to those who are close to him. As a result of his loyalty, he has many fond memories of his childhood friends who have remained his best friends for his entire life. His sense of loyalty also caused him to become a very dependable personality as well as enjoy playing team sports as he really enjoys being a good teammate. Robert has always thought his loyalty caused his fond memories of his childhood friends, his love of team sports, and his dependability.	Robert has many fond memories of his childhood spent with his childhood friends who have remained his best friends for his entire life. As a result of these friendships, he became a a very dependable person and a very loyal person; he is extremely loyal to those who are close to him. As he played many sports with his childhood friends, these memories also caused his love of playing team sports as he really enjoys being a good teammate. Robert has always thought his memories of his childhood friendships caused him develop his dependable personality, loyalty, and his love of team sports.

Common Effect Vignettes	
Version C	Version D
Jack has a few salient characteristics. Among them are that he is very shy, he likes solitary activities, he has an awkward demeanor, and he has very strong memories from his childhood.	Jack has a few salient characteristics. Among them are that he is very shy, he likes solitary activities, he has an awkward demeanor, and he has very strong memories from his childhood.
Jack didn't have many friends as a child so many of his memories from childhood are of him playing alone. These memories and experiences were a result of the fact that he has always been shy, has an awkward demeanor, and his favorite activities are generally solitary activities like building model airplanes so he never really was inclined or able to interact with his peers much. He has always thought his memories of his childhood experiences were caused by his shyness, awkward demeanor, and his preference for solitary activities.	Jack is very shy. Jack has always had an awkward demeanor and preference for solitary activities like building model airplanes so he was never really inclined or able to interact much with his peers which led him to become quiet shy. In addition, he didn't have many friends as a child so many of his memories from childhood are of him playing alone which has also caused him to be quite shy. He has always thought his shyness was caused by these memories of his childhood experiences, his awkward demeanor, and his preference for solitary activities.
Mary has a few salient characteristics. Among them are that she is very ambitious, her professional goal is to become a neurosurgeon, she enjoys doing volunteer work, and many of her important memories are of her various academic accomplishments.	Mary has a few salient characteristics. Among them are that she is very ambitious, her professional goal is to become a neurosurgeon, she enjoys doing volunteer work, and many of her important memories are of her various academic accomplishments.
Mary's professional goal is to become a neurosurgeon. This goal is caused by the fact that Mary has always been very good student and has many important memories are of her various academic accomplishments. Her professional goal is also a result of the fact that she has always had a very high level of ambition and has enjoyed doing volunteer work in hospitals. She has always thought that her desire to become a neurosurgeon was caused these important memories of her accomplishments, her enjoyment of volunteer work, and her ambitious personality.	Mary is very ambitious. Mary's ambition is caused by the fact that she was a very good student so many of her important memories are of her various academic accomplishments and people encouraging her to do things that challenged her. In addition, her ambitious personality was caused by her professional goal of being a neurosurgeon and her enjoyment of volunteering at hospitals. She has always thought that her ambitious personality was caused by her desire to be a neurosurgeon, her enjoyment of volunteer work, and these memories of her academic experiences.
Henry has a few salient characteristics. Among them are that he is extremely honest, he has many memories of the lessons his parents taught him as a child, he is naive, and he has a lot of close friends.	Henry has a few salient characteristics. Among them are that he is extremely honest, he has many memories of the lessons his parents taught him as a child, he is naive, and he has a lot of close friends.
Henry has always been extremely honest; no one can remember an instance of his being dishonest. He also has many close friendships and can be a bit naive as he often thinks that everyone is as honest as he is. As a result of this, he has many memories of his parents praising him for being honest and emphasizing the importance of honesty and making friends with people who are honest like him. Henry has always thought that his memories of his parents emphasis on and praise of honesty were caused by his close friendships, his naiveté, and his honesty.	Henry is an extremely honest person; no one can remember an instance of Henry being dishonest. His honesty was caused by many memories of his parents praising him for being honest and emphasizing the importance of honesty, and by many close friendships that have taught him the value of honesty. In addition, his honesty was also caused by his naiveté which leads him to sometimes believe that everyone is honest so he should be as well. Henry has always thought that his honesty was caused by his memories of the life lessons his parents taught him, his naiveté, and his many close friendships.

Version C	Version D
Jane has a few salient characteristics. Among them are that she is very intelligent, she has many cherished childhood memories of time spent with her parents, she is a very curious person, and her favorite activities generally involve learning new things.	Jane has a few salient characteristics. Among them are that she is very intelligent, she has many cherished childhood memories of time spent with her parents, she is a very curious person, and her favorite activities generally involve learning new things.
Jane is very intelligent. Her intelligence was caused by her early experiences with her parents who taught her about everything from science to music. She has many cherished memories of her parents teaching her all sorts of things about the world. Her intelligence is also caused by her very curious personality and her love for activities that involve learning—one of her favorite activities is going to museums. Both of these characteristics have led Jane to learn a great amount and develop her intellect. Jane has always thought that her intelligence was caused by her memories of her parents teaching her about the world, her curiosity, and her love of learning-related activities.	Jane is a very intelligent and curious person. Her favorite activities generally involve opportunities to learn new things—one of her favorite activities is going to museums. Because of her intelligence, curiosity, and love of learning-related activities, from an early age, her parents taught her all about the world. Jane has many cherished memories of her parents teaching her about everything from science to music. Jane has always thought that her memories of her parents teaching her about the world were caused by her intelligence, her curiosity, and her love of learning.
Anne has a few salient characteristics. Among them are that she is very reliable, she has many memories of her parents giving her a lot of responsibility while she was growing up, she is very organized, and she enjoys being involved in group activities.	Anne has a few salient characteristics. Among them are that she is very reliable, she has many memories of her parents giving her a lot of responsibility while she was growing up, she is very organized, and she enjoys being involved in group activities.
Anne is very reliable. Her reliability was caused by the fact that she is a very organized person and she has been involved in and enjoys organizing group activities that have taught her how to be reliable-- for example, she was the class president in middle school and captain of her basketball team. In addition, Anne's parents gave her lots of responsibilities when she was growing up--she was often asked to take care of her siblings and she often held part time jobs while she was in school. Because she was responsible for so many things, she developed into a very reliable person. She has always thought her reliability was caused by her high level of organization, her participation in group activities, and the many responsibilities her parents gave her growing up.	Anne's parents always gave her a lot of responsibility growing up. She was often asked to take care of her siblings and she often held part time jobs when she was in school. The reason they gave her so much responsibility is because she was always a very organized person and a very reliable person. In addition, she really enjoyed and had a lot of experience organizing group activities-- for example, she was the class president in middle school and captain of her basketball team. Thus, her parents were able to observe that she could handle a lot of responsibility. She has always thought her memories of the responsibilities her parents gave her growing up were caused by the fact that she is very organized, her high level of responsibility and her participation in group activities.
Robert has a few salient characteristics. Among them are that he is very loyal, he has many memories fond memories of childhood, his favorite activity is playing team sports, and he is very dependable.	Robert has a few salient characteristics. Among them are that he is very loyal, he has many memories fond memories of childhood, his favorite activity is playing team sports, and he is very dependable.
Robert has many fond memories of his childhood spent with his childhood friends who have remained his best friends for his entire life. These memories are a result of the fact that Robert is a very dependable and loyal person. These traits allowed him to have close childhood friends. These memories were also caused by the fact that he loved playing team sports. He met many of his childhood friends playing team sports and was always an excellent teammate. Robert has always thought his fond memories of his childhood friendships were caused	Robert is extremely loyal to those who are close to him. His high level of loyalty is caused by both his dependability and his many fond memories of his childhood friends who have remained his best friends for his entire life. His sense of loyalty was also caused by his love of playing team sports. Playing team sports taught him the value of being a good, loyal teammate. Robert has always thought his loyalty was caused by his fond memories of his childhood friends, his love of team sports, and his dependability.

by his dependability, his loyalty, and his love of team sports.

*Common cause vignettes were modified to include only three features in Supplemental Experiment 4.

Results (No Trials or Participants Excluded)

When no trials or participants were excluded, the results were similar to the results reported in the main manuscript. The dependent measure was the average percentage of trials where participants selected the person who was missing the effect feature. For the common cause trials, we predicted that participants should pick the person missing the effect feature as the same person rather than pick the person missing the cause feature. This is what we found. Participants were significantly more likely to select the person missing the effect feature (M = 67%, SD = 31%, 95% CI = [.59, .75]), $t(59) = 4.21$, $p < .001$.

For the common effect trials, however, the two approaches to causal centrality yield different predictions. Our results were more consistent with the number of causal connections approach suggesting that a missing effect should disrupt identity more than a missing cause, because the effect feature participates in more causal relationships than the cause features do. That is, participants were *less* likely to select the person missing the effect feature—the feature that had more connections but was less deep in the causal chain (M = 38%, SD = 28%, 95% CI = [.30, .45]), $t(59) = 3.33$, $p = .002$. The average percentage of missing effect selections was significantly different between the common cause and common effect conditions, (Ms = 67% vs. 38%), $t(59) = 5.18$, $p < .001$.

Supplemental Experiment 4

In Experiment 4, we manipulated the causal centrality of features in a set of vignettes to test whether making a feature more causally central impacts how defining that feature is for identity. These vignettes described how three features of a character's identity related to one another in a common cause structure.

Method

Fifty-nine participants were recruited via Amazon Mechanical Turk. One participant was removed before analysis for answering all comprehension check question incorrectly, leaving 58 participants for analysis. Similar results were found including all participants.

