

What's left of me? The role of self-continuity in decision making and judgments about identity persistence

Stephanie Chen & Oleg Urminsky

For inclusion in Advances in Experimental Philosophy of Identity and the Self

Author biographies

Stephanie Chen is an assistant professor of marketing at London Business School.

Oleg Urminsky is a professor of marketing at the University of Chicago Booth School of Business.

Abstract

In this chapter, we discuss the concept of self-continuity and how it fits into the broader literature on the nature of people's representation of the self and judgments about whether the self persists over time. We then review recent research that examines how beliefs about self-continuity underlie a wide range of decisions.

Introduction

Personal identity is an important driver of choice and behavior. While a large literature in philosophy and psychology has explored questions of how people represent the self-concept and what changes people think will disrupt their identity, a growing literature has focused on understanding how these beliefs underlie decision making. By connecting beliefs about personal identity with common decisions like financial and consumption choices, we explore how people are likely to think about the self in their everyday lives.

In this chapter, we first situate our work on personal identity and the self-concept in the broader literature about personal identity. A central distinction in the philosophical study of personal identity is the distinction between numerical and qualitative identity. Recent accounts have debated whether these two types of identity have been confused in recent research on personal identity (Dranseika 2017; Starmans and Bloom 2018a). We suggest that the types of beliefs about the self that underlie decision making are those that are likely to reflect how people think about the self in everyday life. We further suggest that the type of beliefs about identity that track these decisions, self-continuity beliefs, do not fit neatly into the numerical-qualitative identity dichotomy and challenge current ideas about what identity researchers are and should be studying.

We then review emerging research that we believe leverages a more psychologically realistic conceptualization of how people think about identity to explain decisions that are driven by beliefs about personal identity. First, we review research on how intertemporal choice is influenced by psychological connectedness to the future self, the extent to which a person believes that she shares key psychological characteristics with her future self. Second, we examine choices that are driven by the norms of the various social categories that people belong to. We review work that suggests that how likely a person is to make choices consistent with the norms of a social category is a function of her beliefs about how that social category is causally related to the other aspects of her self-concept.

What is identity?

While the meaning of "identity" may seem intuitively obvious, different scholars in philosophy and psychology have brought different understandings and assumptions to their work on identity, resulting in different operational definitions. For example, social psychologists often treat a person's identity (or identities) in terms of social categories (Markus and Wurf 1987; Tajfel 1978), while work in philosophy and cognitive science has thought of identity in terms of the features that would need to be preserved to maintain identity (Blok et al. 2005; Strohminger and Nichols 2014).

A recent debate has attempted to clarify the meaning of identity and to question whether researchers' operationalizations in studying identity have actually explored the construct that they conceptually intended to study (Starmans and Bloom 2018a; 2018b; DeFreitas et al 2018). Starmans and Bloom (2018a) highlight the important distinction between *numerical* and *qualitative* identity in people's intuitions about the identity of the self and others undergoing change. Numerical identity concerns whether the person literally ceases to exist after a change, such that the person who occupied the body is no longer there, akin to death of the individual, even if the body is still alive. By contrast, qualitative identity relates to the similarity between the person at different points in time (e.g., before vs. after a change), defined as the degree to which the post-change person's features overlap with the original person's features.

Starmans and Bloom (2018a) suggest that researchers aiming to study changes in numerical identity (i.e., whether a person literally ceases to exist after a change) have instead studied changes to qualitative identity (i.e., changes to similarity). Consider a moral person who suffers a cognitive impairment and begins behaving immorally, as in the classic case of railway worker Phineas Gage (Damasio et al. 1994). When family members say that the person no longer seems like the same person underneath or that the person seems like a stranger (e.g., as was found for changes in morality, in Strohminger and Nichols 2015), what does that mean? Do they mean that the person is no longer there, has ceased to exist and there is another person in the physical body? Or do they mean that the same person is still there, but is now so dissimilar from the person they used to be that loss-of-identity metaphors are appropriate for characterizing the large degree of dissimilarity?

While clarity on these concepts is important, the distinctions being drawn in this debate may perpetuate a false dichotomy, between literal cessation of existence at one extreme and "mere" dissimilarity of any kind (e.g., reduction in feature-overlap between the past and present self, Starmans and Bloom 2018a). In our view, much recent research on the role of identity in people's lives instead explores a third type of identity, distinct from both numerical and qualitative identity. This type of identity can be referred to as *self-continuity*¹ (e.g., Urminsky and Bartels 2019; Molouki and Bartels 2017), building on characterizations of change in the self over time and partially overlapping selves in Parfit (1984). We believe that the notion of self-continuity best aligns with how people make subjective judgments about their own and others' identities in everyday situations, particularly ones that underlie common decisions.

