



Personal Change and the Continuity of Identity

Sarah Molouki
Daniel M. Bartels
University of Chicago



In this project we examine how thinking about different types of personal change affects people's perceptions of their own identity.



Before I start talking about self-perceptions, let's think a bit about how change might affect the identity continuity of another person.

I'm sure you all have heard many times about the story of Phineas Gage, a railroad worker who survived a severe injury in the 19th century when a rod flew through his skull and brain. Although there is some debate about the true effects of his injury, at least one account by his physician contains the following statements:

Dr. Harlow wrote: "Previous to his injury...he possessed a well-balanced mind." But after his injury, Harlow stated: "He is fitful, irreverent...manifesting but little deference for his fellows, impatient of restraint or advice when it conflicts with his desires. Harlow concludes by stating that "his mind was radically changed, so decidedly that his friends and acquaintances said he was 'no longer Gage.'"



- Why did these particular changes cause Gage to be regarded as “no longer Gage” after the accident?
- Not all changes are created equal
- Psychological essentialism (Medin & Ortony, 1989; Keil, 1989)
Change in features that define the “essence” of an object disrupt its identity

We may ask why the particular changes mentioned (irreverence, impatience, and lack of social regard) caused people to think of Gage as a different person after the accident.

For example, if Gage had lost the use of his legs, or perhaps even lost his memory, people may not have drawn the same conclusion. This example illustrates that when thinking about identity continuity, not all changes are created equal.

The theory of psychological essentialism suggests that changes to certain features that define the essence of an object are those that would disrupt its identity.

Important elements of change for the identity of (other) individuals



1) Type of characteristic

- **Mental vs. physical qualities** (Blok, Newman, Behr, & Rips, 2001; Nichols & Bruno, 2010)
- **Moral characteristics** (Strohminger & Nichols, 2014)

2) Valence of change (Tobia, 2015; Newman, Bloom, & Knobe, 2014)

- **Declines more disruptive than improvements**
- E.g., Phineas “becoming cruel” vs. “becoming kind”

So what types of changes might disrupt or not disrupt the identity continuity of other people as well as ourselves?

Previous research gives us some idea about how we view changes in other people.

First, the type of characteristic that is changing matters. For example, mental characteristics, and specifically moral characteristics, are seen as more central to the concept of a person.

Thus, changes in these characteristics, such the socio-moral changes seen in Gage, seem to be the most identity-disrupting.

Furthermore, the valence of change, and not just its magnitude, matters. **NEGATIVE** changes of all sorts appear to be more damaging to personal continuity than positive changes. Kevin Tobia has shown in his recent paper that presenting a hypothetical scenario where Phineas Gage goes from cruel to kind after his accident caused people to report that his identity was less disrupted than if he goes from kind to cruel.

However, all of these results have been examined in the context of third parties. What if the changes were happening to you?

Differences between self and other

- **Motivational biases towards self-improvement**
(Busseri, Choma, & Sadava, 2009; Haslam, Bastian, Fox, & Whelan, 2007; Wilson & Ross, 2001)
- **Phenomenological identification with current self**
(Loewenstein, O'Donoghue, & Rabin, 2003)
 - **People underestimate their own future change**
(Quoidbach, Gilbert, & Wilson, 2013)
 - **Large change leads to lower connectedness with future self** (Bartels & Urminsky, 2011)
- **More elaborated self-concept, expectations, desires**

Specifically looking at the self is important since research has shown that people often think quite differently about themselves and others.

For example, people are more motivated to see themselves in a positive light. And there is ample research suggesting that people expect that they will personally improve over time. Thus, in the context of identity-continuity, it could be that people see any positive change as consistent with their sense of self.

However, other research stresses the importance of one's phenomenological identification their current viewpoint and state of mind. For example, research on the projection bias has found that people think the preferences they hold now will carry forward into the future. Other research similarly suggests that people underestimate the degree to which they will change over time, and feel less connected to a future self that has changed a lot.

Finally, people tend to have more knowledge of themselves and a more elaborated-self concept than they do for others. Therefore, they are likely to have more developed specific expectations and desires for themselves over time which may also affect their response to changes that are either consistent or inconsistent with these expectations.