Six vignettes described the causal relationship between three salient features of a person, with one of the features presented as the cause of the other two, which were effects (see Table SU4). Since, in the vignettes, the cause feature has more causal connections than the effect feature, it is relatively more causally central. Each vignette had two versions (A and B), manipulating the causal centrality of two focal features by switching which of the two features was a cause and which was an effect. For example, in one version (A) of the vignette, Jack's lonely memories caused both his shyness and solitary preferences. In the other version (B), Jack's shyness caused both his memories of being a lonely child and his solitary preferences (see Table SU4). So, the exact same features were counterbalanced to play both the cause and effect role, to control for any idiosyncratic influences of specific features.

Each participant read six vignettes, each of which included a diagram summarizing the information (see Table SU4). Participants were randomly assigned to read one of the two versions of each vignette (e.g., Version A or B). A comprehension check was instituted on a separate screen after each vignette to ensure participants understood the causal structure.

Participants then read about two people, one missing the cause feature and one missing the focal effect feature. To measure which feature was seen as more defining to identity, we then asked participants which of the two people was most likely to be the character in the vignette. Participants then reported how plausible they found the vignette on a scale of 0 (“not at all plausible) to 100 (“extremely plausible”). To ensure participants had carefully made their selections, on a separate screen they again reported which person was more likely to be the character in the vignette. Finally, to promote close reading, participants answered a simple multiple-choice question about the details of the story.

Table SU4.

Example Story Used in Experiment 4

Version A

Jack has a few salient characteristics. Among them are that he is very shy, he likes solitary activities, and he has very strong memories from his childhood.

Jack didn't have many friends as a child so many of his memories from childhood are of him playing alone. As a result of these memories and experiences, he has always been shy and his favorite activities are generally solitary ones like building model airplanes. He has always thought his memories of his childhood experiences caused both his shyness and his preference for solitary activities.

This information can be summarized as follows:

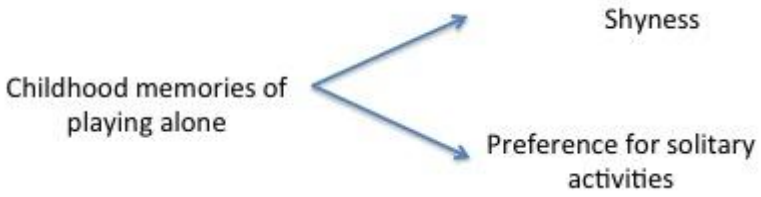

Version B

Jack has a few salient characteristics. Among them are that he is very shy, he likes solitary activities, and he has very strong memories from his childhood.

Jack has always been shy. As a result of his shyness, he didn't have many friends as a child so many of his memories from childhood are of him playing alone and his favorite activities are generally solitary ones like building model airplanes. He has always thought shyness caused both these memories of his childhood experiences and his preference for solitary activities.

This information can be summarized as follows:

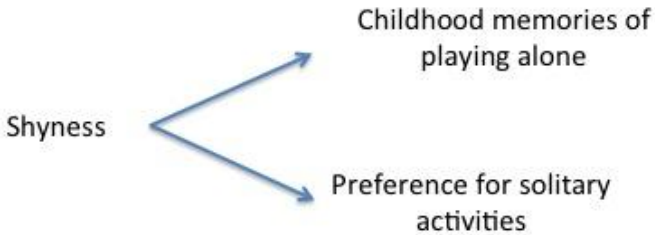

Results

We excluded trials where participants misreported the causal relationship of the features in the comprehension check (9% of trials) or provided inconsistent answers of which person was the character in the vignette when asked a second time (3% of trials). Similar results were found when no trials were excluded.

The dependent measure was the average percentage of trials in which participants selected the person who was missing the effect feature. We predicted that participants will believe that the person missing the causally peripheral effect feature is the same person, rather than the person missing the causally central cause feature. This is what we found. Participants were significantly more likely to select the person missing the effect feature (M = 68%, SD = 23%, 95% CI = [.62, .74]), $t(57) = 5.96$, $p < .001$, $d = .78$, replicating the prior findings using an experimental manipulation of causal centrality. Results were similar when no trials were excluded, (M = 66%, SD = 23%, 95% CI = [.60, .72]), $t(58) = 5.6$, $p < .001$. When a feature was made more causally central, changes in that feature were perceived as more inconsistent with continuity of identity.

As in Experiment 3, we also looked at the effects of vignette plausibility. Representations of concepts, in general, are influenced by our prior knowledge (Murphy, 2002; Murphy & Medin, 1985). This suggests that prior intuitions about what causal relationships are likely to occur among features of identity may moderate identity judgments. The higher the perceived plausibility, the more likely that participants believed that the effect feature was actually causally peripheral and the more likely that they selected the person missing the effect feature.

We found the predicted moderation by plausibility. Story plausibility was correlated with the proportion of selections of the person missing the causally peripheral effect feature (r = .65, p = .023, 95% CI = [.12, .89]). The average proportion of selections of the person missing the effect

feature was significantly higher among the six most plausible vignettes (M = 56.9%, SD = 15.8%) than the six least plausible vignettes (M = 80.2%, SD = 9.1%), $t(10) = 2.55$, $p = .029$, 95% CI = [.03, .44]. These results suggests that participants' use of the causal information from the vignettes was moderated by how it fit with their beliefs about which causal relationships are relatively more likely to occur between features of identity.

Pilot Experiment

The pilot experiment aimed to examine people's beliefs about the causal relations between features of identity and to test whether more causally central features are more essential to the self-concept. Similar to Experiment 1 in the main text, participants completed two tasks: one that elicited beliefs about the causal centrality of features of identity and one that measured beliefs about how disruptive a change to these features would be to identity.

Method

Ninety-two University of Chicago students completed two tasks in randomized order, a concept map task and a survey about personal identity. Twelve participants were excluded due to computer program failures, either to record data or to display randomized features, yielding 80 cases.

Each participant drew a computerized map of the causal links between the 16 features of their identity used in Experiment 1. Participants used ConceptBuilder software to report beliefs about causal relations (Kim & Park, 2009), first in an unrelated practice task and then to draw the causal map of personal identity using the 16 features. The features were initially presented on the screen in random order. Participants could move the features around and draw unidirectional or bidirectional arrows between them to represent cause-effect relationships, as they saw fit (see

Figure SU2). For each link specified, participants also rated the strength of the causal relationship (1=weak, 2=moderate, 3=strong).

In the survey, participants rated each feature on how much (i) a change in the feature would disrupt their identity, and (ii) they expected the feature to change in the next five years.

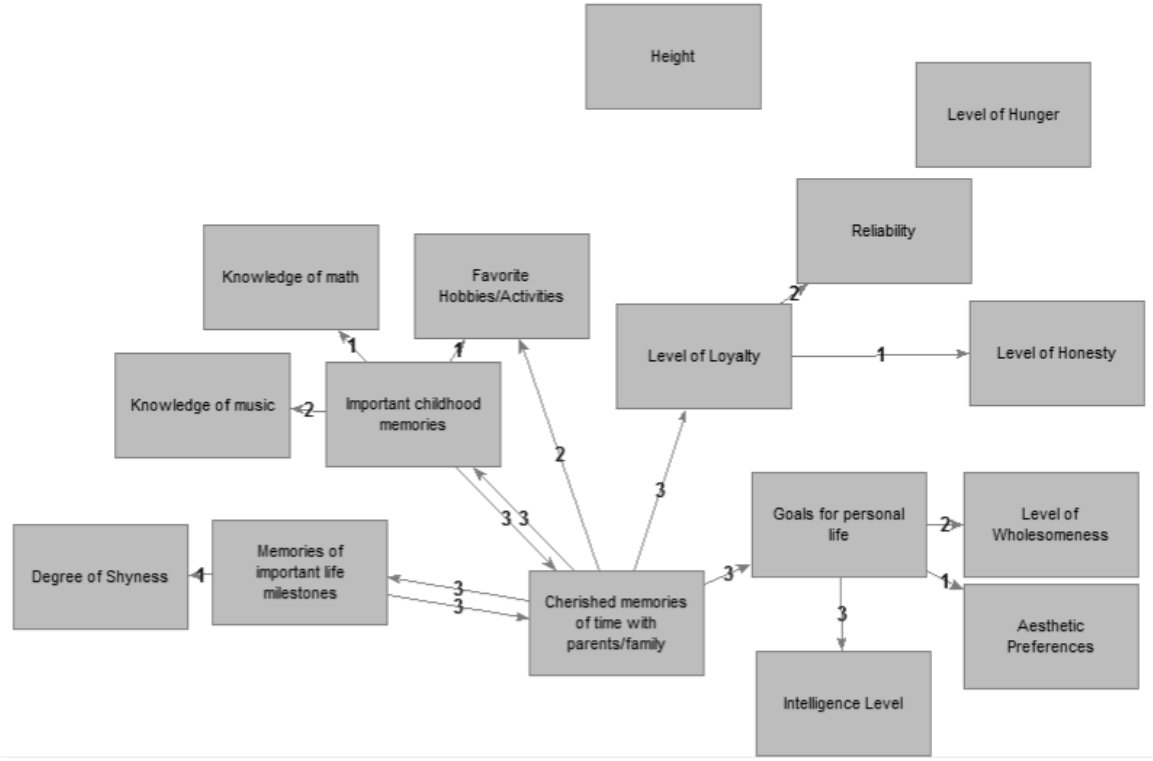

Figure SU2. Example of a self-concept map. Each box contains one of the 16 features of personal identity. The arrows represent causal relationships between features. The arrow starts at the cause feature and points to the effect feature. The numbers that are on each arrow indicate the strength of the causal relationship (1 = weak, 2 = moderate, 3 = strong).