Assessing self-continuity and persistence of subjective identity

Self-continuity judgments are not all-or-nothing numerical identity judgments about whether a change causes the prior self to literally cease to exist. Instead, they are continuous judgments about how much of the self, in terms of people's subjective perceptions of what defines a person as oneself, persists over time (Bartels and Rips 2010; Bartels and Urminsky 2011; Chen, Urminsky and Bartels 2016; Mott 2018). Going from being a shy child to an outgoing teenager is unlikely to make a person feel that they have ceased to exist, but it may make her feel like she's changed into a somewhat different person than she had been, perhaps losing an important part of herself or gaining a new aspect of the self in the process.

¹ The notion of self-continuity we introduce here is distinct from notions of psychological continuity that are characterized by lack of variability in psychological connections over time (Parfit 1984).

Importantly, self-continuity judgments are not simply defined by feature overlap (e.g., qualitative identity), because changes to some features are more important to the definition of the self than others. In fact, we have found that changing the same number of features (i.e., equivalent changes in feature overlap/qualitative identity) can result in different judgments about self-continuity (Chen et al. 2016), depending on which features change. While becoming an outgoing teenager may make a previously shy child feel a sense of being a somewhat different person, going from always having long hair as a child to having short hair as a teenager likely doesn't make one feel like a significantly different person. That said, which features disrupt subjective self-continuity is idiosyncratic—giving Rapunzel or Samson a haircut might make them feel more like a different person than would giving the average person a haircut. In our view, much recent research should be understood as aiming to examine which features are weighted most heavily in such subjective judgments of self-continuity (e.g., Blok et al. 2005; Haslam et al. 2004; Strohminger and Nichols 2014) and, importantly, why.

Consistent with this conceptualization, many explorations of personal identity use continuous measures of self-continuity judgments, asking participants to rate to what extent a change will make them into a different person (see Figure 1 for an example). This is similar to the notion in cognitive psychology that categorization judgments are, ironically, continuous, not categorical. Items vary on how likely they are to be categorized in a given category and how typical they are of the category. Typical items are those that not only share a lot of properties with other category members but also do *not* share properties with the members of other categories. Atypical members of a category are the converse: they share fewer properties with category members and tend to resemble members of other categories (Rosch and Mervis 1975).

For example, a platypus is a terrible member of the mammal category, in part, because it has properties that make it resemble the bird category (e.g., lays eggs, has a beak). Nonetheless, it is still (generally) categorized as a mammal. In the same way, some changes to personal identity may be significant enough that an individual starts to resemble a “new” individual and become a less good example of the previous version of the self, despite remaining the same individual in numerical terms. Thus, we argue that characterizations of change in terms of loss of identity (e.g., no longer the same person, a stranger) often do not constitute metaphorical exaggeration, but instead may be accurate representations of the decline in perceived self-continuity over time: some of what made you the person you were then is no longer part of who you are now.

If self-continuity is not mere similarity, but is instead determined by change specifically in those features of the self that are perceived as defining of the person, what determines the degree to which a given feature is identity-defining? Building on the idea that people fundamentally reason in causal terms, including in reasoning about numerical identity in objects (Rips et al 2006), a growing body of research suggests that causal beliefs shape self-continuity judgments (Chen et al 2016). Unlike most people (who likely see their hair as causally linked to few features of the self), Rapunzel likely believes that her hair is causally linked to many aspects of the self—e.g., her relationship with the prince, her role as a prisoner, and her magical powers. For Rapunzel, changing that feature would therefore likely be more disruptive to her self-continuity than it would be to most others' self-continuity.

Furthermore, deviations from a person's expected causal trajectory of the self are seen as more disruptive of self-continuity than expected changes (Molouki and Bartels 2017). For a person who expected to remain a shy wallflower forever, becoming outgoing will disrupt self-continuity more than if that person expected to eventually become popular and outgoing. Thus, the exact same change (and the same difference in mere similarity between an individual before and after the change) will influence self-continuity to the degree that it is expected. This can help explain research suggesting that the direction of change matters to persistence of identity judgements—improvements in moral qualities are seen as more consistent with the self than declines in these traits (Newman, De Freitas, and Knobe 2015). If one expects to improve over time (as most do, including in their moral qualities), then improvements will match expectations and preserve self-continuity. However, in those cases where people expect to worsen, anticipating improvements disrupts self-continuity (Molouki and Bartels 2017).

If the goal is to understand how people commonly think about their own and to explore the role of identity change in decision making, we believe that self-continuity is the most relevant and useful construct to study. Both researchers and lay-people often tend not to think in terms of numerical identity, despite its philosophical significance, because literally ceasing to exist rarely occurs. For example, Starmans and Bloom (2018b) suggest that even the case of Phineas Gage's dramatic transformation is not a change in numerical identity. It follows that most changes that people experience and think about are not about numerical identity. However, we believe that people have a richer and more nuanced understanding than is captured by qualitative identity, that is better captured by self-continuity. Framing identity in terms of self-continuity enables us to study the role of identity beliefs in decision-making and explore why changes to some features impact self-continuity more than others.

