



Personal change and the continuity of the self



Sarah Molouki^{*}, Daniel M. Bartels^{*}

University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637, United States

ARTICLE INFO

Article history:

Accepted 14 November 2016

Available online 28 December 2016

Keywords:

Psychological essentialism

Personal change

Identity

Self-continuity

Self-concept

Psychological connectedness

Feature centrality

ABSTRACT

Five studies explore how anticipating different types of personal change affects people's perceptions of their own self-continuity. The studies find that improvements are seen as less disruptive to personal continuity than worsening or unspecified change, although this difference varies in magnitude based on the type of feature being considered. Also, people's expectations and desires matter. For example, a negative change is highly disruptive to perceived continuity when people expect improvement and less disruptive when people expect to worsen. The finding that some types of change are consistent with perceptions of self-continuity suggests that the self-concept may include beliefs about personal development.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Studies find that people distinguish between central, immutable features of a concept and features that are allowed to change without fundamentally altering its identity (Slovan & Ahn, 1999; Slovan, Love, & Ahn, 1998). Belief in immutable features is one element of psychological essentialism: the idea that an object possesses a deep underlying character or “essence” that cannot be changed without altering the definition of the object (Gelman & Hirschfeld, 1999; Haslam, 1998; Medin & Ortony, 1989; Medin & Rips, 2005). For example, the genetic makeup of a skunk is treated as an essential feature: most people believe that the organism would shift to a different natural kind if this feature were to change (Keil, 1989). In contrast, the skunk's stripe is characteristic but non-essential, as it could change without altering basic category membership (i.e., a stripe-less skunk could still be a skunk; Keil, 1989).

Other research examines how people track the continuity of individual objects across time and transformation, and how they differentiate between changes that disrupt versus maintain the identity of an individual object (Blok, Newman, & Rips, 2005; Gutheil, Bloom, Valderrama, & Freedman, 2004; Gutheil, Gelman, Klein, Michos, & Kelaita, 2008; Gutheil & Rosengren, 1996; Rips, Blok, & Newman, 2006). This research asks questions like “how might we determine whether a *particular* skunk is the same skunk we previously observed?” This judgment is based partially on whether any changes in the object are consistent with it still being a skunk. However, it is also based on beliefs about what types of changes are consistent with the continuity of an individual object of this class over time (see, e.g., Rips, 2011). For example, did the skunk get older (consistent with our beliefs about skunks) or younger (inconsistent with beliefs about skunks) since the last time we saw it?

^{*} Corresponding authors.

E-mail addresses: smolouki@uchicago.edu (S. Molouki), bartels@uchicago.edu (D.M. Bartels).

The current research explores ideas about change and continuity in the domain of the self-concept. We examine people's beliefs about which features of themselves may change, and in what ways they may change, while still preserving the continuity of the self.¹

1.1. Stability, change, and the self-concept

What is the essence of a person? People seem not to hold a pure physicalist view of personal continuity, but rather ascribe special importance to mental characteristics (Blok, Newman, Behr, & Rips, 2001; Nichols & Bruno, 2010; Tierney, Howard, Kumar, Kvaran, & Nichols, 2014). As such, people judge that significant changes in mental content undermine the relation between a current and future self (Bartels & Rips, 2010). Also, the type of mental change matters. Perceiving changes in another person's moral qualities leads to attributions of greater personal change (Heiphetz, Strohminger, & Young, 2016; Strohminger & Nichols, 2014, 2015) and greater disruption of established social relationships (Strohminger & Nichols, 2015) than perceiving changes in other mental features, like cognitive abilities or memories.

However, rather than thinking of others' moral values and mental characteristics as uniformly immutable, people consider positive changes in these qualities to be more allowable than negative changes (Newman, Bloom, & Knobe, 2014; Newman, De Freitas, & Knobe, 2015; Tobia, 2015). This pattern of judgments could arise from beliefs that the human essence is fundamentally positive. On this view, negative change deviates from this essence whereas positive change enables its expression (Newman et al., 2014; Tobia, 2015). This view presents the idea that some aspects of who a person really is seem best expressed via a trajectory of (positive) change.

Is this also how people think about themselves? There are many differences in how people think about themselves versus others (Molouki & Pronin, 2015), underscoring the importance of specifically examining first-person judgments of self-continuity over time. We think there are at least three possibilities for how people think about their own continuity: (i) People's thoughts about their self-continuity might be similar to their thoughts about others. Recent research suggests patterns consistent with this view in at least some situations (Bench, Schlegel, Davis, & Vess, 2015; Heiphetz et al., 2016). (ii) Differences in the accessibility of mental contents could cause differences in judgments of continuity for the self versus others. People can directly access and vividly experience their own (but not others') current mental contents (see Loewenstein, O'Donoghue, & Rabin, 2003; Pronin, 2008). So, nuanced changes in mental contents that are more easily observable for the self might seem more disruptive to one's own self-continuity than similar changes would be for the continuity of others. (iii) People might expect more improvement in themselves than in others, for any number of reasons. For example, if people have a stronger motivation to see themselves (versus others) in a positive light (e.g., Kunda, 1987), then they might view specifically positive changes as more continuity-preserving for the self than for others.

At present, it is unclear from existing research how people integrate ideas about the direction of change into their own self-concept. On the one hand, people expect to undergo many positive changes throughout their lives (Busseri, Choma, & Sadava, 2009; Haslam, Bastian, Fox, & Whelan, 2007; Newby-Clark & Ross, 2003; Wilson & Ross, 2001). They also report that positive change is more consistent with a process of "self-discovery" (Bench et al., 2015), which could mean that anticipating improvement might not undermine self-continuity. On the other hand, some studies find that imagining general (i.e., neither explicitly positive nor negative) personal change to one's core psychological characteristics reduces perceived continuity between the present and future self (Bartels & Urminsky, 2011; Ersner-Hershfield, Garton, Samanez-Larkin, & Knutson, 2009). This latter finding suggests that people may in fact treat any significant change in themselves as disruptive to self-continuity, regardless of valence.

The current studies aim to inform a handful of questions that remain unresolved in the literature to date: Do people think about their own continuity similarly to how they think about the continuity of others? Do people regard some aspects as strictly immutable in themselves, or might they view all improvements as continuity-sustaining? Previous research suggests that changes to moral characteristics (Strohminger & Nichols, 2014, 2015) and negative changes (Newman et al., 2014, 2015; Tobia, 2015) in others undermine *their* continuity, but it is less clear how such changes might affect *one's own* sense of self-continuity. Also, although people expect to change for the better over time (Busseri et al., 2009; Haslam et al., 2007; Newby-Clark & Ross, 2003), they nevertheless react to the general idea of personal change as if it disrupts self-continuity (Bartels & Urminsky, 2011; Ersner-Hershfield, Garton, et al., 2009), posing somewhat of a puzzle as to how people integrate change into the self-concept. To address these questions, we examine how people think about their own continuity across time and different types of personal change, and in doing so, we provide new insight and synthesis across these various streams of theorizing.

In this paper, we explore the relationship between personal change and perceived continuity of the self by systematically varying aspects of the changes considered. As we have noted, previous research about change in different domains (i.e., tracking the continuity of objects and other people) has suggested the importance of (i) the specific nature of the feature changing and (ii) the valence of change. Also, people tend to have specific goals and thoughts about how they themselves will change over time (Brunstein, 1993; Miceli & Castelfranchi, 2015). Therefore, we suggest that expectations and desires related to future change may also be particularly relevant in the domain of the self-concept. The following studies explore

¹ In this paper we measure people's degree of relatedness between the current and future self in a graded fashion, an approach taken by related past research (e.g., Bartels & Rips, 2010; Strohminger & Nichols, 2014). To avoid confusion between the concept of graded continuity and other identity-related concepts, we will use the terms "continuity" or "self-continuity" rather than "identity" to describe the construct measured in our research.

the impact of the type of characteristic that is changing, the direction of change, and one's expectations and desires on perceptions of self-continuity.

1.2. The present studies

Five studies investigate the influence of various kinds of change on perceptions of self-continuity. In Study 1, we explore whether positive changes are seen as more consistent with one's self-concept than negative changes, and also whether this difference varies depending on the type of feature changing. In Study 2, we explore whether the perceived effect of a given change is moderated by people's expectations about the likelihood of this change. In Study 3, we examine whether desire for a particular change affects the impact of this change on self-continuity. Studies 4 and 5 replicate the findings related to expectations and desires while controlling for two possible confounds. We also include an additional condition in Studies 4 and 5 to assess whether people are answering questions about future self-continuity based on their image of the ideal self.