Pre-test Comparing Concept Map and Listing Causal Relationships Tasks

In a separate pre-test, we confirmed that the "listing causal relationships" task used in Experiments 1 and 2 of the main text and the concept map task yielded similar causal centrality scores. Thirty subjects performed both tasks with the same 16 features of identity used in Experiment 1 and the pilot experiment. For each subject, we calculated the Spearman correlation between the causal centrality of the 16 features generated by the two tasks. The average Spearman correlation was significantly greater than 0 for both the causal connections approach (M = .67, SD = .45, 95% CI = [.51, .84]), $t(29) = 8.27$, $p < .001$ and the causal depth approach (M = .44, SD = .38, 95% CI = [.30, .59]), $t(29) = 6.36$, $p < .001$. The Spearman correlation was positive for the vast majority of subjects—29 out of 30 for the causal connections approach, and 26 out of 30 for the causal depth approach. So, the two tasks appear to measure causal centrality in a similar way.

Results

On average, participants drew 20.0 causal links between the features. We performed our analyses with both the causal connections and causal depth approaches to causal centrality.

Causal Connections. Consistent with our hypothesis, changes in features with more causal connections were rated as more disruptive to identity. We found a significant overall Spearman correlation between the average causal connections and rated disruption to identity (r_s = .79, $p <$ .001). This positive relationship between causal connections and disruption to identity was observed at the individual level for 80% of participants. The mean individual-level correlation between feature centrality and disruption to identity[4] was significantly positive ($\overline{r_s}$ = .33, 95% CI = [.24, .42]), $t(79) = 7.43$, $p < .001$.

[4]All correlations reported are Spearman's rho. Fisher transformations were performed prior to t-tests.

Similar results were found when analyzing only the 12 high-importance features. Changes to features with more causal connections on average were rated as moderately more disruptive to identity, r_s = .56, p = .063. The majority of participants in all conditions had a positive correlation between features' causal connections and rated disruption (67% of participants). The mean individual-level correlation between feature centrality and disruption was significantly positive, $\overline{r_s}$ = .18, 95% CI = [.10, .27]), $t(78) = 4.26$, $p < .001$. One participant gave the same answer to the disruption to identity question for all high importance features and had to be excluded from this analysis.

These results suggest that causal connections consistently influence identity judgments. The more causally central a feature was, the more strongly participants believed that a change to that feature would be disruptive to their identity.

Causal Depth. We found evidence that this measure of causal centrality influenced identity judgments. The Spearman rank correlation between causal depth and disruption to identity ratings was significant, r_s = .65, p = .008. The mean individual-level correlation between feature centrality and disruption to identity was significantly positive (M = .23, 95% CI = [.14, .31]), $t(79) = 5.27$, $p < .001$. This positive relationship between causal depth and disruption to identity was observed for 78% of subjects.

When analyzing only the 12 high-importance features, changes to features with higher causal depth scores (see SOM-R for details on how this score was calculated) on average were not rated as more disruptive to identity, r_s = .25, $p > .250$. However, the majority of participants in all conditions had a positive correlation between features' causal depth and rated disruption (61% of participants). The mean individual-level correlation between feature centrality and disruption was significantly positive, $\overline{r_s}$ = .12, 95% CI = [.02, .22]), $t(78) = 2.44$, p = .017. One participant

gave the same answer to the disruption to identity question for all high importance features and had to be excluded from this analysis.

Relationship between Causal Centrality and Expected Change

We found that the two measures of causal centrality, casual connections and causal depth, did not relate to the ratings of expected change in each feature ($\bar{r}_{Connections}$=.07, p = .061, 95% CI = [0.0, .15]; $\bar{r}_{Depth}$ = -.03, p = .437, 95% CI = [-.11, .05]). In fact, some of the most central features (childhood memories) and least central features (height) were expected to remain the most stable. This suggests that the causal centrality of a feature, although strongly related to the importance of the feature's stability for identity, is distinct from the anticipated likelihood of change in that feature.

References

Kim, N. S., & Park, E. Y. (2009). ConceptBuilder: An open-source software tool for measuring, depicting, and quantifying causal models. *Behavior Research Methods*, *41*, 128-136.

Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

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

WE ARE WHAT WE THINK: REPRESENTATIONS OF THE SELF-CONCEPT AND IDENTITY-BASED CHOICE

STEPHANIE Y. CHEN

OLEG URMINSKY

Working Paper

AUTHOR NOTE:

Stephanie Y. Chen (stephanie.chen@chicagobooth.edu) is a postdoctoral fellow, and Oleg Urminsky (oleg.urminsky@chicagobooth.edu) is a professor of marketing at the University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637.

Abstract

We propose a novel approach to identity-based choice that focuses on consumers' representations of the cause-effect relationships that exist among features of their self-concepts. More specifically, we propose that people who believe that a specific aspect of identity, such as a social category, is *causally central* (linked to many other features of the self-concept) are more likely to engage in behaviors consistent with that aspect than those who believe that the same aspect is *causally peripheral* (linked to fewer other features). Across five studies, we provide evidence for our approach to identity-based choice. We demonstrate that among consumers who belong to the same social category, those who believe that the associated identity is more causally central are more likely to engage in behaviors consistent with the social category. Additionally, we show that even aspects of identity that are not necessarily associated with well-defined social categories (e.g., honesty) are also related to behaviors consistent with that aspect when causally central.

Keywords: causal reasoning, identity, identity-based choice, self-concept

"We do what we do, because of who we are. If we did otherwise, we would not be ourselves." - Neil Gaiman, The Kindly Ones

Personal identity, the aspects of the self that people see as defining themselves, is central to the decision making process and has been implicated in a broad range of consumer behaviors (see Reed et al. 2012 for a detailed review). Psychologists and consumer researchers have long noted that the social categories we belong to provide us with norms to follow, scripts for behaviors, and ways to interpret our actions (Kleine, Kleine, and Kernan 1993; Markus and Wurf 1987; Turner 1987). Additionally, consumers choose products that are associated with their identities (Escalas and Bettman 2003), using their brand and product choices to build and express their identities (Belk 1988; Berger and Heath 2007).

Much research has examined how situational factors influence identity-consistent behaviors. This area of research has found that the degree to which a specific aspect of identity influences behavior depends on the salience of that aspect (Benjamin, Choi, and Strickland 2010; Broughs et al. 2016; Cohn et al. 2014; Forehand et al. 2002; LaBeouf, Shafir, and Bayuk 2010; Reed 2004). Further, economists have begun to incorporate identity into utility models of decision-making, under the assumption that deviations from the norms prescribed by a salient identity yield disutility (Akerlof and Kranton 2000, 2010). Thus, prior research concludes that a given identity aspect influences behavior among people who hold that identity, when that identity is salient.

Within a given context, however, people who hold the same salient identity often differ in how identity-consistent their behavior is. For example, the Super Bowl makes being a football fan very salient, but fans differ in their identity-consistent behavior (e.g.,

spending money on merchandise and tickets). Prior research has extensively studied the effects of identity salience on how people respond to identity-consistent and identity-inconsistent stimuli, and how they make identity-based choices (Reed et al. 2012). However, beyond these situational factors, prior approaches to studying identity-based behaviors provide little insight into why a given identity would exert more influence on some identity-holders than others.

We propose that individuals who maintain a given identity will differ in their identity-consistent behaviors, based on their beliefs about how that identity fits into their broader self-concept. Moving beyond previously studied situational determinants of identity-consistent behavior, we examine how identities are causally connected to other important features of people's own subjective self-concept (e.g., memories, moral qualities, and personality traits). Specifically, we draw on research from cognitive psychology on concepts, which suggests that the aspects that are most defining of a concept are those that are more *causally central* (i.e., participate in more cause-effect relationships with other features; Rehder and Hastie 2001).

We hypothesize that a given aspect will exert more influence on behaviors among consumers for whom that aspect is more causally central, compared to consumers who believe the aspect is more causally peripheral. For example, a person who believes being a football fan is causally linked to many other features of her identity (such as her preferences, values and principles, relationships with others, occupation, or childhood memories) would be more likely to act in ways consistent with being a fan than another football fan who sees her football-fan identity as causally linked to few other features of her self-concept.

Across five studies, we find—consistent with our hypothesis—that when an aspect of identity is more causally central, people engage in more identity-consistent behaviors than when it is peripheral. In Study 1, we test whether differences in the causal centrality of environmental goals explains differences in willingness to spend on environmentally-friendly products. In Studies 2, 3a, and 3b, we examine situations in which an aspect of identity that all participants share is made salient by a real-world event—the 2016 presidential election for Democrat and Republican identities (Study 2) and the Super Bowl for football fan identities (Studies 3a and 3b). In these studies, we measure the causal centrality of an aspect of identity and self-reported identity-relevant behaviors (e.g., voting for the presidential candidate nominated by their party). In Study 4, we test the effect of a causally central honesty identity on behavior in an incentive-compatible cheating task, over and above self-rated importance of honesty.

THEORETICAL DEVELOPMENT

Social Categories and Choice

Theories in psychology and economics hold that people are more likely to behave in ways that are consistent with their identities. In particular, these theories hold that people have multiple identities with potentially conflicting norms (LaBeouf, Shafir, and Bayuk 2010; Markus and Wurf 1987; Oyserman 2009). As a result, increasing the salience of an identity through priming, identity threat, or social distinctiveness (i.e., making members of that identity the numerical minority in the decision context) prioritizes the norms associated with that identity, making people with the identity more

likely to perform behaviors consistent with the social group's norms than when the identity is not salient (Broughs et al. 2016, Cohn et al. 2014; LaBeouf, Shafir, and Bayuk 2010).