Psychological connectedness and future-oriented decisions

Many of the decisions that people face can be characterized as trade-offs between immediate and future benefits. The decision to save money for the future means forgoing the enjoyment that you would get from spending that money in the present. Working hard today to ensure that you have better opportunities in the future means giving up all the fun that you could have now if you weren't working. A large literature on such intertemporal choices documents that people are generally impatient; they prefer immediate rewards and discount future rewards in ways that cannot be explained by rational (economic) considerations (see Frederick, Loewenstein, and O'Donoghue 2002 and Urminsky and Zauberman 2016 for reviews).

While people's intertemporal choices are influenced by a number of factors, recent research has identified an aspect of self-continuity, *psychological connectedness*, as a key driver of these choices. Psychological connectedness was first defined by philosopher Derek Parfit (1984) as the degree of overlap in psychological traits (e.g., one's values, memories, ambitions, disposition, etc.) that people believe exists between their current and future selves. Parfit (1984) makes a normative argument that people ought to discount future outcomes more when they are less psychologically connected to their future selves. This because people who are more connected to the future self will be the ones reaping future benefits that they have sacrificed immediate benefits for, more so than people who are less connected to the future self. While such normative claims are controversial, Parfit's ideas provide a useful framework to examine the descriptive psychological drivers of intertemporal choice. The research reviewed in this section leaves aside the question of whether people *ought to*

discount future rewards to the extent that they are psychologically connected to the future self and instead asks the question of whether people *do* act in this way.

Recent research has found that people do, in fact, discount future rewards more when they are less connected to their future selves (see Urminsky 2017 for a review). For example, Bartels and Urminsky (2011) manipulated psychological connectedness among a set of graduating college seniors by having them read passages about the stability of personal identity. Those in the low connectedness condition read that graduation was likely to change many of their important psychological traits (e.g., “The characteristics that make you the person you are”). Those in the high connectedness condition read that personal identity was relatively stable and graduation was unlikely to change their core psychological traits. Participants then made a series of trade-offs about a real lottery, between receiving a \$120 gift certificate in one week and a larger reward later in one year. Consistent with the hypothesis that people who are more connected to the future self will be more willing to forgo more immediate rewards for future ones, participants in the high connectedness condition were significantly more willing to wait for the larger reward (requiring \$16 less to wait, on average) than those in the low connectedness condition.

A number of researchers have also found that lower measured psychological connectedness to the future self predicts greater discounting of both real and hypothetical future rewards (Bartels and Rips 2010, Bartels and Urminsky 2011, Ersner-Hershfield et al. 2009; but not in Frederick 2003). For instance, Bartels and Urminsky (2011) measured psychological connectedness with continuous measures. First, participants rated how much they felt that the future self would be the same person as they are today on a scale of 0 (completely different to the current self) to 100 (exactly the same as the current self). Second, participants completed two visual measures in which they saw pairs of circles that overlapped to different extents, representing their perceived overlap with their future self (see Figure 1). Participants either selected the set of circles (Figure 1A) or marked the position on the line (Figure 1B) that they felt best represented how much overlap they saw between their current and future self. Bartels and Urminsky (2011) found that a composite of these measures of psychological connectedness predicted participants’ willingness to forgo a payment in the present for a larger payment in the future. Further, using similar measures of psychological connectedness, Ersner-Hershfield et al. (2009) found that the degree of connectedness to the future self predicted differences in accumulated financial assets (controlling for age and education).

