

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/234047936>

The End of History Illusion

Article in Science · January 2013

DOI: 10.1126/science.1229294 · Source: PubMed

CITATIONS

204

READS

9,776

3 authors, including:



Jordi Quoidbach

University Pompeu Fabra

47 PUBLICATIONS 4,136 CITATIONS

SEE PROFILE



Timothy D. Wilson

University of Virginia

147 PUBLICATIONS 39,853 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



not entered [View project](#)



Emotional intelligence and health [View project](#)



The End of History Illusion
Jordi Quoidbach *et al.*
Science **339**, 96 (2013);
DOI: 10.1126/science.1229294

This copy is for your personal, non-commercial use only.

If you wish to distribute this article to others, you can order high-quality copies for your colleagues, clients, or customers by [clicking here](#).

Permission to republish or repurpose articles or portions of articles can be obtained by following the guidelines [here](#).

The following resources related to this article are available online at www.sciencemag.org (this information is current as of January 3, 2013):

Updated information and services, including high-resolution figures, can be found in the online version of this article at:

<http://www.sciencemag.org/content/339/6115/96.full.html>

Supporting Online Material can be found at:

<http://www.sciencemag.org/content/suppl/2013/01/03/339.6115.96.DC1.html>

This article **cites 12 articles**, 1 of which can be accessed free:

<http://www.sciencemag.org/content/339/6115/96.full.html#ref-list-1>

This article appears in the following **subject collections**:

Psychology

<http://www.sciencemag.org/cgi/collection/psychology>

The End of History Illusion

Jordi Quoidbach,^{1,2} Daniel T. Gilbert,^{2*} Timothy D. Wilson³

We measured the personalities, values, and preferences of more than 19,000 people who ranged in age from 18 to 68 and asked them to report how much they had changed in the past decade and/or to predict how much they would change in the next decade. Young people, middle-aged people, and older people all believed they had changed a lot in the past but would change relatively little in the future. People, it seems, regard the present as a watershed moment at which they have finally become the person they will be for the rest of their lives. This “end of history illusion” had practical consequences, leading people to overpay for future opportunities to indulge their current preferences.

At every stage of life, people make decisions that profoundly influence the lives of the people they will become—and when they finally become those people, they aren’t always thrilled about it. Young adults pay to remove the tattoos that teenagers paid to get, middle-aged adults rush to divorce the people whom young adults rushed to marry, and older adults visit health spas to lose what middle-aged adults visited restaurants to gain. Why do people so often make decisions that their future selves regret?

One possibility is that people have a fundamental misconception about their future selves. Time is a powerful force that transforms people’s preferences, reshapes their values, and alters their personalities, and we suspect that people generally underestimate the magnitude of those changes. In other words, people may believe that who they are today is pretty much who they will be tomorrow, despite the fact that it isn’t who they were yesterday. In the studies we describe here, we showed that people expect to change little in the future, despite knowing that they have changed a lot in the past, and that this tendency bedevils their decision-making. We call this tendency to underestimate the magnitude of future change the “end of history illusion.”

To investigate this phenomenon, we asked samples of people who varied widely in age to predict how much they would change over the next 10 years, we asked similar samples to report how much they had changed over the past 10 years, and we compared the predictions of people aged a years to the reports of people aged $a + 10$ years. We expected people aged a years to predict less change over the next 10 years than people aged $a + 10$ years reported over the past 10 years. We used this strategy to study how much people thought they would change in the domains of personality (a person’s characteristic patterns of behavior), core values (a person’s ideals and principles), and preferences (a person’s likes and dislikes).

In study 1, we sought to determine whether people underestimate the extent to which their

personalities will change in the future. We recruited a sample of 7519 adults ranging in age from 18 to 68 years [mean (M) = 40 years, standard deviation (SD) = 11.3 years, 80% women] through the Web site of a popular television show and asked them to complete the Ten Item Personality Inventory (I), which is a standard measure of the five trait dimensions that underlie human personality (i.e., conscientiousness, agreeableness, emotional stability, openness to experience, and extraversion). Participants were then randomly assigned either to the reporter condition (and were asked to complete the measure as they would have completed it 10 years earlier) or the predictor condition (and were asked to complete the measure as they thought they would complete it 10 years hence). We then computed the absolute value of the difference between participants’ ratings of their current personality and their reported or predicted personality and averaged these across the five traits to create a measure of reported or predicted change in personality. Additional methodological details about study 1 can be found in supplementary text 1 to 3.