Such identity-salience effects have been shown for a wide range of social categories including (but not limited to) gender (Broughs et al. 2016; Shang, Reed, and Croson 2008), race (Benjamin et al. 2010; Chen, Ng, and Rao 2005; Forehand, Despande, and Reed 2002; Oyserman et al. 2007), occupation (Cohn et al. 2014; LaBeouf, Shafir, and Bayuk 2010), and family roles (Reed 2004; LaBeouf, Shafir, and Bayuk 2010). This research has explored the situational factors that influence the impact of a single social category on behavior relevant to the norms of that category.

While situational factors powerfully influence people's tendency to display identity relevant behaviors, two people may be confronted with the same situational constraints and demonstrate widely different behaviors. Several theories suggest that people who are members of a given social category may differ in their likelihood of displaying identity-relevant behaviors. For example, Ackerlof and Kranton (2010) describe a simple economic model of work incentives that incorporates identity into the utility function. In their model, employees of the same company fall into two groups (or identities), "insiders" and "outsiders." Insiders believe that they should follow the norms of the company and gain utility when they do. So, insiders do not need much change in incentive to increase work efforts. In contrast, outsiders do not gain utility by following the norms of the company and require larger changes in incentives to increase work efforts. This model assumes that members of the same social category (employees of a company) have adopted the company identity to different degrees and that these differing

identities may explain people's behavior. However, these economic models do not provide an account of why people differ in their adoption of identities or identify who will more strongly adopt the identity.

Some researchers have attempted to better understand this link between identity and behavior. Empirical research investigating the effect of identities on behavior suggests that the importance of an identity may be a moderator the effect of identity on behavior (Kihlstrom 1992; Markus and Wurf, 1987; Reed et al. 2012). Reed (2004) found that people who rate a social identity as more important react more strongly to products geared towards that identity. Furthermore, the strength of identification with the social group has been shown to moderate the effect of salience on behavior. In particular, the people who believe a social group is more important to who they are show larger priming effects on behavior (LaBeouf, Shafir, and Bayuk 2010).

However, what exactly importance means in the context of the self-concept remains vague. For example, scales that measure identity importance ask how much the identity reflects or describes who the person is (LaBeouf, Shafir, and Bayuk 2010; Reed 2004). While these measures seem to capture useful differences in how people think about a given social identity, at least in some contexts, this minimally defined importance construct provides little theoretical insight into why or how an identity becomes important. Furthermore, by treating importance as an external attribute of an identity, this approach does not provide insight into how people who think a given identity is important differ from those who do not in their overall conceptualization of the self-concept.

Representation of the Self-Concept

We suggest that understanding for whom a given identity will exert more or less of an influence on identity-consistent behaviors requires investigating that identity in relation to the broader self-concept. In the social psychology, consumer behavior, and economics literature, an *identity* of a person generally refers to a social category that the person belongs to. However, a broad literature on people's beliefs about what defines the self instead focuses on individual-level psychological properties.

Approaches to understanding what defines the self have often focused on individual-level properties (such as memories and moral qualities) that are not necessarily associated with social categories. Philosophers have long suggested that continuity of memories allows for continuity in identity (Locke 1694/1979) and psychological studies have shown that disruption to memories disrupts identity judgments not only for one's self but also in judgments of others (Blok, Newman, and Rips 2005; Nichols and Bruno 2010). Recent research has suggested that lay perceptions of identity instead put moral qualities at the center of the self-concept (Strohminger and Nichols 2014, 2015). Strohminger and Nichols (2014) compared how changes to moral features and various other types of features (e.g., memories, personality, preferences and desires) impacted identity continuity judgments. They found that changes to the moral features of identity were most disruptive to identity judgments.

Research on psychological connectedness to the future self also suggests that a wide range of psychological traits define a person's identity (see Urminsky 2016 for a review). It has been argued that the extent to which one is psychologically connected to their future self (i.e., shares psychological features like memories, intentions, beliefs and

desires) should determine how willing one is to give up current rewards for future ones (Parfit 1984). Indeed, psychological research has found that inducing people to think that the psychological characteristics that make up their identity will change leads to less psychological connectedness to the future self and less willingness to forgo immediate rewards for delayed ones (Bartels and Urminsky 2011; 2015; Ersner-Hershfield et al. 2009). These results suggest that individual-level characteristics are an important part of how people think about the self and have significant impacts on their behavior.

To characterize the self-concept more broadly, we will explore these common attributes of personal identity that extend beyond social categories. We adopt terminology from the concepts and categories literature (e.g., Smith and Medin 1981; Tversky 1977), which uses the term *feature* to refer to properties or aspects of a concept. We use the terms *feature* and *aspect* interchangeably to refer to any property of the self-concept, including social categories as well as other properties of the self, such as memories, personality traits, and moral qualities.

We propose that how people represent their self-concepts will clarify why people who share that feature may nevertheless vary in how likely they are to display identity-consistent behaviors and what it means for a feature of identity to be important. Rather than focusing solely on the situational factors that influence identity-driven behaviors, we instead focus on internal representations of the self and how an identity fits into the broader self-concept. In particular, we propose that beliefs about the cause-effect relationships that exist between aspects of the self-concept influence the degree to which that aspect of personal identity impacts behavior.

Causally Central Aspects of Identity

Research in cognitive psychology has long emphasized the role of causal relationships in our understanding of the world (Gopnik and Wellman 1994; Keil 2006; Murphy and Medin 1985). As understanding cause-effect relationships allows organisms to intervene on the world (Sperber, Premack, and Premack 1995), it seems natural that causal relationships would play a privileged and pervasive role in human cognition. Research in developmental psychology suggests that even very young children have extensive causal knowledge, designing and performing sophisticated interventions on the world that allow them to disambiguate what the true cause of a desirable effect is (e.g., what action makes a toy turn on, Gopnik et al. 2001; Schulz and Gopnik 2004).

Causal reasoning not only allows for effective interventions, but also provides a structure to organize knowledge and shape learning. It has long been suggested that knowledge is represented as intuitive theories about the world that include causal relationships (Keil 1989; Murphy and Medin 1985). For example, our knowledge of Apple products not only includes the knowledge that the products are high quality, have great customer service, and are expensive but also theories about how these features are causally related—e.g., Apple products are expensive because they are high quality and are supported by great customer service.

These theories shape what people learn. People try to incorporate new knowledge into their existing conceptualization by generating theories about their relationship (Kaplan and Murphy 2000)—e.g., if the new iPhone is going to have a better camera, maybe that is because Apple feels that it needs to add better features if they want to keep charging so much for the iPhone. Additionally, new knowledge is easier to acquire if it is

consistent with people's theories about the world (Murphy and Allopenna 1994; Pazzani 1991).

Recent research has found that causal beliefs about aspects of the self-concept are also a critical part of how people think about the stability of the self. Inspired by a large literature in cognitive psychology that has found that the features of a concept that are more causally central (i.e., causally linked to many other features of the concept) are perceived as most defining of simple and artificial concepts (Ahn 1999; Ahn et al. 2000; Rehder and Hastie 2001; Rehder 2003; Sloman et al. 1998), Chen, Urminsky, and Bartels (2016) examined the role of causal beliefs in a much richer and complex setting, the self-concept. They had participants report the cause-effect relationships that they believed existed between features of their self-concept and calculated the causal centrality of a feature by summing the number of other features a given feature was causally linked to (either as a cause or as an effect). The results suggested that people perceive features to be defining of the self-concept to the extent that those features are causally central. That is, participants believed that changing more causally central features was more disruptive to their self-concept than changing causally peripheral features.

To illustrate how causal centrality impacts beliefs about what is most defining of the self-concept, imagine two pre-school teachers, Anne and Jane. Anne believes that it is her high-level of patience that caused her love of young children and caused her to choose her profession, preschool teacher. Jane instead believes that it is being a preschool teacher that has caused her to develop a high-level of patience and love of young children. As a result, even if Anne and Jane's profession, patience, and love of children are identical, their self-concepts will be fundamentally different. That is, patience will be

more central to Anne's self-concept than Jane's, while Jane's profession will be relatively more central. As a result, Jane would experience a change in profession as more of a disruption to her personal identity than Anne would.

We propose a novel causal centrality account of identity-consistent behavior. Our account integrates prior work on how social categories impact behavior, how people broadly construct their self-concept from multiple aspects of personal identity, and how causal relations structure the self-concept. In this view, each person's self-concept is a unique network of subjective causal associations among aspects of identity, including not only social categories, but memories, goals, moral values, preferences and personality traits. As changes to more causally central features of the self cause more disruption to the self-concept, consumers for whom a feature is more causally central will be more likely to show identity-consistent behaviors related to that feature.

Our main hypothesis is that features of identity influence behavior to the extent that they are considered causally central. Specifically, we propose that people who believe that a specific aspect of identity is causally linked to more other features of the self-concept are more likely to engage in a behavior consistent with that aspect than those who believe that the same aspect is relatively less causally linked to other features. For example, if a person believes that being a football fan is causally linked to other features of her self-concept (e.g., it has been caused by her upbringing, and has shaped her relationships with family and friends), our account predicts that she will be more likely to act in a way that is consistent with her identity as a football fan—e.g., purchase merchandise and tickets to see her favorite team play.