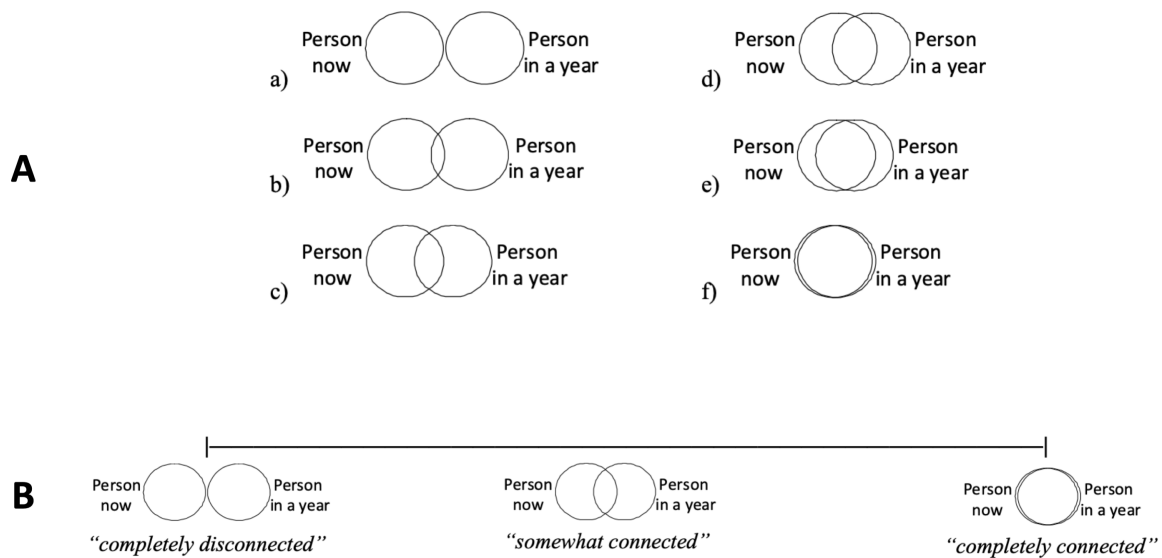


Figure 1. Visual measures of psychological connectedness used in Bartels and Urminsky (2011, study 5). Participants chose the set of circles (panel A) or marked the position on the line (panel B) that best represented their belief about how much their current self's defining psychological traits overlap with their future self's psychological traits.

Recent research has examined the conditions under which psychological connectedness is more likely to predict differences in intertemporal choice. More specifically, Bartels and Urminsky (2015) examined the relationship between psychological connectedness and consideration of opportunity costs. As people often neglect opportunity costs when making purchasing decisions (i.e., do not think about alternative uses for unspent money; Frederick et al. 2009, but see Spiller 2011 for exceptions), the authors hypothesized that consideration of opportunity costs may be necessary for psychological connectedness to influence decisions. That is, neglecting opportunity costs means not being aware that spending in the present often means forgoing future resources or consumption. Thus, when neglecting opportunity costs, people may not be motivated to forgo immediate benefits, no matter how connected they are to the future self, since they are not thinking in terms of depriving the future self of a benefit.

Bartels and Urminsky (2015) found that people did, in fact, need to consider opportunity costs for psychological connectedness to influence intertemporal choice. In one study, psychological connectedness and opportunity cost salience were independently manipulated and the dependent measure was participants' choices between more expensive and cheaper products in a number of product categories. When opportunity costs were made salient (by having participants rank the importance of the product categories), those in the high connectedness condition were more likely to choose cheap products than those in the low connectedness condition. However, when opportunity costs were not made salient, there was no difference in choices between the high and low connectedness conditions. Thus, consideration of opportunity costs may be necessary for psychological connectedness to influence choices with non-explicit intertemporal consequences. As these findings illustrate,

anticipated self-continuity can be an important input into decision-making, but its role depends on how the decision is conceptualized.

Causal centrality approach to the self-concept

What features define the self?

Thus far, we have suggested that people see some features as more defining of the self, and therefore as more relevant to assessments of self-continuity. However, this begs the question of which features are more defining and why. Notably, the empirical research on psychological connectedness discussed above essentially punts on this question, relying on people's idiosyncratic definitions and interpretations of what constitutes their "core" or "important" psychological traits.

A large literature has explored which features of the self people think defines their self-concept, such that a change in that feature would reduce self-continuity. Theories of identity-based consumption and choice have focused on social categories as central to the self-concept. In contrast, other disciplines have debated which individual-level aspects of the self are most defining of the self-concept. Autobiographical memories have long been suggested by philosophers to be defining of personal identity as they are unique to each individual. Consistent with this proposal, psychological studies have found that disrupting a person's memories leads to perceptions that the person has become a different person (Blok et al. 2005; Nichols and Bruno 2010). Other accounts have instead suggested that personality traits are particularly important in defining the self-concept (Haslam, Bastian, and Bissett 2004; Gelman, Heyman, and Legare 2007), while more recent research has emphasized moral qualities (Strohlinger and Nichols 2014, 2015). While these approaches select a different feature to place at the center of the self-concept, they all conceptualize the self-concept as defined by a set of core features in general.

Our *causal centrality* approach to the self-concept instead suggests that the representation of the self-concept is not simply a set of features but also critically includes beliefs about the causal relationships between these features. Our approach was inspired by research on conceptual representation in cognitive psychology which has found that features are defining of a concept to the extent that the features are seen as causally central (Rehder and Hastie 2001; Sloman et al. 1998)—causally connected to many other features of the concept. Following Rehder and Hastie (2001), we define causal centrality of a feature as the total number of other features of the self-concept a given feature is seen as causally linked to, as either a cause or an effect. This definition has most consistently explained participants' judgments across our explorations of self-concept representation and identity-based choice (and in particular better than a common alternative definition that suggests only causing other features counts in the calculation of causal centrality², Ahn et al. 2000; Sloman et al. 1998).