2. Study 1

Study 1 explores the effects of both the valence of change (positive, negative, or unspecified), as well as the type of characteristic that is changing, on perceptions of one's own self-continuity. Exploring valence allows us to assess whether people differentiate between positive and negative change when making these judgments. To examine the influence of the type of characteristic that is changing, we presented people with characteristics from five different categories (morality, personality, preferences, experiences, and memories). These categories were chosen based on previous research suggesting that they represent a range of mental characteristics of varying centrality for perceptions of individual continuity in others (Strohming & Nichols, 2014). In addition to extending the examination of valence of change and type of characteristic changing into the domain of the self-concept, Study 1 also covers new territory by testing for an interaction between these two factors.

2.1. Method

2.1.1. Participants

Three hundred participants from Amazon Mechanical Turk (MTurk) completed an online survey in return for monetary compensation. Eight participants were excluded for failing an attention check question, leaving a final N of 292 (49% male, mean age = 38.3 years, 97% native English speakers).

2.1.2. Stimuli

Participants were presented with a list of 40 personal characteristics and were asked to imagine each of these in turn changing within themselves over time. These characteristics were chosen based on pre-testing and classified into one of the following categories: "preferences", "personality", "morality", "experiences", and "memories" (see Appendix A for full list of stimuli and more information about their selection).

2.1.3. Procedure

Participants were randomly assigned to one of three experimental conditions where they imagined that these characteristics would either (i) change (valence unspecified), (ii) improve, or (iii) worsen in the future.² They then reported the perceived impact of each imagined change on their perceived self-continuity, by responding to the following statement using a sliding scale ranging from 0 to 100:

"If my (characteristic) [changes/improves/worsens], I will ..."
Really be myself [0] – Not at all be myself [100]³

Therefore, a rating of 0 meant that the stated change would cause very little disruption (or no disruption) to the participant's perceived self-continuity. In contrast, a rating of 100 meant that the change would cause a lot of disruption to the participant's self-continuity.

² A specific future time period was not provided, to allow participants to consider their own interpretation of the natural time course of change for each characteristic. However, a follow-up study ($N = 67$) found a similar pattern of results when participants were asked to consider a time period of 1 year.

³ Selection of the dependent variable was based on a pretest where participants ($N = 74$) were asked to imagine that their personality would either improve or worsen and then answered the following four questions, all designed to measure the general concept of immutability (see Sloman et al., 1998):(a) *If my personality [improves/worsens] ... I will be fundamentally the same person[0] – I will be fundamentally a different person[100]*;(b) *If my personality [improves/worsens] ... I will really be myself[0] – I will not at all be myself[100]*;(c) *How easily can you imagine yourself with [improved/worsened] personality? Not at all easy to imagine[0] – Extremely easy to imagine[100] (reverse scored)*;(d) *How surprised would you be if your personality [improved/worsened]? Not at all surprised[0] – Extremely surprised[100]*. These four items had a Cronbach's alpha of 0.72, and item-total correlations ranged from 0.255 to 0.716. Because Item B had the highest item-total correlation, we judged it to be the most representative question to assess immutability and selected it as our dependent variable for these studies.

2.2. Results

2.2.1. Effect of category on perceived self-continuity

We first tested whether change is perceived as more or less disruptive to continuity for characteristics falling into different categories (e.g., morality, preferences, etc.). To do so, we determined whether general (valence unspecified) change in some categories would be more threatening to continuity (causing participants to report that they would be less themselves) than change in other categories. Within the general change condition, a repeated measures ANOVA examining the effect of category on perceived discontinuity revealed an overall difference across the five categories of features, $F(4, 368) = 103.22$, $p < 0.001$, $\eta_p^2 = 0.53$.

Tukey pairwise contrasts revealed significant differences in the effect of change on continuity between all pairs of categories used in our experiment (all $|z| > 2.8$, $p < 0.04$). Changes in items classified in the category of morality were judged to disrupt self-continuity the most ($M = 63.12$, $SD = 33.60$), followed by items from the domains of personality ($M = 48.43$, $SD = 32.85$), preferences ($M = 37.08$, $SD = 31.00$), experiences ($M = 31.13$, $SD = 30.42$), and memories ($M = 26.19$, $SD = 31.14$). These results are generally in agreement with findings for third-party judgments, where changes in morality and personality also had the greatest impact on continuity (Strohming & Nichols, 2014). This suggests that traits within these categories are perceived to be central to one's own self-concept, much as they are to the general concept of a person.

2.2.2. Valence of change

A mixed two-way valence \times category ANOVA revealed a main effect of valence, with positive changes ($M = 27.04$, $SD = 30.14$) seen as less continuity-disrupting than negative changes ($M = 50.37$, $SD = 33.78$) or unspecified changes ($M = 43.93$, $SD = 34.43$); $F(2, 289) = 48.23$, $p < 0.001$, $\eta_p^2 = 0.25$. This perceived difference between positive and negative change held for all categories of characteristics (all $t_s > 3.35$, $p_s < 0.001$).

However, the two-way valence \times category interaction was also significant, indicating that the perceived difference between positive and negative change was larger for some categories than for others, $F(8, 1156) = 28.39$, $p < 0.001$, $\eta_p^2 = 0.16$. This effect appears to be driven by the much larger perceived discontinuity associated with negative (vs. positive) change for morality and personality, compared to a smaller effect of valence within the categories of preferences, experiences, and memories (see Fig. 1).

2.3. Discussion of Study 1

Study 1 suggests that certain categories of features, specifically those pertaining to morality and personality characteristics, are more central to the self-concept than other categories such as memories, experiences, and preferences. Paralleling results of research investigating perceptions of other individuals (Strohming & Nichols, 2014), people perceive the greatest discontinuity in themselves if they imagine moral and personality characteristics changing. This suggests that morality and personality are important for defining one's own self-concept, because altering them causes a relatively large amount of disruption (see also Heiphetz et al., 2016).

However, we also see that this relationship differs based on the valence of imagined change. Positive change has a relatively small effect on perceived continuity across all mental categories, whereas negative change seems progressively more threatening to self-continuity as characteristics become more central. This pattern is consistent with the idea that people might think of themselves as being on an upward developmental trajectory. Positive changes are thereby allowable even in one's most central characteristics, whereas negative changes are particularly threatening to continuity when they occur in core parts of the self.

The pattern of results for unspecified changes more closely mirrors that observed for negative changes. This is consistent with previous research that finds disruptions in continuity associated with general (non-valenced) descriptions of change (e.g., Bartels & Urminsky, 2011). However, it remains unclear why people would react to unspecified change much as they would to negative change if people's predominant expectations lean toward positive change (e.g., Haslam et al., 2007). One possibility is that because people may have already incorporated expectations of positive change into their self-concept, they might interpret specific mentions of personal change as implying unexpected—and thus predominantly negative—changes. However, this issue remains to be examined more closely in future research.

3. Study 2

Study 1 suggests that people see morality and personality characteristics as most central to their self-concept. In Study 2, we examine whether people think these central characteristics are less likely to change than peripheral ones. We also test whether the different reactions to positive and negative change observed in Study 1 are moderated by individual expectations of change. We predict that in addition to the influence of valence, alignment of changes with expectations may exert an independent influence on perceptions of continuity. To explore these questions, we used a 3 (expectation: improve, worsen, or stay the same) \times 3 (imagined change: improve, worsen, or stay the same) repeated measures design. First, participants

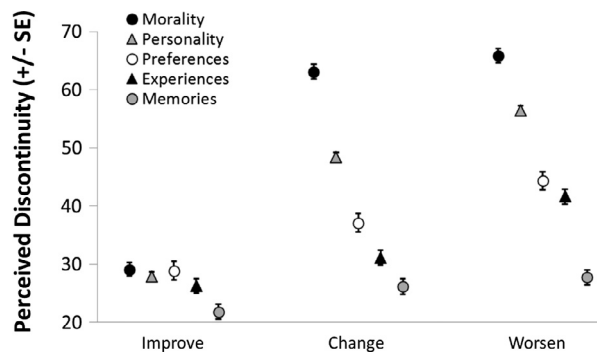


Fig. 1. Perceived self-discontinuity ratings by valence and category of change (Study 1).

reported their expectations of change for a subset of characteristics,⁴ and we then examined the perceived impact on continuity that would result from different types of change in these features.