Research Question

- **What features can change, and how can they change, while still preserving one's own sense of personal identity?**

- 1) Type of characteristic (central/peripheral?)
- 2) Valence of change
- 3) Expectations of and desires for change

Taken together, the existing research on the self is not entirely clear on how people see themselves changing over time and how this might affect their identity. This inspired us to specifically explore what features can change, and how they can change, while still preserving a sense of your OWN identity continuity.

In the studies that follow we explore the following things, all specifically in the context of perceptions of one's own identity. We vary the type of characteristic changing to see whether people view certain categories of traits as more central or peripheral to their identity, as they do for others.

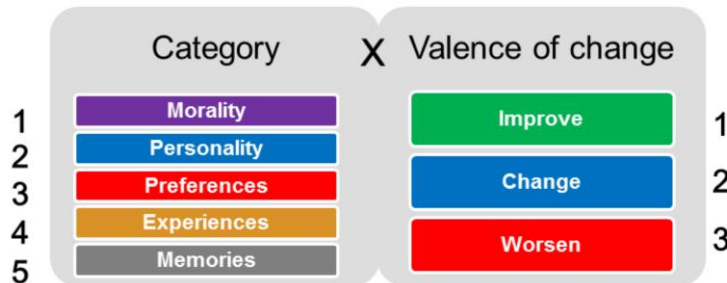
We also look at possible differences in consequences for views of identity based on the valence of change, that is, improvement versus decline

Finally, we also look at possible interactions of the changes with specific expectations and desires, which are likely to be important for the self-concept.

Study 1: Valence and category

Builds upon previous research by:

- a) Examining effects of valence and type of change on perceptions of *self-continuity*
- b) Exploring possible interaction between two factors



Categories:
Strohlinger & Nichols, 2014

The first study looks at valence – whether improvement is treated differently from decline when judging the continuity of one’s own identity.

It also looks at the category of the feature, in order to determine whether changes in certain types of features appear to be more disruptive to identity-continuity. If so, this would suggest that people view certain categories features as more central to their own identity than other features. Although previous research has addressed the effect of these two factors (valence and category) for judgments of others, we extend it to the self. We also examine these in the same experiment, in order to detect a possible interaction between the two, whereas previous research has looked at these separately.

Thus, we employ a mixed design in which we present all participants with characteristics from 5 different categories used in previous research examining perceptions of others’ identity. These categories are morality, personality, preferences, experiences, and memories. We also varied the valence of the change between subjects : in one condition, participants were told to imagine that the listed features improved, in another they changed (with no valence specified), and in the third condition, the features worsened.

Study 1: Valence and category

Stimuli:

40 traits chosen based on pre-testing

- -provide features “important for personal identity”


Morality	Personality	Preferences	Experiences	Memories
Honesty	Impulsiveness	Professional goals	Friendships	Of time spent with family
Values	Sense of humor	Major likes/dislikes	Life experiences	Bad memories

The specific features changing were chosen based on pre-testing, where a group of participants provided groups of features relevant to identity and a different group sorted them into the 5 categories we specified. A few examples of each are listed in the chart.