We analyzed these measures by first assigning a value to each of the 41 10-year periods between ages 18 and 68. We called this variable “decade.” For each decade, we compared the predictions of predictors aged a to the reports of reporters aged $a + 10$ years. So, for example, when decade = 1, we compared 18-year-old predictors and 28-year-old reporters; when decade = 2, we compared 19-year-old predictors and 29-year-old reporters; and so on. We did not collect data from reporters who were younger than 28 years, because in our sample there were no predictors younger than 18 years with whom to compare them, and we did not collect data from predictors who were older than 58 years, because in our sample there were no reporters older than 68 years with whom to compare them.

We entered participants’ reported or predicted changes in personality into a multiple regression analysis with three predictor variables: decade (coded 1 through 41), condition (coded 1 for predictors and -1 for reporters), and a “decade X condition” interaction. First, the analysis revealed an effect of decade [beta coefficient (β) = -0.13 , $P < 0.001$], indicating that the older the participants were, the less personality change they reported or predicted. This finding is consistent with a large body of research showing that personality becomes

more stable as people age (2). Second, the analysis revealed the expected effect of condition (β = -0.14 , $P < 0.001$). The top panel of Fig. 1 shows this end of history illusion: Predictors aged a predicted that they would change less over the next decade than reporters aged $a + 10$ years reported having changed over the same decade. Finally, there was no decade X condition interaction (β = 0.01 , $P = 0.68$), indicating that the magnitude of the end of history illusion did not change across decades. Next, we conducted follow-up studies to answer three questions.

First, is it possible that the discrepancy between participants’ reports and predictions in study 1 was due entirely to the erroneous memory of reporters, who may have overestimated how much they had changed in the past 10 years, rather than to the erroneous predictions of predictors, who may have underestimated how much they would change in the next 10 years? To investigate this possibility, we compared the magnitudes of the predicted and reported personality changes in our sample to the magnitude of actual personality change observed in an independent sample of 3808 adults ranging from 20 to 75 years old (M = 47.2 years, SD = 12.4 years, 55% women), whose personalities had been measured as part of the MacArthur Foundation Survey of Midlife Development in the United States (MIDUS). These adults completed the MIDUS Big Five scale (3) for the first time in 1995–1996 (MIDUS 1) and for a second time in 2004–2006 (MIDUS 2). The MIDUS Big Five scale has good construct validity and correlates with other similar scales (4, 5). Because the personality measures used in the MIDUS study and in our study were scored on different scales, direct comparison of the data was not possible. To estimate the magnitudes of actual, reported, and predicted personality change, we computed intraclass correlations (ICC-A1), which account for both absolute and rank-based change (6). Specifically, we computed (i) the ICC between the two administrations of the personality test in the MIDUS sample, which was 0.52; (ii) the ICC between current and reported personality for participants in our sample, which was 0.51; and (iii) the ICC between current and predicted personality for participants in our sample, which was 0.65. Larger ICCs, of course, indicate less personality change. As inspection of these ICCs reveals, the magnitude of reported personality change in our sample was almost identical to the magnitude of actual personality change in the MIDUS sample, suggesting that participants in our sample were relatively accurate when reporting the amount of change they had experienced in the past. However, the magnitude of actual personality change in the MIDUS sample was substantially larger than the magnitude of predicted personality change in our sample, suggesting that participants in our sample were relatively inaccurate when predicting the amount of change they would experience in the future. In short, it seems likely that the discrepancy between the reported and predicted personality

¹National Fund for Scientific Research, Brussels, Belgium.

²Department of Psychology, Harvard University, Cambridge, MA 02138, USA. ³Department of Psychology, University of Virginia, Charlottesville, VA 22904–4400, USA.

*To whom correspondence should be addressed. E-mail: gilbert@wjh.harvard.edu

changes of participants in study 1 is due at least in part to errors of prediction and not merely to errors of memory. Study 3 provides further support for this claim.