² For the self-concept specifically, there is also a theoretical reason to believe that a given feature's causal centrality should not only be determined by the number of other features it is the cause of. According to such models, causes are always more central than their effects. Since causes always occur before their effects, these models imply that people would always be more defined by things that happened and or the features that developed earlier in life than those that happened/developed later. By also including the number of features that a given feature is an effect of in our calculation of causal centrality, what is most defining of the self can change over time—i.e., features that develop later in life can become more defining of the self-concepts than their causes.

However, the issue of the role of causes vs. effects in defining causal centrality is still debated (Chen and Urminsky 2019; Rehder and Kim 2010).

While ideas about causal centrality have been influential in the study of people's concepts of categories, there are at least two important ways that the self-concept differs from the concepts in which causal centrality has previously been explored. First, the self-concept is a concept of a single individual (the self) and is not a concept of a category (a set of items). Second, the self-concept, unlike many of the artificial, biological, and artefact categories commonly studied, is one which people have a wealth of knowledge about and may have very idiosyncratic beliefs about. Thus, unlike the commonly studied concepts of categories, the self-concept is both a concept of an individual and a highly individualized concept.

The key prediction of our approach is that features of the self-concept would be seen as defining of the self-concept to the extent that they were seen as causally central. That is, changes to causally central features will disrupt self-continuity more than changes to causally peripheral features. To illustrate, imagine two academics who went to Princeton, Stephanie and Oleg. Stephanie believes that her experiences at Princeton shaped her choice of profession and many of her academic interests. Oleg instead believes that it was his academic interests that led him to both become an academic and to become a Princetonian. As a result, even though the features of Stephanie and Oleg's self-concepts are identical, the differences in their causal beliefs lead to differences in what they believe defines their self-concept. Because she believes that it is connected to more features of her self-concept (her interests *and* her profession), Stephanie will see being a Princetonian as more defining her self-concept than Oleg does (since he sees being a Princetonian as connected to his academic interests *only*). In contrast, Oleg's interests will be more defining of his self-concept because he sees them as relatively more causally central (connected to both being a Princetonian and to his profession). And, to foreshadow the next section on identity-consistent behavior, the causal centrality approach to the self-concept also predicts that Stephanie will be more likely to follow the norms of being a Princetonian (e.g., donating to the school) than Oleg will be.

As the above example demonstrates, causal beliefs about the features of the self-concept are subjective—people have different beliefs about how their features fit together. So, the causal centrality approach to the self-concept can explain why a given feature may be more important for some people than to others (i.e., because of differences in the subjective causal beliefs as in the example above). It can also explain why some features or feature types can, on average, be perceived as more defining of the self (i.e., because they are seen as more causally central by more people).

Chen, Urminsky, and Bartels (2016) tested the hypothesis that changes to features that are more causally connected are perceived as more disruptive to self-continuity than changes to features that are seen as less causally connected (i.e., more causally peripheral). They had participants report which cause-effect relationships they believed existed between 16 aspects of the self-concept that had been identified as important to the self-concept in previous research (memories, moral qualities, personality traits, preferences/desires; Strohminger and Nichols 2014, see Figure 2). As discussed above, the causal centrality of a feature was calculated by summing the number of other features a given feature was causally linked to as either a cause or an effect (Rehder and Hastie 2001). Participants also reported how defining each feature was to their self-concept by stating to what extent changing each feature would impact their self-continuity—i.e., to what extent they felt that they would still be the same person vs. a different person after the change. The results supported our hypothesis:

Participants reported that changes to causally central features were more disruptive to self-continuity (transformed them into a different individual to a greater extent) than changes to causally peripheral features. Further, people also saw causally central features as more defining of the self-continuity of *other* people, suggesting that causal relationships are an important part of not the self-concept but also of concepts of other people.