Study 1: Valence and category

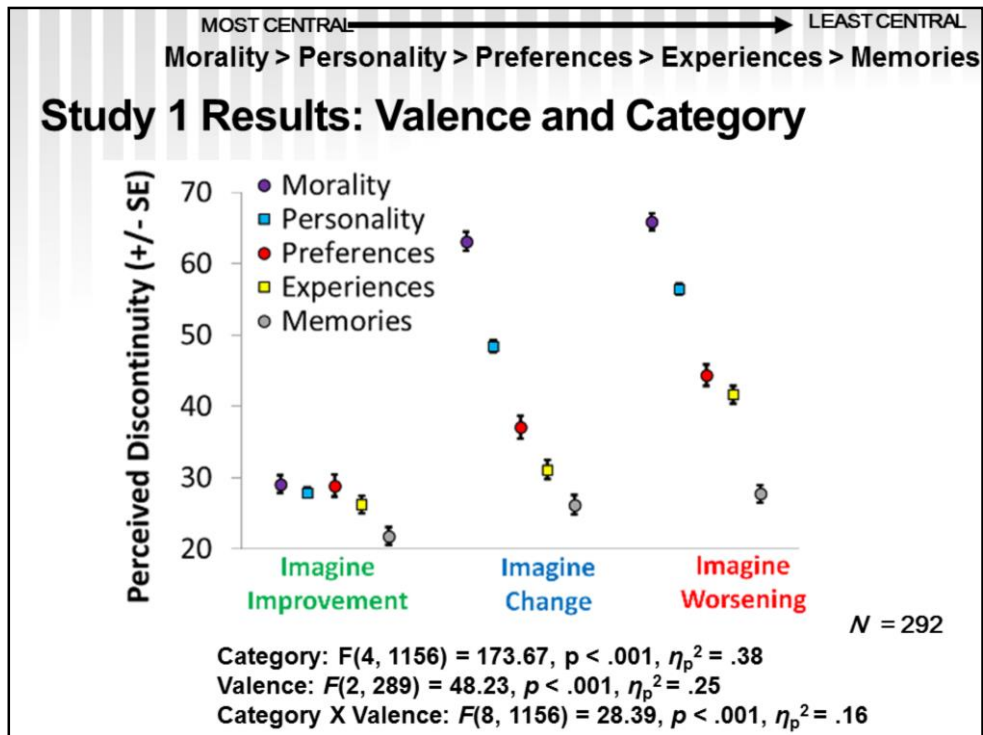
Dependent measure:

- reaction to change in characteristic reflects
centrality to identity (Sloman, Love, & Ahn, 1998)



The dependent measure in these studies was inspired by work on feature centrality. We used a sliding scale designed to measure perceived disruption to self-continuity in response to a change. Ratings at the right end of the scale, “I will not at all be myself”, suggest greater perceived discontinuity. And responses on the left side of our scale “I will really be myself” suggest perceived continuity of one’s identity even after the stated change.

Note that in response to the fact that we are presenting a variety of features, we asked participants to interpret “improvement” or “worsening” however it made sense to them for each one.



Here are the results of the first study.

- 1) first of all, we find a main effect of the different categories of features we used. If we look only at the imagine change condition (where we didn't provide a specific valence), we see that the effect of the categories is consistent with previous research on perceptions of others. That is, just as changes in morality are the most disruptive to our perceptions of other people's identity continuity, as with Phineas Gage, these types of changes also seem to cause the most disruption in perceptions of our own identity continuity. Following morality, changes in personality lead to the most discontinuity, followed by preferences, experiences, and memories.
- 2) We did find a main effect of valence, such that worsening overall led to greater perceptions of discontinuity than improvement. However, another new contribution of this study, in addition to looking at perceptions of the self, was that it specifically found an interaction of category with the valence of change. As you can see from the figure, improvement does not disrupt identity continuity very much for any category of characteristics. However, the increase in discontinuity between improvement and worsening is especially large for central characteristics, and less pronounced for more peripheral characteristics.

- 3) Finally, it is interesting to note that the imagine change condition looks just like the imagine worsening condition. This could explain why some previous research that mentioned general change (without valence) led to lower connectedness with the future self. Therefore, it is possible that people in fact imagine negative changes when we mention “change” in general, though this needs to be explored further in future research.

Study 2: Valence and expectation

Does *prior expectation* about a particular characteristic also affect how we react to a change in it?



In study 2 we examined valence along with prior expectation.

We have seen in previous research that people tend to hold fairly well-formed expectations about their own future, usually in the direction of self-improvement. Could there be an additional interaction between prior expectations of change and valence?

For example, it could be that changes in line with expectations cause less discontinuity independently of the effect of valence.

In this study, we varied both prior expectation and valence completely within subjects. Also note that in this case instead of looking at unvalenced change we ask participants about their reaction to staying the same.

1) Out of the list on the left, please select only **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. Move each of these items to the appropriate box on the right.

Items	My characteristics I expect will IMPROVE
My health	
My helpfulness	
My level of humility	
My professional goals	
My religious or spiritual faith	
My temperament	
My level of happiness	
My confidence level	
My level of selflessness	
My intelligence level	
My friendships	
My dominance	
My personality	
My reliability	
My curiosity	
My degree of independence	
My everyday activities	
My morality	
My impulsiveness	

In order to capture participants actual expectations, we had them select specific characteristics out of our list of 40.