3.1. Participants and stimuli

Three hundred participants (46% male, mean age = 36.2 years, 98% native English speakers) from Amazon Mechanical Turk (MTurk) completed an online survey in return for monetary compensation. One participant was excluded for giving incomplete responses, leaving a final N of 299. Participants were presented with the same 40 characteristics used in Study 1.

3.2. Procedure

Participants were given the list of personal characteristics and selected fifteen in total that met the following criteria for themselves: “five items that you most strongly expect will improve over time,” “five items that you most strongly expect will stay the same over time,” and “five items that you most strongly expect will worsen over time.” After selecting five characteristics for each criterion, participants ranked these characteristics by how closely they matched the criterion (i.e., for the “improve” list, 1 = the feature most expected to improve; 5 = the feature least expected to improve).

Participants were then asked to imagine that their top-ranked characteristic for each type of expectation would in turn improve, worsen, or stay the same (repeated measures, in randomized order). Participants made continuity judgments using the “self/not self” slider question from Study 1; however, in this study, the unspecified change condition was replaced with a “stay the same” condition. Participants made nine such judgments, crossing type of imagined change (improvement vs. staying the same vs. worsening) with participants’ expectation of change (expected to improve vs. expected to stay the same vs. expected to worsen).

3.3. Results

3.3.1. Category and expectations of change

A chi-square test revealed that participants systematically differed in their expectations about different categories of traits, $\chi^2(8) = 927.99, p < 0.001$ (see Table 1). Participants were more likely to state that they expected characteristics from relatively central categories to remain stable (morality, personality, preferences). In contrast, they were more likely to state that characteristics from peripheral categories were those where they expected improvement (experiences) or worsening (memories). Nevertheless, we note that as a whole the sample of participants chose a variety of characteristics from across the range of categories. This suggests that people did consider their idiosyncratic expectations rather than consulting a generalized schema of how they might expect a third party to change.

3.3.2. Expectations of change and judgments of continuity

A 3×3 repeated-measures ANOVA revealed a main effect of type of change imagined ($F(2, 596) = 293.74, p < 0.001, \eta_p^2 = 0.50$), such that imagining a feature worsening (vs. improving or staying stable) caused the greatest disruption of self-continuity (see Fig. 2). There was a nonsignificant trend for expectations of change $F(2, 596) = 2.27, Greenhouse-Geisser adjusted p = 0.11$ (unadjusted $p = 0.10$), $\eta_p^2 = 0.01$, but a significant change \times expectation interaction ($F(4, 1192) = 117.58$,

⁴ Participants in Studies 2–5 responded to imagined changes in three (in Studies 2 and 4) or two (in Studies 3 and 5) self-selected features, to better ensure that they engaged their self concept. This methodological change addresses a possible concern with Study 1, where thinking about so many different features (40 in total) could have induced people to think about changes more abstractly rather than specifically in themselves. We thank a reviewer for pointing us to this possible concern.

Table 1

Probability of stating that an item is most likely to improve, stay the same, or worsen (Study 2).

	Morality	Personality	Preferences	Experiences	Memories
Improve	26%	35%	8%	49%	10%
Same	59%	38%	60%	18%	14%
Worsen	15%	27%	32%	33%	76%
χ^2	197.3***	37.3***	55.3***	158.3***	479.8***

Note: Results listed are conditional on participant having selected a characteristic from a given category, such that percentages in vertical columns add up to 100%.

* $p < .05$; ** $p < .01$; *** $p < .001$.

Bolded values are significantly greater than 33%.

$p < 0.001$, $\eta_p^2 = 0.28$). This interaction reveals that changes consistent with expectations led to less disruption in continuity than changes that were inconsistent with expectations (see circled means in Fig. 2).

Further exploring the main effect of imagined change revealed no difference in ratings when participants imagined improvement ($M = 29.06$, $SD = 29.90$) versus staying the same ($M = 28.64$, $SD = 29.35$, paired- $t(298) < 1$). However, ratings of discontinuity for the conditions where participants imagined worsening ($M = 60.31$, $SD = 33.63$) were significantly higher than in the other two conditions ($ts(298) > 18.26$, $ps < 0.001$).

To further examine the change \times expectation interaction, for each participant, we calculated the correlation between her nine discontinuity ratings and a contrast code reflecting whether the imagined change matched or mismatched expectations. A match was coded as -1 because we expected less discontinuity in these conditions (see circled means in Fig. 2), and a mismatch was coded as $+1$ because we expected more discontinuity for these conditions. The average of these within-subjects correlations was significantly greater than zero ($M = 0.28$, one-sample $t(298) = 15.14$, $p < 0.001$), and 245 of our 299 participants expressed a positive correlation. Matched expectations mitigated the effect of an undesirable change. For example, if people already expected a trait to get worse (versus expecting stability), reported identity discontinuity related to the worsening was reduced by 45% ($M = 39.59$ vs. 72.37 , paired- $t(298) = 13.25$, $p < 0.001$). Comparisons between means of all nine conditions can be seen in Table 2.

3.4. Discussion of Study 2

Study 2 used an individualized trait-selection procedure to assess people's actual expectations about their own developmental trajectory for various types of characteristics. These expectations were then examined to see if they systematically differed based on the category of the feature. Results confirmed that people believe that personal characteristics identified as central in Study 1 (such as moral qualities) are less likely to change than other more peripheral characteristics (such as memories or experiences). Violations of these expectations may contribute to the greater relative disruption to self-continuity reported when people imagine characteristics from the categories of morality and personality changing (Study 1).

We also examined how violations of expectations contribute to differential perceptions of self-continuity associated with positive changes, stability, and negative changes. Because positive change over time tends to be consistent with overall expectations (Haslam et al., 2007; Wilson & Ross, 2001), we independently varied valence and expectation to help disentangle their influence. We crossed all types of imagined change (improve, stay the same, worsen) with all types of expectations of change (improve, stay the same, worsen). We found that although valence exerts a main effect on judgments, greater continuity is also reported when the type of change considered is consistent with expected change. Valence and expectation exert an interactive effect: positive change and stability promote greater perceived continuity than negative change, but this greater continuity is most pronounced when the change is consistent with expectations.

4. Study 3

The first two studies suggest people consider improvement to be more consistent with self-continuity than worsening, and that this difference is especially pronounced when considering characteristics that people expect to improve. In Study 3, we test whether a similar interaction emerges with respect to desires. Do people view positive changes in traits that they want to improve on as especially important for sustaining continuity? If so, this may suggest that people hold an aspirational definition of the self: who they are is partially defined by who they want to be. So, self-continuity will be threatened if their developmental trajectory is interrupted. Alternatively, people's beliefs that they will improve over time could instead be guided by broader ideas about how people in general will change over time, rather than one's own idiosyncratic desires (see, e.g., Newman et al., 2014).

Study 3 used the same design as Study 2, but with the substitution of questions about desires for change rather than expectations of change. Participants evaluated the impact of changes in traits they either desired to improve or desired to keep the same. We did not ask about traits participants desired to worsen, because we thought most participants would not have desires falling into this category. So, the study used a 2 (desire: improve or stay the same) \times 3 (imagined change: improve, stay the same, or worsen) repeated measures design.




Fig. 2. Perceived self-discontinuity ratings by imagined change and expected change (Study 2). Note: Circled values are those where imagined change matches expected change.

4.1. Participants and stimuli

Three hundred and ten participants (56% male, mean age 34.8 years, 95% native English speakers) from Amazon Mechanical Turk (MTurk) completed an online survey in return for monetary compensation.

4.2. Procedure

Participants were presented with the same 40 personal characteristics used in previous studies and were asked to select ten in total that met the following criteria: “five items that you most strongly desire to improve over time,” “five items that you most strongly desire to stay the same over time.” After selecting five characteristics for each criterion, participants ranked these items by the strength of the corresponding desire. They then evaluated the effect of the three different types of change (improve, stay the same, worsen) on perceived self-continuity for the characteristic they selected as most representative of each type of desire. As in the previous two studies, these ratings were made on a scale of 0 (*I will really be myself*) to 100 (*I will not at all be myself*).