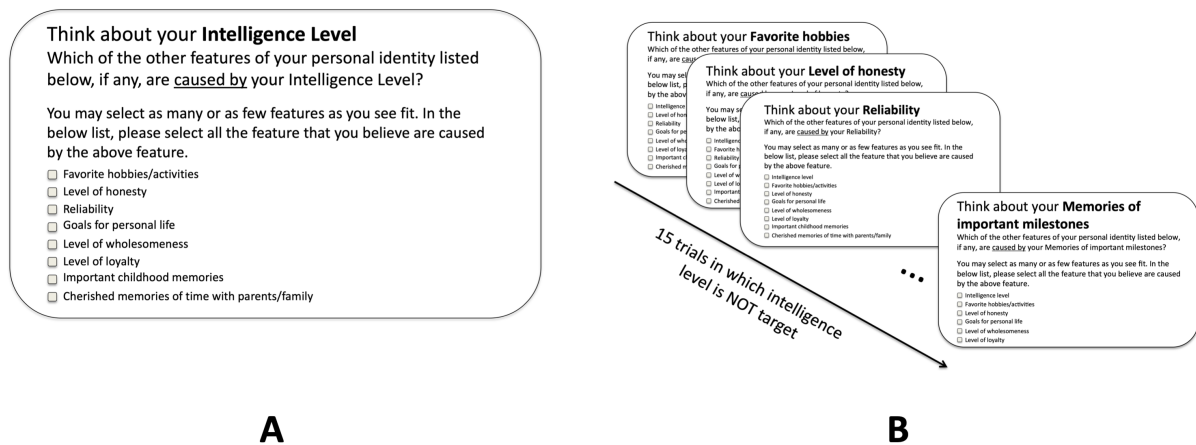


Figure 2. Illustration of task used in Chen et al. (2016) to measure the causal centrality of 16 features of personal identity. Each of the 16 features was the target feature for one trial, so participants completed 16 total trials. What follows is an explanation of how causal centrality was calculated for all 16 features using intelligence level as the example feature. Panel A. Trial in which intelligence level was the target feature. Participants selected all the other features that intelligence level caused from the checklist which included the 15 other features (truncated in the illustration). From this trial the number of features that intelligence level *caused* is calculated by summing the number of features selected as being caused by intelligence level in this trial. Panel B. In 15 trials, intelligence level was not the target and was listed in the checklist of possible effects of the target feature. The number of times intelligence level was selected as an effect of other features across these 15 trials was the number of things that intelligence was *caused by*. The sum of the number of features that intelligence level caused and the number of features that caused intelligence level was the causal centrality of intelligence level.

While the results of the experiments described above support the hypothesis that features that are seen as causally central are perceived as more defining of the self-concept, a perhaps more interesting question is whether the exact same feature can be made more or less defining by simply changing the feature's causal connections with other features—i.e., by changing people's beliefs about how causally central the feature is. This is an important question because many of the accounts of the self-concept discussed earlier emphasize feature type (moral quality, memory, social categories, etc.) as the main determinant of how defining the feature is. However, if changing the perceived causal connections of a feature with other features makes the exact same feature more or less defining, there has to be more that determines how defining a feature is than just feature type—i.e., its causal relationships.

To isolate the effect that causal centrality has on how defining a feature is, Chen et al. (2016) manipulated the causal centrality of a feature in descriptions of people's self-concepts. For example, half of the participants read that a person named Jack believes that his memories of being a lonely child caused him to develop a shy personality *and* a preference for solitary activities. Here, Jack's memories were relatively causally central because they were connected to two other features (his shy personality and his preferences for solitary activities). The other half of the participants read that Jack's memories were relatively causally peripheral; they read that Jack had the same three features but believed that his memories were causally connected to only one other feature (his shy personality). Participants were then introduced to a person who had all of Jack's features except for one (e.g., his memories) and judged whether the person was "still Jack". Supporting the theory, people were less likely to say a version of Jack without his memories was "still Jack" when those memories were manipulated to be causally central (as in the first scenario) than when they were manipulated to be causally peripheral. Across a number of features types, manipulating the changed feature to be causally central disrupted the (hypothetical person's) self-continuity more than manipulating the exact same feature to be causally peripheral.

Implications for identity-consistent behavior

The notion that more causally central aspects are seen as more defining of the self and more necessary for self-continuity provides a new perspective on the role of social category memberships (i.e., "*social identities*") in people's decisions. A large literature across social psychology, economics, and marketing suggests that behaviors and choices are influenced by the social categories that people belong to (e.g., Akerlof and Kranton 2000, 2010; Markus and Wurf 1987). This research suggests that people who belong to a social category are more likely to act in ways consistent with the norms of that category—display identity-consistent behaviors—than those who do not belong to the social category.

While social identity provides a good explanation for differences in behavior at a *category* level (e.g., people who are vs. are not in the category), it does not provide much insight into variance in identity-consistent behaviors at an *individual* level, among people who belong to the same social category. To explain such variance in identity-consistent behaviors, prior research has largely relied on differences in situational factors, particularly identity salience. In this view, members of a social category are likely to display identity-consistent behaviors to the extent that those social identities are made salient in the environment (LeBoeuf, Shafir, and Bayuk 2010; Reed II 2004). However, two people who belong to the same social category and are in similar situations may still behave differently. We propose that such differences in behavior can be explained, at least in part, by differences in people's internal representations of their self-concept.

The causal centrality approach provides a new way of thinking about differences across people in identity-consistent behaviors. As causally central features are seen as more defining of the self-concept, among people who hold a social identity, those who perceive the social identity as more causally central are predicted to be more likely to act in ways consistent with the norms of the group than those who see the same identity as more causally peripheral. The causal centrality account extends prior theoretical accounts in which aspects of the self influence behavior to the extent that they are seen as important to or central to the self-concept (e.g., Markus and Wurf, 1987) by providing a psychological explanation (based

on a large literature on conceptual representation) of what makes a particular aspect important to the self-concept.