They were told to select:

5 items they most expect to improve

5 items they most expect to stay the same

5 items they most expect to worsen

Once they did that, we had them rank these in order for each box and used only the single most representative characteristic from each expectation bin.

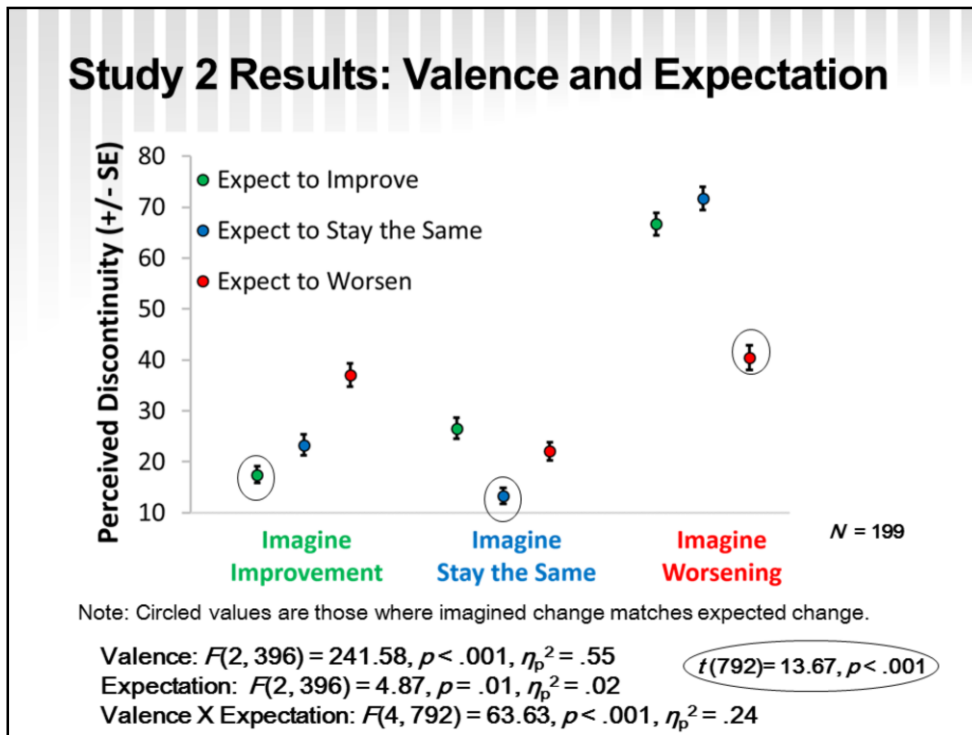
Study 2: Valence and expectation

Dependent measure:

If My reliability improve(s) ...



The dependent measure was exactly the same as it was in study 1.



Here are the results.

1) Again we see an overall main effect of valence. Imagining future worsening leads to a greater disruption in self-continuity than imagining improvement or imagining stability.

2) We also see a main effect of expectation, and an interaction between valence and expectation. If we look at the effect of valence separately for each type of expectation, we can get a better idea about what is driving this interaction.

a) When we expect improvement for a specific characteristic, we experience the least disruption in self-continuity when we imagine that that thing in fact improves, and feel a greater amount of discontinuity when it stays the same or, worst of all, worsens, contrary to our expectations.

b) when we expect stability, we experience the least disruption when this characteristic in fact remains stable. We actually experience significantly more discontinuity if this thing improves, and a much larger disruption in continuity if it worsens.

So, we see that in each of these cases, the type of change that preserves continuity the most is the one that matches prior expectation.

c) If we expect something to worsen, the pattern looks a bit different.

However, overall we see the least disruption to one's own personal continuity if the imagined change matches expectations, as represented by the circled dots. A specific contrast comparing the mean level of discontinuity between the circled versus non matching non circled conditions shows that discontinuity is significantly lower when change matches expectations (which is the circled dots). This, is an important finding as it shows improvement is not **always** the most consistent with identity continuity, but it depends on the specific expectation you held beforehand.