Another unique aspect of our approach to identity-based consumption is that any causally central aspect of the self-concept can influence choice. While the previous discussion has focused on how the causal centrality of a social category determines how much it will influence behavior, we expect that other aspects of the self-concept (e.g., goals, moral qualities) will operate in the same way. For example, if two people both believe that being honest is part of their self-concept, we predict that the person for whom honesty is more causally central will be more likely to act in honest ways. This is in contrast to much of the research on identity and choice which has generally focused on social categories.

Across five studies, we test our causal centrality approach to identity-based consumption. We provide the first evidence that the causal centrality of features of the self-concept is related to identity-based choice. We demonstrate that, among people who share a given aspect of identity, differences in causal centrality of that aspect explain differences in identity-consistent behavior. Additionally, in contrast to the general focus on social categories in the study of identity-based consumption, these studies also demonstrate that the causal centrality of aspects of identity other than social categories is an important determinant of choice.

STUDY 1: CAUSAL CENTRALITY OF ENVIRONMENTAL GOALS AND PURCHASING ENVIRONMENTALLY-FRIENDLY PRODUCTS

As an initial test of our hypothesis, we investigate whether the causal centrality of a goal or desire relates to identity-consistent behavior. We measured the causal centrality of the goal of buying environmentally-friendly products and had participants complete a

series of hypothetical purchase decisions in which they chose between two versions of a product: a more expensive environmentally-friendly version and a cheaper conventional version. We predicted that participants who perceived their desire to buy environmentally-friendly products as causally central would be willing to pay more for environmentally-friendly products and make more identity-consistent choices than those who perceived their desire as causally peripheral.

This study also highlights a unique aspect of our approach to identity-based choice by examining an aspect of identity that is not a social category, the goal to buy environmentally-friendly products. In our causal centrality account, any aspect of identity that is causally central can influence identity-relevant behaviors, regardless of whether the identity aspect is associated with a well-defined social category. This is in contrast to narrower accounts that view the effects of identity on behaviors as driven by social norms associated with and conveyed by membership in social categories.

Measuring Causal Centrality.

We used a "listing causal relationships" task, adapted from the approach used by Chen, Urminsky, and Bartels (2016), to measure the causal centrality of features of the self-concept. In this task, participants reported the causal relationships between a set of features of the self-concept. The features were participant-generated in Study 1, but in later studies we also use standardized experimenter-defined features.

Participants completed one trial for each feature in which that feature was the target (e.g., in Figure 1, "desire to buy environmentally-friendly products" is the target feature). In each trial, participants were asked to select the other features of the self-concept that they thought were caused by the target feature. Participants saw the target

feature at the top of the screen (along with the question text) and all of the other features, with check boxes, listed under it. They also had the option, at the bottom of the list, to select "none of these are caused by my: [target feature]" (see Figure 1). Participants were required to check at least one option but could check as many as they wanted.

Figure 1

EXAMPLE TRIAL OF LISTING CAUSAL RELATIONSHIPS TASK

Think about your **Desire to buy environmentally-friendly products**
Which of the other features of your personal identity listed below, if any, are caused by your Desire to buy environmentally-friendly products?

You may select as many or as few features as you see fit. In the below list, please select all the feature that you believe are caused by the above feature.

- ☐ Getting married
- ☐ Birth of my first child
- ☐ My goal is to positively influence to world
- ☐ My goal is to get a Masters Degree by 35
- ☐ I prefer to be active and healthy
- ☐ I dislike needless waste
- ☐ Loyalty is very important
- ☐ Leave the world the way you found it
- ☐ I'm smart
- ☐ I have a great sense of humor
- ☐ I'm young at heart
- ☐ I'm open-minded
- ☐ My desire to protect the environment
- ☐ None of these are caused by my: Desire to buy environmentally-friendly products

From this series of questions, we calculate the causal centrality—the number of causal relationships that a feature is involved in as either a cause or an effect—of each feature. More specifically, on the trial in which the feature was a target feature, the

number of selected features on that trial is equal to the number of causal relationship in which the target feature is the cause. The number of causal links that the feature participated in as an effect is calculated based on the number of times it was selected from the list of features in all the trials in which the feature is *not* the target. The sum of the number of links that feature was a cause for and the number of links the feature was an effect for provides us with the total number of causal links a feature participated in, our measure of causal centrality.

The listing causal relationships task used to measure causal centrality represents our unique approach to studying identity-driven behaviors by assessing an aspect's role in the broader context of the self-concept. This differs from previous approaches that tend to study a single identity, particularly a specific social category, in isolation from other aspects of identity that an individual may hold. By studying single identities in isolation, prior approaches have missed the information that is held in these causal relationships, which we propose is key to understanding why an aspect of the self-concept is important and how much it influences behavior.

Method

One hundred eleven Mechanical Turk U.S. Amazon Mechanical Turk participants completed the study. Study 1 consisted of two tasks: the listing causal relationships task and a choice task. In the listing causal relationships task, participants listed the two most important features of their identity in each of five categories that had been identified in previous research as being important to identity (memories, goals/desires, preferences, moral qualities, personality traits; Chen, Urminsky, and Bartels 2016; Strohminger and Nichols 2014). They then listed two important features of their identity that had not yet

been reported. After a practice task, participants completed the listing causal relationships task described above with the twelve features of identity that they had reported and two additional features: desire to protect the environment and desire to buy environmentally friendly products (see Figure 1).

In the choice task, participants made hypothetical choices between purchasing an environmentally-friendly and a cheaper conventional version of three products (light bulbs, shopping bags, and batteries, see Figure A1 in Appendix for details). The order of the choices and the placement of the environmentally-friendly and conventional options were randomized for each subject.

Finally, for each of the three products choices, participants reported which item they believed was better for the environment on a scale of 1 (conventional item better) to 5 (environmentally-friendly item better).

Results

Preliminary analysis. Overall, participants reported that the environmentally-friendly items were, on average, significantly better for the environment than conventional items ($M = 4.60$, chance = 2.5, $t(99) = 40.52$, $p < .001$, 95% CI = [4.49, 4.70]). In each pair of options, the environmentally-friendly item was considered significantly better for the environment than the conventional item (Ms ≥ 4.39, ts > 17.78, ps $< .001$).

Relationship between causal centrality of environmental goals and environmentally-friendly choices. There was a significant correlation between causal centrality of the "desire to purchase environmentally-friendly products" feature and the number of environmentally-friendly items selected ($r = .20$, $p < .05$). A comparison of the

STUDY 2: CAUSALLY CENTRAL POLITICAL IDENTITIES

The results of Study 1 support our hypothesis that people who believe that a non-social category aspect of identity is causally central would act in more identity-consistent ways than those who believe that the same aspect is causally peripheral. Our purpose in Study 2 was to expand this finding to salient social categories, political identities. We conducted a two-part study during the 2016 presidential election that tested whether voters for whom political party affiliation was more causally central were more likely to vote for the presidential candidate nominated by their party. All participants belonged to the social category of Democrat or Republican and the study took place at a time when these identities were quite salient, so differences in voting cannot be explained simply by social norms associated with and conveyed by membership in these social categories.

Method

Four hundred eleven Mechanical Turk U.S. Amazon Mechanical Turk participants completed the first wave of a larger study the day before the 2016 U.S. presidential election. Of the participants who completed the first wave, three hundred and fifty-five participants responded to the second wave which was launched the day after the election. Of those three hundred and fifty-five participants, one hundred and sixty-six reported being affiliated with the Democratic party and seventy-seven reported being affiliated with the Republican party. Participants who reported being affiliated with one of these two parties had a relevant political identity and were therefore included in the analyses (n = 243).

In the first wave, participants first answered questions about their demographics (gender, ethnicity, political ideology, political party, religion, income category, and education level). Then participants answered questions about the features of their political identity. There were two sources for the features: self-generated and experimenter-defined. Participants were randomly assigned to one of these two groups. The self-generated group was asked—in an open-ended question—to list eight different aspects of their political identity. The experimenter-defined group reported their position on eight major political issues (e.g., abortion, gun control, immigration, taxes, gay marriage, military spending, social programs, and marijuana legalization).

Participants performed the listing causal relationships task with 20 items: their political party and the six other demographic features they reported, the eight features of their political identity discussed above, and five additional features. These five additional items have been found to be important in previous explorations of personal identity (Chen, Urminsky, and Bartels 2016): childhood memories, personal life goals, friendships, personal values and principles, and personality. They then reported which candidate they supported and their economic and social ideology (conservative vs. liberal). We computed an ideology index by averaging the ratings of economic and social ideology.

The day after the election, participants reported whether they had voted in the election, which candidate they voted for, how satisfied they were that their party nominated Clinton or Trump, how they felt about the outcome of the election (happy vs. unhappy, and afraid vs. hopeful), and how politically involved they thought they would be in the future. We predicted that participants for whom being a Republican or

Democrat was causally central would be more likely to vote for their party's candidate than those for whom these identities were more causally peripheral.

Results

Preliminary Analyses. Participants reported significantly more links to political party when they self-generated the features of their political identity ($M = 9.92$) than among the group who evaluated experimenter-defined features ($M = 7.53$, $t(241) = 2.71$, $p < .001$, 95% CI = [0.65, 4.12]). Similar differences were observed in the total number of links to all features, between those evaluating self-generated features ($M = 156.04$) vs. experimenter-defined features ($M = 111.97$, $t(241) = 3.67$, $p < .001$, 95% CI = [20.44, 67.70]). Because of these differences, we controlled for the source of the features in the regression analyses reported below that use causal centrality to predict voting behavior. No other effects of feature source were found.