Second, is it possible that reporters and predictors in study 1 interpreted the scales differently, so that words such as “conscientious” or “agreeable” meant one thing to reporters and another thing to predictors? To investigate this possibility, we replicated study 1 with an independent sample of 613 adults ($M = 40.5$ years, $SD = 8.4$ years, 86.6% women) recruited through the same Web site and using a design in which each participant was assigned to both the reporter and the predictor conditions, thus ensuring that any idiosyncratic interpretation of the scales would influence both conditions equally. This design required that we restrict our sample to participants aged 28 to 58. Because participants contributed data to both conditions, we performed a multi-level version of the analysis described in study 1. The analysis revealed the expected effect of condition ($\beta = -7.69$, $P = 0.001$), indicating that predictors aged a years predicted that they would change less over the next decade than reporters aged $a + 10$ years reported having changed over the same decade—even though the reports

and predictions were made by the same participants. This finding suggests that idiosyncratic interpretations of the scale are not the cause of the effects seen in study 1.

Third, is it possible that predictors in study 1 knew that they would change over the next 10 years, but because they did not know exactly how they would change, they did not feel confident predicting specific changes? To investigate this possibility, we replicated study 1 with an independent sample of 1163 adults ($M = 38.4$ years, $SD = 12.1$ years, 78% women) recruited through the same Web site. Instead of being asked to report or predict their specific personality traits, these participants were simply asked to report how much they felt they had “changed as a person over the last 10 years” and how much they thought they would “change as a person over the next 10 years.” Because some participants contributed data to both conditions, we performed a multilevel version of the analysis described in study 1. The analysis revealed the expected effect of condition ($\beta = -0.74$, $P = 0.007$), indicating that predictors aged a years predicted that they would change less over the next decade than reporters aged $a + 10$ years reported having changed over the same decade. This finding suggests that

a lack of specific knowledge about how one might change in the future was not the cause of the effects seen in study 1.

In study 2, we sought to determine whether the end of history illusion was limited to the domain of personality, and so we repeated our procedure in the domain of core values. We recruited a new sample of 2717 adults ranging in age from 18 to 68 years ($M = 38.6$ years, $SD = 10.6$ years, 82% women) through the same Web site and asked them to indicate the importance of each of 10 basic values (such as hedonism, success, security, etc.) that were taken from the Schwartz Value Inventory (7). Otherwise, the design was identical to that of study 1.

We performed a regression analysis similar to the one performed in study 1. First, the analysis revealed an effect of decade ($\beta = -0.23$, $P < 0.0010$, indicating that the older participants were, the less change in their core values they reported or predicted. Second, the analysis revealed the expected effect of condition ($\beta = -0.46$, $P < 0.001$). The middle panel of Fig. 1 shows this end of history illusion: Predictors aged a years predicted that they would change less over the next decade than reporters aged $a + 10$ years reported having changed over the




Fig. 1. Standardized predicted and reported changes across decades in study 1 (top panel), study 2 (middle panel), and study 3 (bottom panel). The graph shows moving averages smoothed with a 4-year Gaussian filter. Additional information about this figure can be found in supplementary text 4.

same decade. Finally, the analysis revealed a decade \times condition interaction ($\beta = 0.08$, $P < 0.001$). Although the magnitude of the end of history illusion decreased as participants got older, it was nonetheless present even in the oldest group of participants (aged 50 and up) ($\beta = -0.34$, $P < 0.001$). Further discussion of this decade \times condition interaction can be found in supplementary text 5.

The foregoing studies show that people expect to experience less change in their personalities and core values over the next decade than people a decade older report having experienced over the past decade. The analysis presented in study 1 suggests that this discrepancy represents, at least in part, an error of prediction and is not merely an error of memory. To provide further support for this claim, in study 3 we examined the end of history illusion in a domain in which memory was likely to be highly reliable. Rather than asking reporters to remember how extraverted they had been or how much they had once valued honesty, we asked them to remember simple facts about their strongest preferences, such as the name of their favorite musical band or the name of their best friend. We reasoned that if participants remembered having a different best friend 10 years ago but expected to have the same best friend 10 years from now, then this was probably not due to a pervasive tendency for people of all ages to actually keep their best friends but mistakenly remember changing them.