Participant self-selection of traits

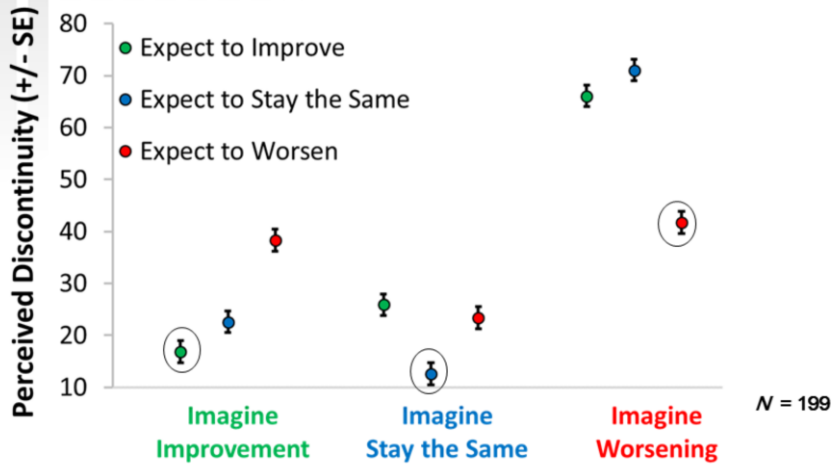
- Important to tap people's genuine, idiosyncratic expectations.
- However, chosen characteristics were not randomly distributed across:
 - Category
 - Less likely to expect change in more central categories
 - Current level
 - Is the characteristics currently at a good or bad level?

Note that although we instructed participants to imagine changes of different valences, they self-selected the traits that were used to define different expectations. This was important because although the overall scenario is hypothetical, we wanted to tap participants' genuine expectations about themselves.

However, this means that the particular stimuli used were **not** randomly assigned. It is possible that there were systematic variations in other aspects of the characteristics placed in each bin.

One of these other aspects was the category of the characteristic, as explored in Study 1. Another was whether people currently perceive themselves to be good or bad on the rated characteristic. This could potentially lead to a ceiling effect if people are selecting characteristics where they already think they are at a very high level.

Study 2: Controlling for level and category



Note: Circled values are those where imagined change matches expected change.

Valence: $F(2, 1579) = 372.40, p < .001$

Expectation: $F(2, 1579) = 0.58, p = .56$

Valence X Expectation: $F(4, 1579) = 57.98, p < .001$

$z = 13.03, p < .001$

When controlled for these factors, we found that results looked pretty much identical, suggesting that the interaction we found between valence and expectation is indeed robust even when controlling for category and level.

Study 3: Valence and desire

Does *specific desire* about a particular characteristic also affect how we react to a change in it?



Study 3 is similar to study 2 but tests for interaction between valence of change and specific desires rather than specific expectations

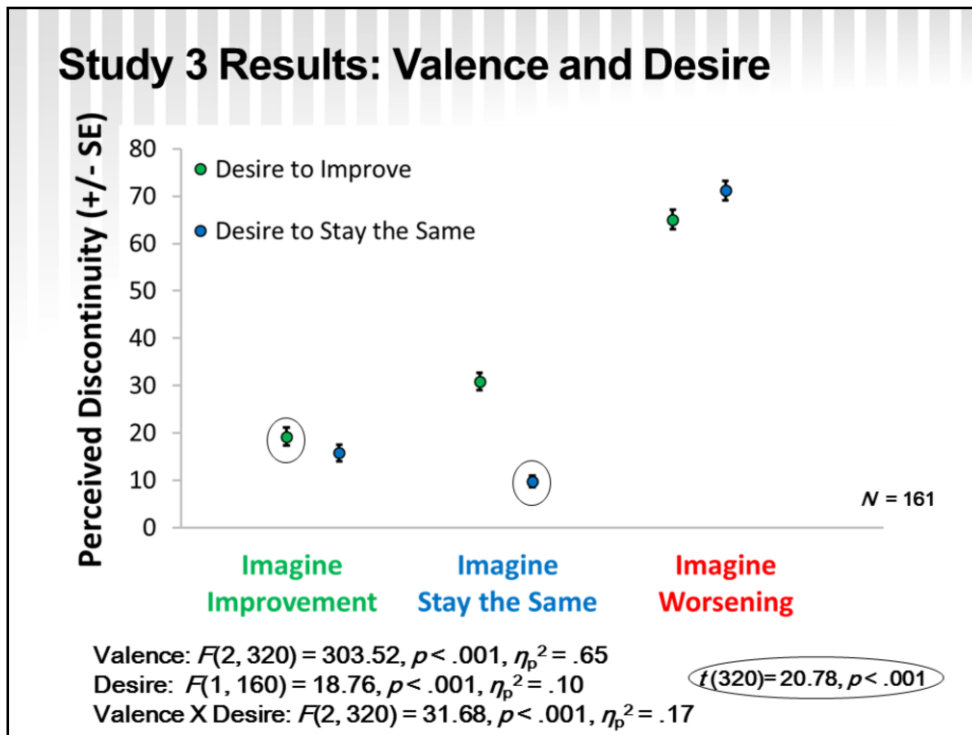
Similar to Study 2, participants selected specific characteristics out of the list of 40 in accordance with their desires.