Voting Behavior. Our sample consisted of 77 Republicans and 166 Democrats. The majority of these participants, who all had a party affiliation, reported that they voted in the election (Democrats: 92%; Republicans: 95%) and that they voted for the candidate that their party nominated (Democrats: 89%; Republicans: 71%). A logistic regression predicting whether participants voted with their party based on causal centrality of political party (the number of causal links political party had to other features) revealed that, as predicted, people who believed that their political party was more causally central were more likely to vote with their party than those who saw political party as more causally peripheral ($B = .099$, Wald $\chi^2(1) = 8.67$, $p = .003$).

We also fit a logistic regression predicting whether participants voted with their party based on causal centrality of political party (the number of causal links political

party had to other features), controlling for total number of links reported (among all features), survey version (experimenter defined vs. self-generated features), ideology index, party (Democrat vs. Republican), and the interaction between ideology and party. As predicted, participants who perceived their political party as more causally central prior to the election were then more likely to vote for their party's candidate (based on the post-election survey; $B = .119$, Wald $\chi^2(1) = 5.29$, $p = .021$) than those who had perceived their party to be more causally peripheral. We also found a main effect of party ($B = -3.96$, Wald $\chi^2(1) = 6.73$, $p = .019$) such that Democrats were significantly more likely to vote with their party. Exploratory analysis revealed that the effect of causal centrality of political party on voting was driven primarily by the voting behavior of Republicans ($B = .264$, Wald $\chi^2(1) = 6.337$, $p = .012$), and no effect was observed for Democrats ($B = .003$, Wald $\chi^2(1) = .359$, $p = .549$). These effects held controlling for the total number of links and survey version (experimenter-generated vs. self-generated features), neither of which significantly predicted voting behavior.

The results of Study 2 extend our findings to a salient social category, political identity during a major election. We found that people who believed that their political identity was causally connected to more aspects of their self-concept were more likely to vote with their party. We have argued that fully understanding which members of a social category will behave in identity-consistent ways requires understanding how the social category fits into the broader self-concept. These results are consistent with our approach to identity-based choice, in which the driver of identity-consistent behaviors is not only membership in a social category (or possessing an aspect of identity) but also how causally connected that aspect is to other important features of personal identity.

STUDY 3A: SPORTS FAN IDENTITY AND SUPER BOWL TICKETS

The results of Studies 1 and 2 support our hypothesis that those who believe that an identity is causally central will be more likely to behave in identity-consistent ways. In Studies 3A and 3B, we extend this finding to a new social category, being a fan of a football team. As in Study 2, we examine this identity during a time when it is likely to be salient, the week of the Super Bowl.

Method

Three hundred sixty-six residents of the two states that were the home of the two teams in the 2016 Super Bowl (North Carolina and Colorado) were recruited from an online paid panel, and completed the study one to three days after the Super Bowl. Twenty-six participants were excluded for not providing valid answers to the focal questions in the survey and thirty participants were excluded for failing an attention check, yielding 310 valid cases.[1]

After completing an unrelated task from a different study (on hedonic adaptation), participants were asked to list the ten things that most define who they are as a person (i.e., features of identity). They then reported whether they were a Denver Broncos fan, a Carolina Panthers fan, a fan of another team (which they specified), or not interested in football. Participants who reported not being football fans (N = 57) were excluded from the analyses, leaving 253 participants who all reported having the identity of football fan.

[1] Participants were excluded for not answering the open-ended identity feature questions (23), providing the same answer for all identity questions (2), or for providing a willingness to pay of $1,000,000,000,000,000 (1).

Participants completed a shorter version of the listing causal relationships task than in Studies 1 and 2, comprised of the ten features they self-generated and six additional features: being a fan of the team they specified, childhood memories, personal life goals, friendships, values and principles, and personality.

The task focused on two features, being a fan of their favorite football team and the fifth feature that the participant had listed. We elicited the causal centrality of the fifth feature as a control, to account for potential differences in the general tendency to report features of the self-concept as causally linked. Participants completed two trials for each of the target features: one that measured the number of other features causing target feature (i.e., the feature's causes) and another that measured the number of other features caused by the target feature (i.e., the feature's effects). For example, a participant who reported being a Carolina Panthers fan would first be asked which other aspects of her identity *caused her* to be a fan of the Carolina Panthers. She would then be asked which other aspects of her identity were *caused by* her being a fan of the Carolina Panthers. The causal centrality of being a Carolina Panthers was calculated by summing the number of features selected across the two trials.

Participants were then asked how much they would be willing to pay for a ticket to see their team play in the Super Bowl if their team made it the following year. Participants then reported measures of sports involvement (whether they knew who had won the Super Bowl, whether they had watched the Super Bowl, their interest in football, and how many hours per week they spent on sports including participating, watching, playing video games, etc).

Results

Football fan analysis. On average, participants reported that three other features were causally linked to being a fan of their favorite team. The average willingness to pay to see their team in the Super Bowl was $475.77. The average interest in football was 1.6 on a 4-point scale (1 = very interested, 4 = very uninterested).

We regressed willingness to pay on the causal centrality of being a fan, controlling for the casual centrality of the control feature. As predicted, those who perceived being a fan of their favorite team as more causally central in their self-concept were willing to pay significantly more than those who perceived being a fan as causally peripheral ($B = 57.742$, $p < .01$). On average, each additional causal link between football fandom and another feature of the self-concept was associated with being willing to pay an additional $58.

According to our account, football fans whose fandom is more causally central are more willing to pay to see their team in the Super Bowl because they perceive acting in identity-consistent ways as more congruent with who they are than those who perceive fandom as causally peripheral. However, it also possible that the causal centrality measure is merely capturing differences in the strength of people's preferences. To examine this, we tested whether the effect of causal centrality was explained by differences in the participants' level of interest in football. We added interest in football as an additional predictor in the regression reported above. As expected, participants who reported greater interest in football (indicated by lower numbers) were willing to pay more to see their team play in the Super Bowl ($B = -280.58$, $p < .001$). More importantly, participants for whom being a fan was more causally central were willing to pay more to

see their team in the Super Bowl ($B = 39.26$, $p < .05$), even controlling for self-reported interest in football (Table 1).

Table 1

MULTIPLE REGRESSION PREDICTING WILLINGNESS TO PAY FOR TICKET TO WATCH FAVORITE TEAM IN SUPER BOWL

Factor	Beta	Std Error	Wald	*p*
Constant	871.52	172.87	5.04	<.001
Fan Causal Centrality	39.26	129.70	1.99	.047
Control Causal Centrality	-8.53	12.00	-.71	.478
Interest	-280.58	78.03	-3.60	<.001

Note. Interest was coded as: 1 = Very interested, 2 = Somewhat interested, 3 = Somewhat uninterested, 4 = Very uninterested.

The results of Study 3A suggest that football fans who believe being a fan is causally central are more willing to spend in identity-relevant ways. We further demonstrated that causal centrality predicts identity-consistent behavior beyond simple measures of preference and frequency of participation in identity-related activities. This is consistent with the proposal that people for whom being a football fan is causally central act in identity-consistent ways because they believe it is more congruent with who they are in a broad sense, compared to those for whom fandom is causally peripheral.

STUDY 3B: THE SUPER BOWL REVISITED

Study 3B provides a direct replication test of Study 3A, and allows us to further explore the relationship between causal centrality and identity importance. As previously

described, the importance of an identity has been shown to moderate the effect of that identity's salience on identity-relevant choice (Reed 2004; LaBeouf, Shafir, and Bayuk 2010) and has been theorized to predict which identities will exert more influence on behavior (Markus and Wurf 1987). As causally central aspects of the self-concept are perceived to be more defining of the self-concept, one of the consequences may be that causally central aspects of identity seem more important. Thus, we predict that importance will partially mediate the impact of an identity's causal centrality on identity-relevant behavior.

Method

We recruited 398 Amazon Mechanical Turk participants from throughout the U.S. approximately 4.5 to 2.5 hours prior to the 2017 Super Bowl. Five participants were excluded for failing an attention check and one participant was excluded for providing the same answer for all the features of identity, yielding three hundred ninety-two cases.

Participants were asked to list the ten things that most define who they are as a person and, separately, whether they would describe themselves as a fan of NFL football. Only participants who answered yes to the football fan question moved on to the rest of the survey (N = 242). Participants then reported which team they considered themselves a fan of.

We measured the importance (Reed 2004) and esteem (Shang, Reed, and Croson 2008; Luhtanen and Crocker 1992) that participants associated with being a fan of their favorite NFL team. The importance scale is composed of three items which ask how much participants felt being a fan of a team describes who they are, how much they identify with that group, and how much they admire the group. The esteem scale is

intended to measure the participants' perceived standing in the group. It is composed of four items that ask how worthy a group member they are, how much they have to offer the group, if they are a cooperative participant in group activities, and if they often feel as if they are a useless group member (reverse scored).

We included the esteem measure because previous literature on identity-based choice has found that it moderates identity congruency effects (Reed, Shang, and Croson 2008). That is, people who have high identity esteem are more likely to imitate the actions and choices of someone who shares that identity (even if these actions are not associated with the social category) than those with low identity esteem. As our Study 3B does not involve social congruency effects, the addition of esteem was exploratory and we did not have an a priori prediction about its relationship with causal centrality.

Participants performed the same abbreviated version of the listing causal relationships task from Study 3A, in which they reported the causes of and effects of being a fan of their team (from a list of the ten self-generated features and six additional features: being a fan of the team they specified, childhood memories, personal life goals, friendships, values and principles, and personality). In this study, we added an additional feature, "Level of Hunger," as the control feature because in previous studies it consistently participated in very few causal relationships, making it a good indicator of whether participants were inclined to report relationships merely because that is what the task involved.