4.3. Results

4.3.1. Centrality and desires for change

A chi-square test revealed that participants systematically differed in their desires to improve different categories of traits, $\chi^2(4) = 327.36, p < 0.001$. Participants were more likely to state that they desired to keep moral characteristics and personality the same, whereas they preferred to improve preferences and experiences and equally desired to improve or maintain memories (see Table 3). Nevertheless, as in Study 2, the traits chosen by participants were varied and spanned all categories, suggesting that participants indeed seem to be reporting their idiosyncratic desires.

4.3.2. Desire for change and judgments of continuity

A 3×2 repeated-measures ANOVA revealed a main effect of imagining improvement, stability, or worsening ($F(2,618) = 723.34, p < 0.001, \eta_p^2 = 0.70$), with imagining a feature worsening leading to greatest perceived discontinuity (see Fig. 3). We also observed a main effect of desire ($F(1,309) = 10.05, p = 0.002, \eta_p^2 = 0.03$), and a significant interaction between change and desire ($F(2,618) = 72.67, p < 0.001, \eta_p^2 = 0.19$). This interaction reveals that changes consistent with desires led to less perceived discontinuity than undesired changes (see circled means in Fig. 3).

Table 2

Means and standard deviations of self-discontinuity ratings by expected and imagined change (Study 2).

	Expectation			F-value
	Improve	Stay the same	Worsen	
<i>Imagined</i>				
Improve	18.81 _a (24.1)	25.87 _b (27.7)	42.52 _c (32.3)	67.27***
Stay the same	34.16 _d (29.8)	18.48 _a (25.0)	33.27 _d (30.2)	40.42***
Worsen	69.12 _e (29.7)	72.34 _e (29.8)	39.45 _{c, d} (31.1)	133.20***
F-value	268.00***	322.39***	7.58***	

Note: Values with different subscripts differ at $p < 0.05$ in post hoc tests controlling for multiple comparisons.

* $p < .05$; ** $p < .01$; *** $p < .001$.

Table 3
Probability of stating that one desires an item to improve or stay the same (Study 3).

	Morality	Personality	Preferences	Experiences	Memories
Improve	21%	47%	68%	71%	49%
Same	79%	53%	32%	29%	51%
χ^2	166.1***	5.69*	36.6***	118.9***	0.03

Note: Results listed are conditional on having selected a characteristic from a given category, such that percentages in vertical columns add up to 100%. **Bolded** values are significantly greater than 50%.

* $p < .05$; ** $p < .01$; *** $p < .001$.




Fig. 3. Perceived self-discontinuity ratings by type of change and desire for change (Study 3). Note: Circled values are those where imagined change matches desired change.

We further examined the effect of imagined change using paired t -tests as in Study 2. We found that people perceived somewhat less discontinuity when imagining improvement ($M = 16.36$, $SD = 22.11$) versus staying the same ($M = 23.38$, $SD = 26.57$, paired- $t(309) = 5.52$, $p < 0.001$). Ratings of perceived discontinuity for the conditions where participants imagined worsening, however, were much higher than in the other two conditions ($M = 71.00$, $SD = 27.87$, paired- $t(309) > 28.43$, $ps < 0.001$).

We also examined the interaction between change and desire much like in Study 2. We found that the average of within-subjects correlations between ratings and contrast codes ($-1 =$ consistent with desires, $+1 =$ inconsistent with desires) was significantly greater than zero ($M = 0.49$, one-sample $t(307)^5 = 35.34$, $p < 0.001$), and 294 of our 310 participants expressed a positive correlation.

The mismatch between desires and imagined change produced greater perceptions of discontinuity. This pattern was most pronounced when people imagined that desired changes would not happen. If a person imagined that a characteristic she desired to improve would instead remain stable, perceived discontinuity was increased by 123% ($M = 14.70$ vs. 32.77 , paired- $t(309) = 9.77$, $p < 0.001$). However, if a person imagined that a trait she desired to stay the same would actually improve, perceived discontinuity only increased by 29% ($M = 13.99$ vs. 18.03 , paired- $t(309) = 2.56$, $p = 0.01$). The size of these two differences was significantly different, $F(1, 618) = 37.71$, $p < 0.001$. Comparisons between means of all six conditions can be seen in Table 4.

4.4. Discussion of Study 3

Study 3 revealed that people were most likely to select morality and personality-related characteristics as those that they preferred to stay the same, and experiences and preferences as those they preferred to improve. This pattern is consistent with the findings from Study 1 that the former group of characteristics are less mutable than the latter.

Although we once again observed a main effect of valence, with worsening causing the greatest discontinuity, specific desires also influenced perceptions of self-continuity. Changes that matched desires (i.e., improvement on traits desired to improve, stability on traits desired to keep stable) were more consistent with people's views of their future selves than were changes that were contrary to desires. But the negative impact of a mismatch on continuity was much more severe when the imagined outcome was worse than desired than when the imagined outcome was better than desired. This suggests that both the general valence of change as well as specific desires for change have an impact on continuity judgments.

⁵ Two participants are omitted from this test because they gave zero variance in their ratings across the six conditions, resulting in an undefined correlation.

Table 4

Means and standard deviations of self-discontinuity ratings by valence of imagined change and desire for change (Study 3).

	Desire		F-value
	Improve	Stay the same	
<i>Imagined</i>			
Improve	14.70 _a (20.4)	18.03 _a (23.6)	5.96*
Stay the same	32.77 _b (28.2)	13.99 _a (21.1)	114.95***
Worsen	67.75 _c (27.6)	74.25 _d (27.8)	14.42***
F-value	379.03***	620.06***	

Note: Values with different subscripts differ at $p < 0.05$ in post hoc tests controlling for multiple comparisons.

* $p < .05$; ** $p < .01$; *** $p < .001$.

Changes that were congruent with specific personal desires increased perceived self-continuity, which suggests that the self-concept may have an aspirational character. People define themselves both by who they expect to be and who they want to be, and they will feel less confident in their self-continuity if this trajectory is interrupted.

5. Studies 4 and 5

Studies 1–3 suggest that people do not consider all types of change to be equally disruptive to self-continuity. Rather, our findings suggest that changes in a positive (but not negative) direction are largely consistent with the self-concept, especially when congruent with expectations and desires. We suspect that people entertain a dynamic, rather than static idea about who they are as a person, in that their self-concept encompasses ideas of personal development. Studies 4 and 5 explore additional possible influences on the patterns observed in Studies 1–3.

First, because characteristics in Studies 2 and 3 were chosen by the participants, it is possible that effects on self-continuity were influenced by confounding factors. For example, one might imagine that characteristics that participants did not expect (or desire) to improve were those that they already felt were at a high, positive level. So, maybe improvement is associated with less discontinuity than worsening because the potential improvement in a characteristic near ceiling is necessarily smaller in magnitude than potential decline. To isolate the influence of valence from any such scaling artifacts, Studies 4 and 5 repeat the procedures of Studies 2 and 3 while controlling for self-reported current levels of the characteristics. We also controlled for the category (i.e., morality, personality, preferences, experiences, memories), of each characteristic, to assess whether relationships between category and expectations or desires were driving the effects we observed in Studies 2 and 3.

Second, it is possible that participants were evaluating changes in comparison to a vision of their *ideal self*, rather than their self-concept as they realistically see it. According to self-discrepancy theory (Higgins, 1987, 1990), people have representations of their actual self-concept, but they also have aspirational self-guides. The “ideal self” is a representation of the attributes that one would ideally like to possess – that is, a person’s hopes, aspirations, or wishes (Evans & Petty, 2003; Higgins, 1987). Although achieving this ideal may or may not be realistic, it serves to guide goal-directed behavior as people pursue a long-term pattern of trying to reduce the discrepancy between their current state and the desired self.

In our studies, it is possible that mentioning personal change—and perhaps, specifically, improvements—to our participants could have activated representations of the ideal self. If participants were rating personal changes based on level of congruence with the ideal self, they would likely view improvements to be more congruent (and staying the same or worsening as less congruent) than they would if they were thinking about their (actual, more realistic) future self. (After all, staying the same or declining in some feature does not usually reduce the discrepancy between the current self and the ideal self.) To assess whether people in our “future self” conditions are instead referencing the “ideal self”, half of the participants in Studies 4 and 5 are explicitly asked to make judgments in reference to their ideal self. If the responses of these participants differ from responses of those who answer our original self-continuity questions, we will infer that people differentiate between representations of the actual future self and the ideal self.