They selected

5 items they most desire to improve

5 items they most desire to stay the same

We did not ask about desires for worsening since we figured most people would not have such a desire. However, we still included worsening in our valence variable.



We see a similar effect as we did when looking at expectations. Even though there is less disruption to the self-concept when things improve in general compared to when they worsen, people also feel a greater sense of self-continuity when changes match their specific desires.

- 1) Again we see a strong main effect of valence repeated, where imagining future worsening leads to a greater disruption in self-continuity than imagining improvement or imagining stability.

- 2) We do see a main effect of desire, and also a valence by desire interaction.
 - a) When we desire improvement for a specific characteristic, we experience the least disruption in self-continuity when we imagine that that thing in fact improves, and then feel a greater amount of discontinuity when it stays the same or, worst of all, worsens.
 - b) when we desire stability, we experience the least disruption when this characteristic in fact remains stable. We experience significantly more discontinuity if this thing improves, but a much larger disruption in continuity if it worsens.

Just as for expectations, a specific contrast comparing the mean level of discontinuity between the matching versus non matching conditions

shows that discontinuity is significantly lower when change matches expectations. Again, whether improvement or stability is the most consistent with identity continuity depends on your specific desire for that trait.

Study 3 : Controlling for category and level



Valence: $F(1, 792) = 412.91, p < .001$
Desire: $F(1, 792) = 6.30, p = 0.01$
Valence X Desire: $F(2, 792) = 24.55, p < .001$

$z = 18.2, p < .001$

We also repeated the analysis controlling for category and level and again found the same pattern indicating that the effect of desires is not in fact due to either of these factors.

Conclusions

- Imagined change can, but does not always, reduce perceptions of self-continuity over time
- Varies by category (certain traits more central)
- Improvement more consistent with identity than worsening
 - people can hold both stable sense of self and positive expectations
- Change least disruptive when it matches expectations/desires for the self
- distinct from and interactive with valence

Overall, our research helps to clarify which types of change threaten self-continuity, and which ones seem to be consistent with identity. In doing so, we clarify some contradictions in previous research, where some findings suggested change can disrupt the self concept, and others suggesting that people expect to undergo positive change.

First, we found that characteristics of morality and personality traits, do appear central for self-continuity, as they do for the continuity of others. However, we also discovered that even these characteristics are not viewed as completely immutable. Instead, people believe that improvements, even in core features of themselves, are still consistent with a continuous sense of self, and often even more consistent than stability. This suggests that the self-concept likely includes some views about personal development rather than representing only one's current state. This might be one way that people are able to hold positive expectations while still preserving a stable sense of self.