Participants then reported how much they were willing to pay to watch their team play in the Super Bowl, how likely they were to watch the 2017 Super Bowl, how

interested they were in football, and how many hours per week they spent on sports. Participants then completed a set of questions from an unrelated study.

Results

On average, participants reported four links to being a fan of their favorite team and a willingness to pay to see their team in the Super Bowl of $538.18. The average interest in football was 1.5 on a 4-point scale (1 = very interested, 4 = very uninterested).

We regressed willingness to pay on the causal centrality of being a fan, controlling for the casual centrality of the control feature. As predicted, those who perceived that being a fan of their favorite team was more causally central in their self-concept were willing to pay significantly more than those who perceived that being a fan was causally peripheral ($B = 33.74$, $p = .027$). On average, each additional causal link between football fandom and another feature of the self-concept was associated with being willing to pay an additional $34.

We added interest in football as an additional predictor in the regression reported above. As expected, participants who reported greater interest in football were willing to pay more to see their team play in the Super Bowl ($B = -208.87$, $p < .001$). When controlling for interest in football, participants for whom being a fan was more causally central were marginally more willing to pay more to see their team in the Super Bowl ($B = 27.94$, $p = .066$). While we find a slight reduction in the effect of causal centrality in this study when controlling for interest in football, interest did not significantly mediate the effect of causal centrality.

Next, we tested how the effects of causal centrality relate to prior findings suggesting that identity-consistent behaviors are influenced by identity importance and

esteem (Reed 2004; Shang, Reed, and Croson 2008). We first regressed willingness to pay on importance and esteem. Esteem and importance were calculated by averaging the answers to the questions in each scale. While importance was significant in this regression ($B = 95.49$, $p = .015$), there was no significant relationship with esteem ($B = 44.88$, $p = .166$), so the rest of our analysis focuses on importance.

We conducted a mediation analysis to test whether some of the relationship between causal centrality and valuation operates via importance of being a football-team fan. We found a significant indirect effect of causal centrality on willingness to pay via fan identity importance ($B = 11.34$, 95% Bootstrapped CI = [4.76, 20.87]). The relationship between causal centrality and willingness to pay was mediated by importance and no longer significant when controlling for importance ($B = 17.80$, $p = .198$).

The results of Study 3B replicate the findings from Study 3A. Further, we found that importance of being a football fan mediates the effect of causal centrality on willingness to pay. These results suggest that one way causal centrality can influence identity-consistent behavior is by making an aspect of identity seem more important. The next study seeks to further clarify the relationship between a feature's causal centrality and its importance.

STUDY 4: CHEATING AND THE CAUSAL CENTRALITY OF HONESTY

In Study 4, we examine whether the causal centrality of a moral quality, honesty, explains differences in identity-*inconsistent* behavior, cheating. To do so, we measured

the causal centrality of honesty and had participants complete an incentivized task in which they had an incentive to be dishonest (cheat) because doing so would increase their compensation for the study. We predicted that participants who perceived honesty as causally central would display less cheating than those who perceived honesty as causally peripheral.

This study also allows us to clarify the relationship between the causal centrality of a feature and feature importance. Study 3B demonstrated that importance at least partially mediated the relationship between causal centrality of being a fan and willingness to pay, in a context where there is a fairly direct relationship between importance of the identity and the role of the identity. We have proposed that this occurs because importance is one of the consequences of causal centrality.

In our view, causal centrality represents how people think about their self-concept. In contrast, importance of an identity aspect involves a more complex judgment that incorporates not only the aspect's actual role in the current self-concept, but other factors, including the role people would like the aspect to have and social expectations. As a result, we predict that when causal centrality and importance diverge, causal centrality will be more strongly related to behavioral outcomes. Honesty is an appropriate feature to test whether the causal centrality of a feature predicts behavior over and above importance, as moral qualities are generally perceived to be very important but people vary in how causally central they believe these moral qualities to be (Chen, Urminsky, and Bartels 2016).

Method

Eighty-one U.S. Amazon Mechanical Turk participants completed the study. Five participants were excluded for failing an attention check, yielding seventy-six cases.

Participants completed two ostensibly unrelated studies: a study about identity (listing causal relationships task) and a study about perceptions of probability (a coin flip task). In the study about identity, participants listed the two most important features of their identity in each of five categories that had been identified in previous research as being important to identity (memories, goals/desires, preferences, moral qualities, personality traits; Chen, Urminsky, and Bartels 2016; Strohminger and Nichols 2014). After practicing the listing causal relationships task with an unrelated concept, participants completed the listing causal relationships task with the ten features of their identity that they had reported and three additional features: honesty, relationships with family and close friends, and relationships with significant others (past or present). The two relationship features were included because, in previous studies, they were commonly listed when participants were asked to list the most important features of their identity.

In the second part of the study, participants were asked to flip a coin ten times and then report the number of times the coin came up heads. Participants were told that they would win a five cent bonus every time the coin came up heads. As a result, they could infer that being dishonest would increase their compensation. While we cannot evaluate whether specific individuals cheated, this paradigm allows us to evaluate the average level of cheating, by comparing the average reported wins (heads on the coin flip) to the statistically predicted rate of wins (50%). This method has been used in previous research to study cheating (Bryan, Adams, and Monin 2013).

Participants then answered two questions about their beliefs about the likelihood of the outcome of the coin flip task (the probability of their reported pattern of coin flips occurring, and the average number of wins other participants got). These questions were unrelated to the causal centrality task and did not enter into any of our hypotheses. Finally, participants reported their own level of honesty on a scale of 1 (not at all honest) to 7 (extremely honest) and how important honesty was to them on a scale of 1 (not at all important) to 7 (extremely important).

Results

The mean number of heads reported (6.19 out of 10) was significantly greater than chance (5 out of 10) ($t(75) = 4.738$, $p < .001$, 95% CI = [5.69, 6.70]). So, overall there was evidence of cheating. Nevertheless, on average participants reported high honesty levels ($M = 6.37$) and high importance of honesty ($M = 6.47$). In fact the majority of participants reported a score of 6 or 7 (on a 7-point scale) for level of honesty (87%) and importance of honesty (84%).

A comparison of the means of those who perceived honesty as causally central vs. those who perceived honesty as causally peripheral (based on a median split) illustrates the degree to which people who believe that honesty was linked to more features of the self-concept were less likely to cheat. Participants who perceived honesty as causally central reported significantly fewer heads ($M = 5.5$) than those who perceived honesty as causally peripheral ($M = 6.9$, $t(74) = 3.014$, $p < .01$, 95% CI = [2.40, -0.49]). The average number of heads reported by those who perceived honesty as causally central was not significantly greater than what is expected by chance (5 heads) ($M = 5.5$, $t(37) = 1.35$, $p = .186$, 95% CI = [4.67, 6.19]), while the average number of heads reported by

those who perceived honesty as causally peripheral was significantly greater than change (M = 6.9, t(37) = 5.867, p < .001, 95% CI = [6.26, 7.59]).

To confirm the observed relationship, we regressed the number of claimed wins on causal centrality of honesty, controlling for self-reported level and importance of honesty. Participants who reported higher levels of honesty were no less likely to cheat than those who reported a lower level of honesty (B = .021, p < .01). Similarly, participants who reported that honesty was more important were no less likely to cheat than those who reported that honesty was less important (B = .007, p < .01). Consistent with the causal centrality account, however, participants who listed more links between honesty and other aspects of their identity claimed fewer wins than those who reported fewer links to honesty (B = .061, p < .01) even when controlling for importance and level of honesty.

The results of Study 4 demonstrate that the casual centrality of a feature can predict behavior above and beyond the importance of that feature. This suggests that causal centrality of features is a particularly useful predictor of behaviors in situations in which there is little variation in the importance of a feature or when people may be unable or unwilling to provide accurate importance ratings.

GENERAL DISCUSSION

Our approach to identity-based choice extends beyond individual aspects of identity to how these aspects fit within the broader self-concept. More specifically, we argue that an aspect of identity will exert more influence on behavior when it is perceived

as more causally central. We find that among people who belong to the same social category, those for whom the social category is perceived as more causally central are more likely to act in identity-consistent ways, compared to those for whom the same social category is more causally peripheral (Studies 2, 3a, and 3b). These studies were done during times when the social categories we examined—political identity and being a football fan—were very salient, during the 2016 Presidential election and the 2016 and 2017 Super Bowls. So, difference in the salience of these identities cannot explain the difference in identity-consistent behavior.

Unlike previous approaches to identity and choice, which have focused on social categories, a unique aspect of our causal centrality account is that any aspect of identity that is causally central can influence identity-consistent behaviors, regardless of whether the identity aspect is associated with a well-defined social category. Indeed, we found that features of identity such as goals (the desire to be environmentally-friendly) and moral qualities (honesty) were more influential on identity-consistent behaviors among people for whom these aspects were causally central than people for whom they were causally peripheral (Studies 1 and 4). Furthermore, Study 4 finds that the causal centrality of honesty predicted honesty behavior (i.e., less cheating) above and beyond measures of importance. This suggests that causal centrality is a useful tool for predicting identity-consistent behaviors when differences in measured importance are not fully diagnostic.