5.1. Study 4

5.1.1. Participants and stimuli

Four hundred participants (53% male, mean age = 32.8 years, 99% native English speakers) from Amazon Mechanical Turk (MTurk) completed an online survey in return for monetary compensation. One participant was excluded for providing incomplete data, leaving a final N of 399.

5.1.2. Procedure

Participants completed the same procedure as in Study 2, however, those in the “ideal self” condition responded to the changes presented using a scale of 0 (“I will be exactly the person I would ideally like to be”) to 100 (“I will be not at all the person I would ideally like to be”). Also, at the end of the study, participants specified their perceived current level of each charac-

teristic considered (on a scale ranging from “extremely negative, it could not possibly get any worse” to “extremely positive, it could not possibly get any better”).

5.1.3. Results

5.1.3.1. Perceived discontinuity with actual versus ideal self-concept. We conducted a 3 (imagined change: improve, stay the same, worsen) \times 3 (expectation: improve, stay the same, worsen) \times 2 (question type: actual self, ideal self) mixed ANOVA, with the first two factors as repeated measures and the latter between-subjects. A main effect of question type was found, $F(1, 397) = 34.24$, $p < 0.001$, $\eta_p^2 = 0.08$, such that people gave higher overall discontinuity ratings when answering questions about continuity with the ideal self rather than the actual self. More importantly, however, this was qualified by (i) a significant question type \times imagined change interaction, $F(2, 794) = 40.21$, $p < 0.001$, $\eta_p^2 = 0.09$, (ii) a question type \times expectation interaction, $F(2, 794) = 11.83$, $p < 0.001$, $\eta_p^2 = 0.03$, and (iii) a 3-way question type \times imagined change \times expectation interaction, $F(4, 1588) = 2.77$, *Greenhouse-Geisser adjusted* $p = 0.04$ (*unadjusted* $p = 0.03$), $\eta_p^2 = 0.01$. These differences by question type (ideal vs. actual)—both the main effect as well as the fact that question type interacts with every other term in the ANOVA—reveal that when people are responding to questions about the future self, they are providing different answers than those who are asked to think about ideal self (see Fig. 4). In particular, the question type \times imagined change interaction reflects the fact that those in the ideal self condition reported greater discontinuity than those in the actual self condition when imagining stability or worsening. However, participants answering questions about the ideal self reported less discontinuity when imagining improvement than those considering the actual self (see Table 5). This pattern of results suggests that people are differentiating between the ideal self and the (more realistic) actual future self.

5.1.3.2. Expectations, valence, and judgments of continuity. A 3 \times 3 repeated-measures ANOVA including only those participants in the actual self condition revealed a similar pattern of results to those in Study 2 (see Fig. 5). We reanalyzed our results using a mixed-effects model to ensure that our findings are robust when accounting for variation in responses due to the current level (extremely positive to extremely negative) or category (morality, personality, preference, experience, memory) of the relevant feature. We found that current levels varied by expectations of change. Features that were expected to remain stable were reported to be at higher current levels ($M = 73.89$, $SD = 22.48$) than those that were expected either to improve ($M = 60.75$, $SD = 23.09$) or to worsen ($M = 48.82$, $SD = 25.32$), $t(596) > 11.08$, $ps < 0.001$. Expectations were also associated with category, $\chi^2(8) = 1207.18$, $p < 0.001$. As in Study 2, participants more frequently expected moral characteristics to stay the same and less central categories such as experiences and memories to either improve or worsen (see Table 6). However, these associations do not account for our findings: The interaction between valence and expectation remains significant when controlling for both level and category in our model; $F(4, 1579) = 57.98$, $p < 0.001$. So, the observed effects of expectation, here and in Study 2, do not seem to be the product of extraneous factors related to one's current level of a characteristic or the category from which it was drawn.

5.2. Study 5

5.2.1. Participants and stimuli

Three hundred and sixty-one participants (55% male, mean age = 33.2 years, 96% native English speakers) from Amazon Mechanical Turk (MTurk) completed an online survey in return for monetary compensation.

5.2.2. Procedure

Participants completed the same procedure as in Study 3, however, those in the “ideal self” condition responded to the changes presented using a scale of 0 (“I will be exactly the person I would ideally like to be”) to 100 (“I will be not at all the person




Fig. 4. Differences between judgments of continuity with ideal versus actual self plotted by imagined change and expectation of change (Study 4).

Table 5
Discontinuity ratings for actual self vs. ideal self based on valence of imagined change (Study 4).

Imagined change	Actual self	Ideal self	F value
Improve	25.94 _a (29.83)	20.06 _b (25.81)	11.71***
Stay the Same	20.63 _b (26.14)	36.57 _d (30.68)	81.92***
Worsen	59.62 _c (35.04)	71.96 _e (29.63)	50.82***

Note: Values with different subscripts differ at $p < 0.05$ in post hoc tests controlling for multiple comparisons.

* $p < .05$; ** $p < .01$; *** $p < .001$.



Fig. 5. Perceived self-discontinuity ratings by imagined change and expected change (Study 4). Note: Figure displays marginal means when controlling for current level and category of each characteristic. Circled values are those where imagined change matches expected change.

Table 6
Probability of stating that an item is most likely to improve, stay the same, or worsen (Study 4).

	Morality	Personality	Preferences	Experiences	Memories
Improve	23%	37%	51%	44%	8%
Same	58%	36%	30%	21%	14%
Worsen	19%	27%	19%	35%	78%
χ^2	264.8***	48.0***	83.7***	82.3***	728.3***

Note: Results listed are conditional on having selected a characteristic from a given category, such that percentages in vertical columns add up to 100%.

Bolded values are significantly greater than 33%.

* $p < .05$; ** $p < .01$; *** $p < .001$.

I would ideally like to be). Also, at the end of the study, participants specified their perceived current level of each characteristic considered (on a scale ranging from “*extremely negative, it could not possibly get any worse*” to “*extremely positive, it could not possibly get any better*”).

5.2.3. Results

5.2.3.1. Perceived discontinuity with actual versus ideal self-concept. Similar to Study 4, we conducted a 3 (imagined change: improvement, no change, worsening) \times 2 (desire: improvement vs. same) \times 2 (question type: actual vs. ideal self) mixed ANOVA, with the first two factors as repeated measures and the latter between-subjects. Again, as in Study 4, we found a main effect of question type (people’s discontinuity responses were higher for the ideal self). But, we also found interactions between question type and (i) the effects of imagined change, (ii) the effects of desired change, and the (iii) interaction of these two factors ($F_s \geq 15.05$, $p_s < 0.001$, $\eta^2_p s \geq 0.04$, see Fig. 6). This suggests that when people are responding to questions about the actual self-concept, they are providing different answers than those who are contemplating the ideal self.

5.2.3.2. Desires, valence, and judgments of continuity. A 3 \times 2 repeated-measures ANOVA including only those participants in the actual self condition revealed a similar pattern of results to those in Study 3 (see Fig. 7). We ran a mixed-effects model similar to that in Study 4 to account for any variance in responses due to the current level or category of the feature being judged. Characteristics people desired to improve were currently at a lower level ($M = 51.73$, $SD = 24.40$) than characteristics they desired to keep the same ($M = 76.32$, $SD = 19.20$), paired- $t(482) = 18.49$, $p < 0.001$. Participants also expressed differences in their desires to improve characteristics based on category ($\chi^2(4) = 386.9$, $p < 0.001$), stating that they desired more central categories to remain the same and less central categories to improve (see Table 7). However, a significant interaction

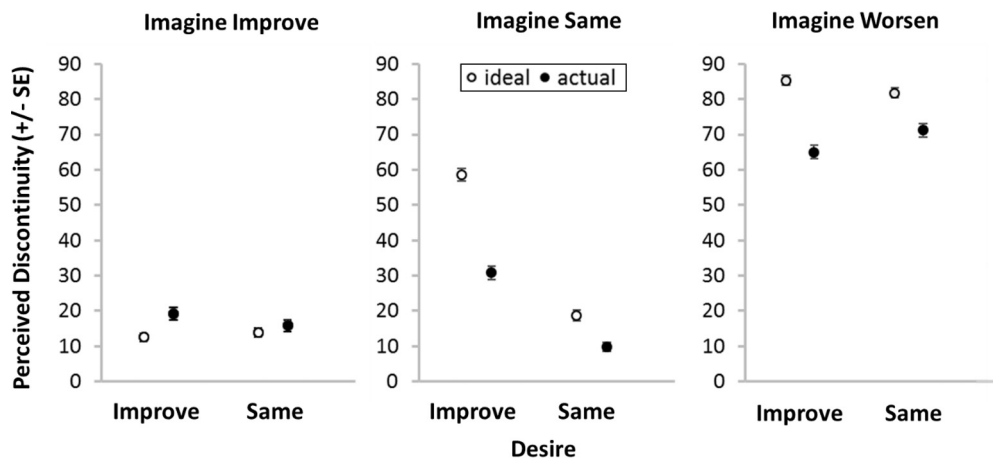


Fig. 6. Differences between judgments of continuity with ideal versus actual self plotted by imagined change and desire for change (Study 5).




