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

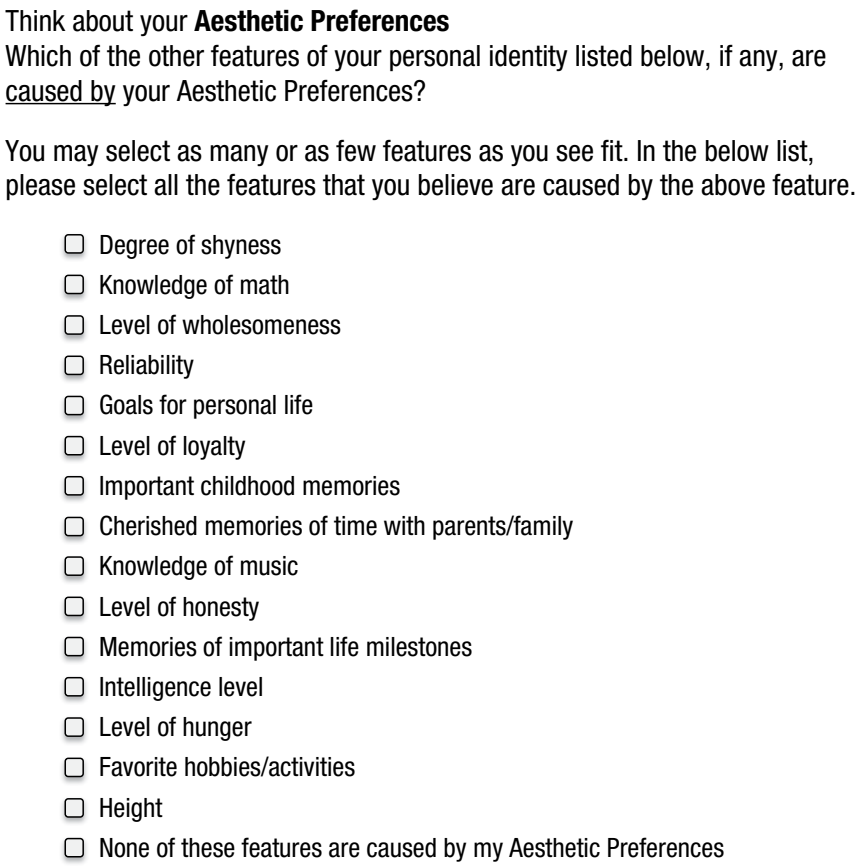

Fig. 1. Example of a question from the causal-relationships task used in Experiment 1.

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

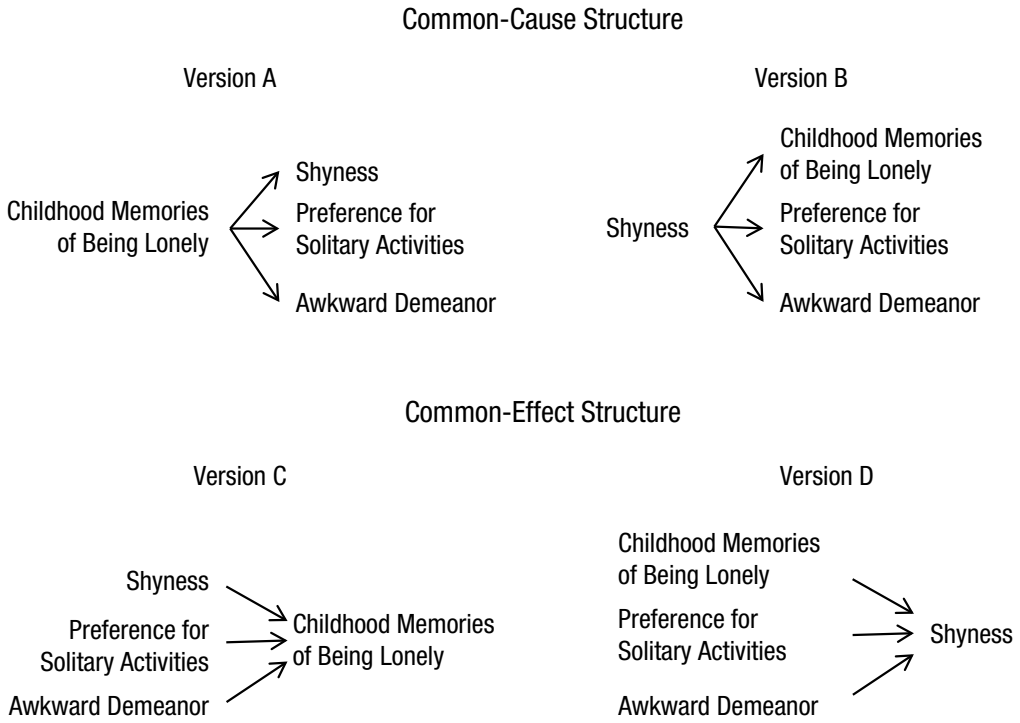

Fig. 2. Illustration of the structure of the vignettes used in Experiment 3. There were four versions of each vignette, two versions for each causal structure. The two versions for the common-cause structure referred to the same four features but differed in which feature was described as the cause of 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.