Theoretical Implications

The results of the present research have significant theoretical implications. Our approach to identity-based behavior identifies which consumers are most likely to behave in identity-consistent ways. While prior research in consumer behavior has examined

how membership in social categories influences choice (e.g., Brough et al. 2016, LeBoeuf, Shafir, and Bayuk 2010; Reed 2004), this research is the first to demonstrate that the causal relationships that the social categories participate in are related to identity-based choice. In doing so, we provide a more nuanced understanding of how membership in social categories relates to choice and why membership of the same social category differ in their identity-consistent behaviors. Additionally, we demonstrate that a wide-range of aspects of personal identity, not necessarily associated with a social category, relate to identity-consistent choice. In doing so, we connect research on identity-based choice and the self-concept.

As discussed earlier, some economic models of utility incorporate identity by grouping people based on how much they accept the norms of their social category. That is, how much utility an individual gains by acting in identity-relevant ways depends on how much they have embraced the social category. These models start with the assumption that different people embrace a social category to different degrees and do not attempt to explain these differences in adoption of an identity. We demonstrate, consistent with the model assumptions, that people who belong to the same category do indeed integrate social categories into their self-concepts to different degrees and that these differences have implications for choice. Further, our approach to identity provides a psychological explanation for what it means to adopt an identity and how adopters differ from non-adopters.

Implications for Marketers

The findings have important implications for how marketers target and communicate with consumers. In Study 1, we demonstrated that consumers for whom a

preference for buying environmentally-friendly products was causally central indicated greater willingness to pay for these products, compared with consumers for whom this desire was causally peripheral. This finding provides a potentially more effective means to identify the most receptive consumers as a targeting strategy. Furthermore, this suggests that causally linking a preference to other features of consumers' self-concepts could provide additional motivation for making choices consistent with those preferences. For example, for an Apple-user for whom that identity is central, buying a Dell is perceived negatively not only because she prefers the features of Apple computers to those of Dell computers, but also because this purchase feels incongruent with who she is. Thus, effective marketing may need to not only establish a preference for the product but also causally link this preference to important features of the consumers' self-concepts.

REFERENCES

Akerlof, George A., and Rachel E. Kranton (2000), "Economics and Identity," *The Quarterly Journal of Economics,* 115, 715–753.

Akerlof, George A., and Rachel E. Kranton (2010), *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being.* Princeton, NJ: Princeton University Press.

Ahn, Woo-Kyoung (1999), "Effect of Causal Structure on Category Construction," *Memory & Cognition*, 27 (6), 1008-1023.

Ahn, Woo-Kyoung, Nancy S. Kim, Mary E. Lassaline, and Martin J.Dennis (2000), "Causal Status as a Determinant of Feature Centrality," *Cognitive Psychology,* 41, 361-416.

Bartels, Daniel M., and Lance J. Rips (2010), "Psychological Connectedness and Intertemporal Choice," *Journal of Experimental Psychology: General,* 139, 49-69.

Bartels, Daniel M., and Oleg Urminsky (2011), "On Intertemporal Selfishness: How the Perceived Instability of Identity Underlies Impatient Consumption," *Journal of Consumer Research*, 38 (1), 182-198.

Bartels, Daniel M., and Oleg Urminsky (2015), "To Know and to Care: How Awareness and Valuation of the Future Jointly Shape Consumer Spending," *Journal of Consumer Research* 41, 1469-1485.

Belk, Russell W (1988), "Possessions and the Extended Self." *Journal of Consumer Research,* 15, 139-168.

Benjamin, Daniel J., James J. Choi, and A. Joshua Strickland (2007), "Social Identity and Preferences," rep., w13309. National Bureau of Economic Research.

Blok, Sergey, George Newman, and Lance J. Rips (2005), "Individuals and their concepts," in *Categorization Inside and Outside the Lab,* eds. Woo-Kyoung Ahn, Robert. L. Goldstone, Brad C. Love, Arthur B. Markman, & Phillip Wolff, Washington, DC: American Psychological Association, 127-149.

Brough, Aaron R., James E. B. Wilkie, Jingjing Ma, Mathew S. Isaac, David Gal (2016), Is Eco-Friendly Unmanly? The Green-Feminine Stereotype and Its Effect on Sustainable Consumption. *Journal of Consumer Research*, 43(4), 567-582.

Bryan, Christopher J., Gabrielle S. Adams, and Benoît Monin (2013), "When Cheating Would Make you a Cheater: Implicating the Self Prevents Unethical Behavior," *Journal of Experimental Psychology: General*, 142(4), 1001.

Chen, Stephanie Y., Oleg Urminsky, and Daniel M. Bartels (2016), "Beliefs About the Causal Structure of the Self-Concept Determine Which Changes Disrupt Personal Identity," *Psychological Science*, 27, 1398-1406.

Cohn, Alain, Ernst Fehr, and Michel André Maréchal (2014), "Business Culture and Dishonesty in the Banking Industry," *Nature,* 516(7529), 86-89.

Ersner-Hershfield, Hal, M. Tess Garton, Kacey Ballard, Gregory R. Samanez-Larkin, and Brian Knutson (2009), "Don't Stop Thinking About Tomorrow: Individual Differences in Future Self-Continuity Account for Saving," *Judgment and Decision Making*, 4 (4), 280.

Escalas, Jennifer Edson, and James R. Bettman (2003), "You are what they eat: The influence of reference groups on consumers' connections to brands." *Journal of Consumer Psychology,* 13, 339-348.

Forehand, Mark R., Rohit Deshpande´, and Americus Reed II (2002), "Identity Salience and the Influence of Activation of the Social Self-Schema on Advertising Response," *Journal of Applied Psychology*, 87 (6), 1086–99.

Gopnik, Alison, and Henry M. Wellman (1994). "The 'Theory Theory,' in *Mapping the Mind: Domain Specificity in Culture and Cognition,* Lawrence Hirschfeld and Susan Gelman, eds. New York, NY: Cambridge University Press, 257–293.

Gopnik, Alison, David M. Sobel, Laura E. Schulz, and Clark Glymour (2001), "Causal Learning Mechanisms in Very Young Children: Two, Three, and Four-Year-Olds Infer Causal Relations From Patterns of Variation and Covariation," *Developmental Psychology*, 37 (5), 620-629.

Kaplan, Audrey S., and Gregory L. Murphy (2000), "Category Learning with Minimal Prior Knowledge," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26, (4) 829.

Keil, Frank, C. (1989), *Concepts, kinds, and cognitive development.* Cambridge, MA: MIT Press.

_____ (2006), "Explanation and Understanding," *Annual Review of Psychology*, 57, 227–254.

Kleine, Robert E., Susan S. Kleine, and Jerome B. Kernan (1993), "Mundane Consumption and the Self: A Social-Identity Perspective," *Journal of Consumer Psychology*, 2 (3), 209-235.

Oyserman, Daphna, Stephanie A. Fryberg, and Nicholas Yoder (2007), "Identity-Based Motivation and Health," *Journal of Personality and Social Psychology,* 93 (6), 1011−1027.

Pazzani, Michael J. (1991), "Influence of Prior Knowledge on Concept Acquisition: Experimental and Computational Results," *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 17 (3), 416-432.

Reed II, Americus (2004), "Activating the Self-Importance of Consumer Selves: Exploring Identity Salience Effects on Judgments," *Journal of Consumer Research*, 31 (2), 286-295.

Reed II, Americus, Mark Forehand, Stefano Puntoni, and Luk Warlop (2012), "Identity-Based Consumer Behavior," *International Journal of Research in Marketing*, 29 (4), 310-321.

Rehder, Bob (2003), "A Causal-Model Theory of Conceptual Representation and Categorization," *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 29 (6), 1141-59.

Rehder, Bob and Reid Hastie (2001), "Causal Knowledge and Categories: The Effects of Causal Beliefs on Categorization, Induction, and Similarity," *Journal of Experimental Psychology: General,* 130 (3), 323-360.

Schulz, Laura E., and Alison Gopnik (2004), "Causal Learning Across Domains. Developmental Psychology, 40 (2), 162–176.

Shang, Jen, Americus Reed, and Rachel Croson (2008), "Identity congruency effects on donations." *Journal of Marketing Research* (45), 351-361.

Sloman, Steven, Bradley C. Love, and Ahn Woo-Kyoung (1998), “Feature Centrality and Conceptual Coherence,” *Cognitive Science,* 22 (2), 189–228.

Sperber, Dan, David Premack, and Ann James Premack, eds. (1995), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Strohminger, Nina, and Shaun Nichols (2014), “The Essential Moral Self,” *Cognition* 131, 159-171.

_____ (2015), "Neurodegeneration and identity." *Psychological Science*, 26, 1469-1479.

Smith, Edward E., and Douglas L. Medin (1981), *Categories and Concepts.* Cambridge, MA: Harvard University Press.

Turner, John C (1985), “Social categorization and the self-concept: A social cognitive theory of group behavior,” in *Advances in group processes: Theory and research (Vol 2)*, ed. E. J. Lawler, Greenwich, CT: JAI Press, 77–122.

Tversky, Amos (1977), “Features of Similarity,” *Psychological Review*, 84 (4), 327-352.

Urminsky, Oleg (2017) "The Role of ‘Psychological Connectedness to the Future Self’ in Decisions Over Time," *Current Directions in Psychological Science*, 26, 34–39.

Appendix

Figure 1A

CHOICES USED IN STUDY 1

Light bulbs		GE 60-watt, 4-pack A19 light bulbs for $4 OR GE LED 4-pack replacement 60-watt light bulbs for $19.99
		
Shopping bags		Single-use plastic bag for $0.10 OR Re-useable canvas bag for $3.99
		
Batteries	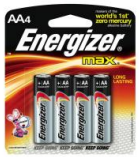	Energizer 4-pack of AA alkaline batteries for $4.99 OR Energizer 4-pack of AA rechargeable batteries for $13.99