Fig. 7. Perceived self-discontinuity ratings by imagined change and desired change (Study 5). Note: Figure displays marginal means when controlling for current level and category of each characteristic. Circled values are those where imagined change matches desired change.

between valence and desire emerged even when controlling for both level and category in our model, $F(2, 792) = 24.55$, $p < 0.001$. So, the observed effects of desired change, here and in Study 3, are not the product of extraneous factors related to one's current level on a characteristic or the category from which it was drawn.

5.3. Discussion of studies 4 and 5

Studies 4 and 5 found the interactive effects of valence and expectation/desire on perceived self-continuity observed in Studies 2 and 3 while controlling for the possible influences of two additional factors. First, we wanted to examine whether people were thinking about their ideal selves, rather than their future selves, when answering the continuity questions in previous studies. We found that those who made judgments about the ideal self gave significantly different answers than those who thought about the actual self-concept.

People acknowledge that improvements lead to somewhat more continuity with an image of the ideal self than with their more realistic views of the future self. In contrast, worsening or remaining the same leads to greater perceived continuity when considering the actual self rather than the ideal self. So, Studies 1–3 appear to be tapping into perceptions of continuity with the future self that are distinct from hopes for the ideal self. However, Figs. 4 and 6 reveal general similarities in the response patterns when considering the ideal versus actual self. This suggests that although people do seem to distinguish the concepts of the future self and the ideal self in their minds, people's conception of the future self may include some characteristics of the hoped-for self.

We also addressed the possibility that people may endorse improvements more than worsening because of a ceiling effect restricting the possible magnitude of improvement in one's current characteristics. However, we confirmed that the observed effects remained strong even when controlling for variations in perceived levels, as well as when controlling for the category of the characteristic in question. Both an overall preference for improvements, as well as a match between expectations/desires and change, play roles in maintaining perceived continuity over time.

Table 7
Probability of stating that one desires an item to improve or stay the same (Study 5).

	Morality	Personality	Preferences	Experiences	Memories
Improve	22%	48%	65%	73%	38%
Same	78%	52%	35%	27%	62%
χ^2	181.7***	3.73*	30.1***	161.5***	9.88**

Note: Results listed are conditional on having selected a characteristic from a given category, such that percentages in vertical columns add up to 100%.

Bolded values are significantly greater than 50%.

* $p < .05$; ** $p < .01$; *** $p < .001$.

6. General discussion

Every person has a unique self-concept, but these studies uncover general similarities across participants in the perceived effects of different types of personal change on self-continuity. Study 1 found that people differentiate between categories of mental change when thinking about their own self-concept over time, perceiving changes to morality and personality to be the most disruptive to self-continuity. However, these differences were attenuated when the changes were improvements, which people treated as relatively congruent with a continuous sense of self regardless of the type of characteristic changing. There was a large main effect of the valence of change across all studies, with positive changes being less disruptive to continuity than negative changes. However, Studies 2–5 found that specific expectations and desires also matter, with changes that are in line with expectations and desires causing less perceived discontinuity than changes that are incongruent with expectations and desires. Studies 4–5 also suggest that although people seem to incorporate improvement into their self-concept, they nevertheless distinguish the future self from their vision of the ideal self.

Previous research suggests that moral traits are seen as central to the identity of a person. People report that another individual would still remain mostly the same after changes to her preferences or memories, but would be fundamentally altered after experiencing changes to moral traits (and, to a somewhat lesser extent, personality traits; Strohminger & Nichols, 2014). Thus, maintaining moral characteristics has been considered as essential in defining and preserving a person's continuity over time. However, it was not clear whether similar views would apply to perceptions of one's own continuity. For example, as Strohminger and Nichols suggest, although people may be particularly attuned to the moral character of another person to help guide interactions with her, a different set of characteristics might come to the fore when introspecting about oneself.

In Study 1, we found that changes in moral traits indeed led to the greatest perceived discontinuity in the self. Changes in personality also led to relatively large perceptions of discontinuity, whereas changes in experiences, preferences, and memory were relatively less disruptive to continuity. This suggests that people treat morality, and to a somewhat lesser extent, personality, as defining features of themselves, much as they do for others.

Beyond the overall centrality of morality and personality, we also find an important interaction between type of characteristic and valence of change. This interaction suggests that people do not treat even their central characteristics as strictly immutable. Negative changes and unvalenced changes were more disruptive for features relating to morality and personality than for other features. But imagined improvement led to relatively low levels of perceived discontinuity regardless of the type of characteristic that changed. That is, characterization as positive change seems to mitigate perceptions of discontinuity stemming from the type of characteristic changing. Rather than considering even their core characteristics (such as morality) to be fixed, people seem to be receptive to positive, but not negative, changes, viewing them as compatible with their self-concept.

How do people reconcile these positive changes with their view of a stable, continuous self? One possibility is that people incorporate a trajectory of improvement into their self-concept. Although a person's future self may seem somewhat different from who she is now, ideas of personal development may be central to her personal identity. As a result, she would perceive discontinuity if this trajectory were to be interrupted. For example, an entering college freshman may consider the prospect of intellectual development to be a fundamental aspect of her identity as a scholar. She might perceive that she would be "more herself" after an increase in intellectual ability. In contrast, she would be "less herself" if her intellectual ability were to remain the same over the next four years. This judgment would occur despite the fact that in the latter scenario she would actually be more similar (i.e., would experience less change) across the two time points. This view is consistent with Peetz and Wilson's (2008) concept of the "temporally extended self" which claims that aspects of the future self are incorporated into one's overall self-concept. And our Studies 4 and 5 suggest that, although people distinguish between their views of the (more realistic) future self and the ideal self, the general pattern of reported discontinuity in response to various changes is similar for these two kinds of judgments.

Our basic finding that self-improvement is more consistent with the self-concept than either worsening or general change helps to reconcile two streams of previous research. We began this paper wondering how people can believe that they will improve over time (Busseri et al., 2009; Haslam et al., 2007; Newby-Clark & Ross, 2003; Wilson & Ross, 2001) while simultaneously reacting to the prospect of fundamental change by feeling disconnected from their future self (Bartels & Urminsky, 2011). After interpreting the results of our studies, we now have a tentative answer: People do not treat improvement as much of a change at all. Study 1's results suggest that participants instructed to think about change (without a specified

valence) made continuity judgments that were similar to those of participants instructed to imagine worsening. In contrast, imagining improvement seems to be a special case (not brought to mind when generally thinking about “change”) that is not perceived as disruptive. We suspect the distinction between improvement and the broader idea of change allows people to simultaneously hold positive expectations for their future development while still maintaining a stable sense of self over time.

However, although the valence of imagined change exerts a strong influence on perceived continuity, expectations and desires also play a role. Changes that match expectations or desires produce less perceived disruption. This interactive effect suggests that people’s endorsement of self-improvement is tempered by reality constraints. Although these influences usually work in tandem, with expectations and desires being aligned toward improvement, sometimes they do not align. For example, although most of us do not look forward to the prospect that our physical and mental characteristics will worsen with age, we all expect this to occur as a part of the normal human lifespan. In this case, an improvement (e.g., looking younger with age, barring external interventions) would likely cause greater disruption to self-continuity than worsening (e.g., looking older with age), because of the fact that the former type of change violates deeply-held expectations.