During the 2016 US Presidential election, Chen and Urminsky (2019) investigated whether the causal centrality of political identities (Democrat and Republican) predicted identity-consistent behavior (voting for the party's candidate) among holders of each identity. The day before the election, participants reported the causal relationships that existed between the features of their self-concept, including political party. The day after the election, participants reported who they voted for and their satisfaction with their party's candidate. Running the study at a time when political identity would have been highly salient to everyone allowed us to focus on differences in internal representations, as differences in choice were unlikely to be due to differences in political identity salience.

Consistent with the hypothesis that the causal centrality of a social identity predicts how likely people are to act in identity-consistent ways, participants who saw their political party as more causally central were more likely to vote for their party's candidate,³ even when controlling for satisfaction with the candidate. These findings suggest that causal centrality of political identity is associated with greater norm-compliance, not just greater preference for the party's candidate (Chen and Urminsky 2019). Thus, among people who personally do not approve of their party's candidate, if their political identity is causally central enough, they may still behave in line with the norms of the group. Chen and Urminsky (2019) also found that among people who identified as British or English, the more causally central they said those identities were, the more likely they were to vote for Brexit, consistent with the perceived norms for those nationalities (as reported by the participants).

More recent research using both measured and manipulated causal centrality has extended these findings to a wide range of consumer behaviors. For example, Chen, Urminsky, and Yu (2020) found that the causal centrality of the environmentalist identity among self-identified environmentalists predicted reported willingness to purchase more expensive environmentally-friendly products (e.g., rechargeable batteries) over their cheaper traditional counterparts. Further, the causal centrality of the environmentalist identities predicted willingness to purchase environmentally-products even when the judgments of identity and the purchase task were separated in time. For example, in one study, the causal centrality of the environmentalist identity predicted willingness to purchase measured one year later. These results suggest that the causal centrality of a social identity predicts a wide range of choices and that an identity's causal centrality may generally be a relatively stable belief that predicts choices over time.

Conclusion

As the introduction to this chapter noted, there are many ways to think about personal identity and many judgments that can be made about the persistence of the self over time. While some researchers have aimed to study numerical identity and investigate what changes make a person literally cease to exist, we have aimed to study how people think about personal identity in their everyday lives. A key criterion that we used to identify the types of beliefs about the self that are involved in everyday thought is that they should underlie

³ To ensure that the relationship between an identity's causal centrality and choice was not driven by the fact that some participants reported more causal relationships between the features of their self-concept, in general, in the analyses for all studies we controlled for the number of total causal relationships reported.

common forms of decision making. That is, to understand how people commonly think about the self, we have put beliefs about the self into the context of decision making and asked: What beliefs drive people's future-oriented decisions and decisions to follow group norms? In both cases, we have found that beliefs about self-continuity, a continuous judgment about how much of the self remains over time, underlie a wide range of choices.

Self-continuity represents a new characterization of beliefs about whether identity persists over time—distinct from numerical and qualitative identity—that has not only allowed for a better understanding of decision making but also provided insight into judgments of the persistence of the self over time and the representation of the self-concept. Unlike qualitative identity, not all features of the self are equal in shaping judgments about self-continuity. Thus, a key question is why some features impact self-continuity more than others. An emerging line of research reviewed in this chapter has identified causal reasoning as a key determinant of how influential features are on self-continuity judgments and as a critical component of self-concept representation. Continuing to answer questions about why changes to some features impact self-continuity judgments more than others and how causal beliefs guide these judgments will be key issues for future research, particularly because judgments of self-continuity have important consequences for decision making.