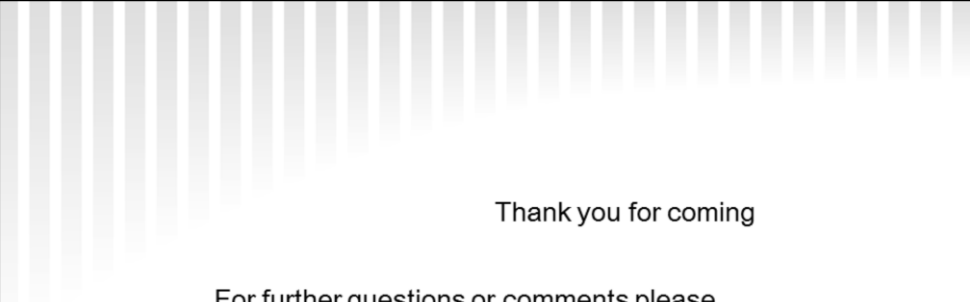
Nevertheless, the self-concept is not always defined by a simple bias towards self-improvement. Rather, people seem to hold a more nuanced view of their identity over time that is shaped by their specific expectations and desires. For example, the conditions that are most conducive towards self-continuity tend to be those that match expectations and desires, which may or may not be

improvement. We showed that this effect was distinct from and interactive with valence. This suggests that people can, for example, adjust their expectations downwards to lessen the disruptive impact of changes that occur.

Implications

- What is the self?
- Effects of imagined and expected change on future-oriented behaviors


This question is important from a philosophical viewpoint, in clarifying how people really think about the self. We find that people don't view the self as a completely unchanging entity, but they also don't view it as blindly improving either. It is also important from a more practical viewpoint, since a disrupted sense of self-continuity can lead to various effects on future-oriented planning and behaviors. We have further research in the works examining this second point about effects of these beliefs.



Thank you for coming

For further questions or comments please
feel free to contact me:

s.molouki@chicagobooth.edu



Study 2 Results: Category and Expectation

MOST CENTRAL



LEAST CENTRAL

	Morality	Personality	Preferences	Experiences	Memories
Improve:	26%	35%	9%	49%	10%
Same:	59%	38%	60%	18%	14%
Worsen:	15%	27%	32%	33%	76%

Bolded values are significantly greater than 33%.

We see that when people placed an item from a more central category into one of the bins, it was more likely they would say the certain characteristic would stay the same. However, when they selected peripheral characteristics they were more likely to say that these would change in one way or another. This makes sense given the results of Study 1, but it also means that we decided to control for category in study 2 to make sure it was not driving the effect of expectations.

Participant self-selection of traits

My CURRENT level of My reliability is...



	Expect Improve	Expect Same	Expect Worsen
Current level	60.75 (23.09)	73.89 (22.48)	48.82 (25.32)

$$F(2,1590) = 248.89, p < .001$$

We also found that expectations varied with participants' perceived current level of the trait. For example, they expected to stay the same on things they were already quite good on. Therefore we also controlled for current level to make sure our expectation results were not due to a ceiling effect.

We pretested a scale of 4 questions and used the question that was most representative of the scale

e.g. [[1) How surprised would you be if your morality improved?

2) How easily can you imagine [a version of] yourself with improved morality?

3) If my morality improves ... I will really be myself – I will not at all be myself

4) How much would an improvement in your morality affect other aspects of yourself?]]

Study 3 Results: Category and Desire

MOST CENTRAL



LEAST CENTRAL

	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

*** = $p < .001$. **Bolded** values are significantly greater than 33%.

Again, we wanted to examine the relationship between category and desire revealed by the self-selection of traits

Categories we had identified as more central in Study 1 are also more desired to stay the same

Levels by expectation, desire and category

	Morality	Personality	Preferences	Experiences	Memories
S2	61.81 (31.17)	64.12 (24.01)	59.77 (19.86)	61.84 (23.08)	51.29 (28.93)
S3	75.25 (21.15)	62.33 (25.34)	47.61 (27.24)	61.74 (23.79)	75.07 (24.77)

S2: $F(4,1588) = 28.15, p < .001$

S3: $F(4,798) = 34.19, p < .001$

EXP: Improve	EXP: Same	EXP: Worsen
60.75 (23.09)	73.89 (22.48)	48.82 (25.32)

$F(2,1590) = 248.89, p < .001$

DES: Improve	DES: Same
51.73 (24.39)	76.32 (19.20)

$F(2,804) = 570.50, p < .001$