6.1. Are people in our studies using an essentialist view of the self when making continuity judgments?

It seems reasonable to expect that people will sometimes think about the existence and stability of certain essential features to identify an object, or maybe even themselves. For example, Haslam, Bastian, and Bissett (2004) state that an individual’s essentialized characteristics are “immutable, informative, discrete, and biologically based entities that inhere with the person” (p. 1662), and Gelman (2003) describes the essence of a living creature as “whatever quality is thought to remain unchanging as an organism grows, reproduces, and undergoes morphological transformations” (p. 8). Are participants in our studies appealing to a similar essentialist view of the self? In Study 1, we find that people consider changes in morality and personality to be quite disruptive to their self-continuity, which is consistent with previous research that identifies these characteristics as central to the identity of an individual (Strohlinger & Nichols, 2014). Nevertheless, we do not find evidence that people consider any class of the personal features we tested to be strictly immutable. Overall, people find improvements in almost every characteristic to be quite compatible with their conception of who they are as a person. So, although people perceive morality and personality as quite important to their self-concept (i.e., generally less mutable than more peripheral characteristics), they do not seem to view these features as strictly immutable.

Newman and Knobe (2016) distinguish between two kinds of essentialist theories—one kind (which they call “Lockean”) is about immutable properties. The other kind (which they call “Platonic”) defines certain concepts not by immutable properties, but by the fact that their properties embody a higher order value (Newman & Knobe, 2016; see also Knobe, Prasada, & Newman, 2013). In other words, this type of essence is not characterized directly by the presence of concrete features, but rather “in terms of the abstract values that these features seem to realize” (Knobe et al., 2013, p.243). In accordance with this latter form of essentialism, the essence of the self might be defined by the value of being fundamentally good (see Newman et al., 2014). In this case, positive changes would not disrupt self-continuity, since they would be consistent with the expression of a person’s fundamentally good essence (Bench et al., 2015; Newman et al., 2014, 2015).

Are participants in our studies using the kind of Platonic essentialist theorizing described by Newman and Knobe (2016)? On one hand, we find that people generally view their own positive changes as less disruptive than negative changes. However, we nonetheless observed situations where people do not view improvement as most compatible with their own self-continuity. For example, when people *expect* a characteristic to worsen, this negative change is just as compatible with their self-concept as improvement or stability (Studies 2 and 4). Overall, results from Studies 2–5 suggest that people’s views about how they will develop in the future are based on idiosyncratic expectations and desires, which are often, but not always, consistent with improvement.

In summary, we find that people do not tend to treat any of the features in our studies as strictly immutable, and not all improvements are seen as sustaining continuity of the self. So, essentialist views are not required to explain our pattern of results, despite their clear importance for understanding identity judgments more generally. Instead, we find that people seem to judge self-continuity by comparing change with an individually-defined trajectory, which calls to mind the causal continuer model of identity (Blok, Newman, & Rips, 2007; Blok et al., 2005; Nozick, 1981; Rips, 2011; Rips & Hespos, 2015; Rips et al., 2006; Sagi & Rips, 2014; see also Rips, 2001). In this framework, two objects are judged to be identical if there is a coherent causal history linking them, which will differ based on the object being judged. For example, an iceberg can only be the same iceberg at two time points if it has shrunken (but not grown over time), yet a human child can be the same child if it has grown (but not shrunken). Similarly, the self-concept may incorporate individual beliefs about features and the causal relations between them, rather than being defined by a single essence (Rips et al., 2006). In line with this idea, recent research suggests that people think changes in causally central personal features (i.e., those that share many causal connections with other features) are more disruptive to self-continuity than changes in more peripheral features (Chen, Urminsky, & Bartels, 2016).

6.2. Conclusions and implications

Overall, we find that people’s sense of their own self-concept tends to be the most disrupted by negative change, especially to moral or personality characteristics. Nevertheless, this pattern may vary based on one’s own individual expectations

and desires. Rather than being strictly defined by a specific set of features, our sense of our own self-continuity over time seems to be maintained to the extent that we follow an expected trajectory of development, which is usually (though not exclusively) defined by positive change.

Identifying the perceived impact of personal change on self-continuity has implications for other thoughts and behaviors. Previous research has already found that expectations of change between the current and future self has been linked to unethical acts (Hershfield, Cohen, & Thompson, 2012), increased splurging with discretionary funds (Bartels & Urminsky, 2015), lower asset accrual (Ersner-Hershfield, Wimmer, & Knutson, 2009), and reduced motivation (Peetz, Wilson, & Strahan, 2009), but may at times lead to positive outcomes such as increased giving to others (Bartels, Kvaran, & Nichols, 2013). However, our results suggest that describing specifically positive changes—because they cause minimal disruption to self-continuity—might not lead to imprudent behaviors or behaviors that harm the future self. Continuing to explore how people think of themselves changing over time can not only help deepen our understanding of basic questions about self-perception, but may also help us to better understand the determinants of self-relevant behaviors across different contexts.

Human subjects statement

The authors assert that all work with human subjects described in this manuscript was carried out in accordance with the Code of Ethics, and that informed consent was obtained from each participant.

Acknowledgements

We thank Halley Bayer, Stephanie Chen, Hal Hershfield, Trevor Kvaran, Art Markman, George Newman, Lance Rips and his lab, Nina Strohminger, Kevin Tobia, Oleg Urminsky, and an anonymous reviewer for very helpful feedback on this project. Financial support from the John Templeton Foundation is gratefully acknowledged.

Appendix A. Stimuli used in studies 1–5

The characteristics used in this study were selected based on a pre-test where we elicited descriptors relevant to personal identity from a separate sample of 35 MTurk participants. Each participant was asked to list “five or more aspects of yourself that are important to you,” “five or more aspects of yourself that distinguish you from others,” “five or more aspects of yourself that define how others see you,” “five or more aspects of yourself that you don’t want to see change,” and “five or more aspects of yourself that you do want to see change.” These questions were chosen in order to elicit wide coverage of characteristics that people deem important to their own identity. The last question (aspects of yourself that you do want to see change) was included in order to ensure that some negative characteristics were included.

The entire list of characteristics was compiled and similar characteristics were combined into a single concept (for example, specific mentions of school or career-related items were grouped under the characteristic “professional goals”). We then selected the most frequently mentioned items and presented this list to a new sample of 70 MTurk participants for classification. These participants placed each characteristic into one of the following categories: “preferences”, “personality”, “morality”, “experiences”, and “memories” (see Strohminger & Nichols, 2014). Characteristics where fewer than half of the sample agreed on category membership were not included in the current studies. We also included some characteristics used in previous research (Strohminger & Nichols, 2014), which are indicated in the table. The resulting list of 40 characteristics, with their corresponding category classifications, are presented below.

Characteristic	Category	Characteristic	Category	Characteristic	Category
Morality	Morality	Confidence Level	Personality	Life Experiences	Experiences
Values	Morality	Imagination	Personality	Friendships	Experiences
Degree of Honesty	Morality	Reliability	Personality	Everyday Activities	Experiences
Beliefs	Morality	Degree of Independence	Personality	Health	Experiences
Level of Humility	Morality	Level of Friendliness	Personality	Ability to Feel Pain*	Experiences
Religious or Spiritual Faith	Morality	Self-Awareness	Personality	Occupation	Experiences
Level of Selflessness	Morality	Level of Happiness	Personality	Cherished Memories of Time Spent with Family	Memories
Personality	Personality	Intelligence Level	Personality	Memories of Time Spent Commuting to Work*	Memories
Sense of Humor	Personality	Level of Disorganization	Personality	Knowledge of How to Ride a Bike*	Memories

(continued on next page)

Stimuli used in studies 1–5 (continued)

Characteristic	Category	Characteristic	Category	Characteristic	Category
Impulsiveness	Personality	Helpfulness	Personality	Knowledge of How to Play the Piano*	Memories
Temperament	Personality	Professional Goals	Preferences	Knowledge of Math*	Memories
Curiosity	Personality	Preferences/ Favorite Things	Preferences	Bad Memories	Memories
Dominance	Personality	Major Likes and Dislikes	Preferences		
Level of Calmness	Personality	Goals for your personal life	Preferences		

* Note: indicates that this characteristic was drawn from [Strohmingner and Nichols \(2014\)](#) rather than provided during pretesting.