References

- Ahn, W., N. S. Kim, M. E. Lassaline, and M. J. Dennis (2000), 'Causal status as a determinant of feature centrality', *Cognitive Psychology*, 41: 361-416.
- Akerlof, G. A., and R. E. Kranton (2000), 'Economics and Identity', *The Quarterly Journal of Economics*, 115: 715–753.
- Akerlof, G. A., and R. E. Kranton (2010), *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being*, Princeton, NJ: Princeton University Press.
- Bartels, D.M., and O. Urminsky (2011), 'On intertemporal selfishness: how the perceived instability of identity underlies impatient consumption', *Journal of Consumer Research*, 38: 182–198.
- Bartels, D. M., and O. Urminsky (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.
- Blok, S., G. Newman, and L. J. Rips (2005), 'Individuals and their concepts', in W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, and P. Wolff (eds), *Categorization Inside and Outside the Lab*, Washington, DC: American Psychological Association, 127-149.
- Chen, S. Y., and O. Urminsky (2019), 'The role of causal beliefs in political identity and voting', *Cognition*, 188: 27-38.
- Chen, S.Y., O. Urminsky, and D.M. Bartels (2016), 'Beliefs about the causal structure of the self-concept determine which changes disrupt personal identity', *Psychological Science*, 27: 1398–1406.
- Chen, S.Y., O. Urminsky, and J. Yu (2020), 'We do what we are: representations of the self-concept and identity-based consumption.'
- Damasio, H., T. Grabowski, R. Frank, A. M. Galaburda, and A. R. Damasio (1994), 'The return of Phineas Gage: clues about the brain from the skull of a famous patient', *Science*, 264(5162): 1102-1105.
- De Freitas, J. et al. (2018), 'Moral goodness is the essence of personal identity', *Trends Cognitive Science*, 22 (9): 739-740.
- Dranseika, V. (2017), 'On the Ambiguity of 'the Same Person'', *AJOB Neuroscience*, 8: 184-186.
- Ersner-Hershfield, H. et al. (2009), 'Don't stop thinking about tomorrow: Individual differences in future self-continuity account for saving', *Judgment and Decision Making*, 4: 280-286.
- Frederick, S. (2003), 'Time preference and personal identity', in G. Loewenstein, D. Read, and R. Baumeister (eds), *Time and Decision*, New York, NY, US: Russell Sage Foundation, 89-113.

- Frederick, S., N. Novemsky, J. Wang, R. Dhar, and S. Nowlis (2009), 'Opportunity cost neglect', *Journal of Consumer Research*, 36: 553-561.
- Frederick, S., G. Loewenstein, and T. O'Donoghue (2002), 'Time Discounting and Time Preference: A Critical Review', *Journal of Economic Literature*, 40 (2): 351-401.
- Gelman, S., G. Heyman, and C. Legare, 'Developmental Changes in the Coherence of Essentialist Beliefs About Psychological Characteristics', *Child Development*, 78 (3): 757-774.
- Haslam, N., B. Bastian, and M. Bissett (2004), 'Essentialist beliefs about personality and their implications', *Personality and Social Psychology Bulletin*, 30 (12): 1661-1673.
- LeBoeuf, R. A., E. Shafir, and J. B. Bayuk (2010), 'The conflicting choices of alternating selves', *Organizational Behavior and Human Decision Processes*, 111 (1): 48-61.
- Markus, H., and E. Wurf (1987), 'The dynamic self-concept: A social psychological perspective', *Annual Review of Psychology*, 38 (1): 299-337.
- Molouki, S., and D.M. Bartels (2017), 'Personal change and the continuity of the self', *Cognitive Psychology*, 93: 1-17.
- Mott, C. (2018). 'Statutes of Limitations and Personal Identity', in T. Lombrozo, J. Knobe, and S. Nichols (eds.), *Oxford Studies in Experimental Philosophy Volume 2*, Oxford: Oxford University Press.
- Parfit, D. (1984), *Reasons and Person*, Oxford, England: Oxford University Press.
- Reed II, A. (2004), 'Activating the self-importance of consumer selves: Exploring identity salience effects on judgments', *Journal of Consumer Research*, 31 (2): 286-295.
- Rehder, B., and R. 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.
- Rehder, B., and S. Kim (2010), 'Causal status and coherence in causal-based categorization', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(5): 1171-1206.
- Rosch, E., and C. B. Mervis (1975), 'Family resemblances: studies in the internal structure of categories', *Cognitive Psychology*, 7: 573-605.
- Sloman, S., B. C. Love, and W. Ahn (1998), 'Feature centrality and conceptual coherence', *Cognitive Science*, 22 (2): 189-228.
- Spiller, S. A. (2011), 'Opportunity cost consideration', *Journal of Consumer Research*, 38: 595-610.

Starmans, C. and P. Bloom (2018a), 'Nothing personal: what psychologists get wrong about identity', *Trends Cognitive Science*, 22 (7): 566-568.

Starmans, C. and P. Bloom (2018b), 'If you become evil do you die?' *Trends Cognitive Science*, 22 (9): 740-741.

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

Strohminger, N. and S. Nichols (2015), 'Neurodegeneration and identity', *Psychological Science*, 26 (9): 1469-1479.

Tajfel, H. (1978), *Differentiation between social groups: Studies in the social psychology of intergroup relations*. London: Academic Press.

Urminsky, O. (2017), 'The Role of 'Psychological Connectedness to the Future Self' in Decisions Over Time', *Current Directions in Psychological Science*, 26 (1): 34-39

Urminsky, O., & D. Bartels (2019), 'Identity, personal continuity and psychological connectedness across time and over transformation', in A. Reed II and M. Forehand (eds.), *Handbook of Research on Identity Theory in Marketing*, Northampton, MA: Edward Elgar.

Urminsky, O., and G. Zauberan (2016), 'The psychology of intertemporal preferences', in G. Keren and G. Wu (eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making*, Chichester, West Sussex: John Wiley and Sons.