References

- Bartels, D. M., Kvaran, T., & Nichols, S. (2013). Selfless giving. *Cognition*, *129*(2), 392–403.
- Bartels, D. M., & Rips, L. J. (2010). Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, *139*(1), 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*(1), 182–198.
- Bartels, D. M., & Urminsky, O. (2015). To know and to care: How awareness and valuation of future jointly shape consumer spending. *Journal of Consumer Research*, *41*(6), 1469–1485.
- Bench, S. W., Schlegel, R. J., Davis, W. E., & Vess, M. (2015). Thinking about change in the self and others: The role of self-discovery metaphors and the true self. *Social Cognition*, *33*(3), 169–185.
- Blok, S., Newman, G., Behr, J., & Rips, L. J. (2001). Inferences about personal identity. In *Proceedings of the twenty-third annual conference of the cognitive science society* (pp. 80–85).
- Blok, S., Newman, G., & Rips, L. J. (2007). Out of sorts? Some remedies for theories of object concepts: A reply to Rhemtulla and Xu (2007). *Psychological Review*, *114*(4), 1096–1102.
- Blok, S., Newman, G., & Rips, L. J. (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 lab* (pp. 127–149). American Psychological Association.
- Brunstein, J. C. (1993). Personal goals and subjective well-being: A longitudinal study. *Journal of Personality and Social Psychology*, *65*(5), 1061–1070.
- Busseri, M. A., Choma, B. L., & Sadava, S. W. (2009). Functional or fantasy? Examining the implications of subjective temporal perspective “Trajectories” for life satisfaction. *Personality and Social Psychology Bulletin*, *35*(3), 295–308.
- Chen, S. Y., Urminsky, O., & Bartels, D. M. (2016). Beliefs about the causal structure of the self-concept determine which changes disrupt personal identity. *Psychological Science*, *27*(10), 1398–1406.
- Ersner-Hershfield, H., Garton, M. T., Samanez-Larkin, G. R., & Knutson, B. (2009). Don't stop thinking about tomorrow: Individual differences in future self-continuity account for saving. *Judgment and Decision Making*, *4*(4), 280–286.
- Ersner-Hershfield, H., Wimmer, G. E., & Knutson, B. (2009). Saving for the future self: Neural measures of future self-continuity predict temporal discounting. *Social Cognitive and Affective Neuroscience*, *4*(1), 85–92.
- Evans, L. M., & Petty, R. E. (2003). Self-guide framing and persuasion: Responsibly increasing message processing to ideal levels. *Personality and Social Psychology Bulletin*, *29*(3), 313–324.
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford University Press.
- Gelman, S. A., & Hirschfeld, L. A. (1999). How biological is essentialism. *Folkbiology*, 403–446.
- Gutheil, G., Bloom, P., Valderrama, N., & Freedman, R. (2004). The role of historical intuitions in children's and adults' naming of artifacts. *Cognition*, *91*(1), 23–42.
- Gutheil, G., Gelman, S. A., Klein, E., Michos, K., & Kelaita, K. (2008). Preschoolers' use of spatiotemporal history, appearance, and proper name in determining individual identity. *Cognition*, *107*(1), 366–380.
- Gutheil, G., & Rosengren, K. S. (1996). A rose by any other name: Preschoolers' understanding of individual identity across name and appearance changes. *British Journal of Developmental Psychology*, *14*(4), 477–498.
- Haslam, N. (1998). Natural kinds, human kinds, and essentialism. *Social Research*, *65*(2), 291–314.
- Haslam, N., Bastian, B., & Bissett, M. (2004). Essentialist beliefs about personality and their implications. *Personality and Social Psychology Bulletin*, *30*(12), 1661–1673.
- Haslam, N., Bastian, B., Fox, C., & Whelan, J. (2007). Beliefs about personality change and continuity. *Personality and Individual Differences*, *42*(8), 1621–1631.
- Heiphetz, L., Strohmingner, N., & Young, L. L. (2016). The role of moral beliefs, memories, and preferences in representations of identity. *Cognitive Science*, *10*.
- Hershfield, H. E., Cohen, T. R., & Thompson, L. (2012). Short horizons and tempting situations: Lack of continuity to our future selves leads to unethical decision making and behavior. *Organizational Behavior and Human Decision Processes*, *117*(2), 298–310.
- Higgins, E. T. (1987). Self-discrepancy: A theory relating self and affect. *Psychological Review*, *94*(3), 319–340.
- Higgins, E. T. (1990). Self-state representations: Patterns of interconnected beliefs with specific holistic meanings and importance. *Bulletin of the Psychonomic Society*, *28*(3), 248–253.
- Keil, F. C. (1989). *Concepts, kinds and cognitive development*. Cambridge, MA: Bradford Books/MIT Press.
- Knobe, J., Prasada, S., & Newman, G. E. (2013). Dual character concepts and the normative dimension of conceptual representation. *Cognition*, *127*(2), 242–257.
- Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology*, *53*(4), 636–647.
- Loewenstein, G., O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *The Quarterly Journal of Economics*, *118*(4), 1209–1248.
- Medin, D. L., & Ortony, A. (1989). Psychological essentialism. *Similarity and Analogical Reasoning*, 179–195.
- Medin, D. L., & Rips, L. J. (2005). Concepts and categories: Memory, meaning, and metaphysics. In K. J. Holyoak & R. G. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 37–72).
- Miceli, M., & Castelfranchi, C. (2015). *Expectancy and emotion*. New York: Oxford University Press.
- Molouki, S., & Pronin, E. (2015). Self and other. In M. Mikulincer & P. R. Shaver (Eds.), *APA Handbook of Personality and Social Psychology* (Vol. 1, pp. 387–414). Washington, DC: APA.
- Newby-Clark, I. R., & Ross, M. (2003). Conceiving the past and future. *Personality and Social Psychology Bulletin*, *29*(7), 807–818.
- Newman, G. E., & Knobe, J. (2016). The essence of essentialism. Unpublished manuscript.

- Newman, G. E., Bloom, P., & Knobe, J. (2014). Value judgments and the true self. *Personality and Social Psychology Bulletin*, 40(2), 203–216.
- Newman, G. E., De Freitas, J., & Knobe, J. (2015). Beliefs about the true self explain asymmetries based on moral judgment. *Cognitive Science*, 39(1), 96–125.
- Nichols, S., & Bruno, M. (2010). Intuition about personal identity: An empirical study. *Philosophical Psychology*, 23(3), 293–312.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- Peetz, J., & Wilson, A. E. (2008). The temporally extended self: The relation of past and future selves to current identity, motivation, and goal pursuit. *Social and Personality Psychology Compass*, 2(6), 2090–2106.
- Peetz, J., Wilson, A. E., & Strahan, E. J. (2009). So far away: The role of subjective temporal distance to future goals in motivation and behavior. *Social Cognition*, 27(4), 475–495.
- Pronin, E. (2008). How we see ourselves and how we see others. *Science*, 320(5880), 1177–1180.
- Rips, L. J. (2001). Necessity and natural categories. *Psychological Bulletin*, 127(6), 827.
- Rips, L. J. (2011). Split identity: Intransitive judgments of the identity of objects. *Cognition*, 119(3), 356–373.
- Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review*, 113(1), 1–30.
- Rips, L. J., & Hespos, S. J. (2015). Divisions of the physical world: Concepts of objects and substances. *Psychological Bulletin*, 141(4), 786–811.
- Sagi, E., & Rips, L. J. (2014). Identity, causality, and pronoun ambiguity. *Topics in Cognitive Science*, 6(4), 663–680.
- Slooman, S. A., & Ahn, W. K. (1999). Feature centrality: Naming versus imagining. *Memory & Cognition*, 27(3), 526–537.
- Slooman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 22(2), 189–228.
- Strohlinger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1), 159–171.
- Strohlinger, N., & Nichols, S. (2015). Neurodegeneration and identity. *Psychological Science*, 26(9), 1469–1479.
- Tierney, H., Howard, C., Kumar, V., Kvaran, T., & Nichols, S. (2014). How many of us are there? In J. Sytsma (Ed.), *Advances in experimental philosophy of mind* (pp. 181–202). Bloomsbury Press.
- Tobia, K. P. (2015). Personal identity and the Phineas Gage effect. *Analysis*, 75(3), 396–405.
- Wilson, A. E., & Ross, M. (2001). From chump to champ: People's appraisals of their earlier and present selves. *Journal of Personality and Social Psychology*, 80(4), 572–584.