To test this hypothesis, we recruited a new sample of 7130 adults ranging from 18 to 68 years old ($M = 40.2$ years, $SD = 11.1$ years, 80% women) through the same Web site and asked them to report their favorite type of music, their favorite type of vacation, their favorite type of food, their favorite hobby, and the name of their best friend. Participants were then randomly assigned either to the reporter condition (and were asked to report whether each of their current preferences was the same as or different than it was 10 years ago) or the predictor condition (and were asked to predict whether each of their current preferences would be the same or different 10 years from now). We then counted the number of items on which participants responded "different" and used this as a measure of reported or predicted changes in preference.

We performed a regression analysis similar to the ones performed in studies 1 and 2. First, the analysis revealed an effect of decade ($\beta = -0.14$, $P < 0.001$). The older participants were, the less change in preferences they reported or predicted. Second, the analysis revealed the expected effect of condition ($\beta = -0.19$, $P < 0.001$). The bottom panel in Fig. 1 shows this end of history illusion: Predictors aged a years predicted that their preferences would change less over the next decade than reporters aged $a + 10$ years reported that their preferences had changed over the same decade. Finally, the analysis revealed a decade \times condition interaction ($\beta = 0.07$, $P < 0.001$). Although the magnitude of the end of history illusion decreased as partic-

ipants got older, it was nonetheless present even in the oldest group of participants (aged 50 and up) ($\beta = -0.08$, $P < 0.01$). Further discussion of this decade \times condition interaction can be found in supplementary text 5, and additional details about study 3 can be found in supplementary text 6.

The foregoing studies suggest that people underestimate the extent to which their personalities, values, and preferences will change in the future. In study 4, we sought to show that this end of history illusion can have practical consequences. Specifically, we sought to show that because people overestimate the stability of their current preferences, they will overpay for future opportunities to indulge them.

In study 4, we recruited a new sample of 170 adults ranging from 18 to 64 years old ($M = 34.9$ years, $SD = 10.6$ years, 52% women) through the Amazon Mechanical Turk Web site (8, 9). Some participants were randomly assigned to the "future concert" condition. These participants were asked to name their current favorite musical band and then to report the maximum amount of money they thought they would be willing to pay today in order to see that band perform in 10 years. Other participants were randomly assigned to the "present concert" condition. These participants were asked to name the musical band that was their favorite 10 years ago and then to report the amount of money that they thought they would be willing to pay today to see that band perform in the coming week.

We performed a regression analysis similar to the ones performed in studies 1, 2, and 3. First, the analysis revealed the expected effect of condition ($\beta = 0.16$, $P < 0.05$). Participants aged a years thought they would pay 61% more to see their current favorite band perform 10 years in the future ($M = \$129$) than participants aged $a + 10$ years thought they would pay to see their once-favorite band perform in the present ($M = \$80$). The analysis revealed no effect of decade ($\beta = -0.06$, $P = 0.41$), indicating that the price participants thought they would pay did not vary with age, and no decade \times condition interaction ($\beta = 0.01$, $P = 0.94$), indicating that willingness to pay more for a future concert than a present concert did not diminish in magnitude as participants got older. In short, participants substantially overpaid for a future opportunity to indulge a current preference.

Across six studies of more than 19,000 participants, we found consistent evidence to indicate that people underestimate how much they will change in the future, and that doing so can lead to suboptimal decisions. Although these data cannot tell us what causes the end of history illusion, two possibilities seem likely. First, most people believe that their personalities are attractive, their values admirable, and their preferences wise (10); and having reached that exalted state, they may be reluctant to entertain the possibility of change. People also like to believe that they know themselves well (11), and the possibility of future change may threaten that belief. In short,

people are motivated to think well of themselves and to feel secure in that understanding, and the end of history illusion may help them accomplish these goals.

Second, there is at least one important difference between the cognitive processes that allow people to look forward and backward in time (12). Prospecion is a constructive process, retrospection is a reconstructive process, and constructing new things is typically more difficult than reconstructing old ones (13, 14). The reason this matters is that people often draw inferences from the ease with which they can remember or imagine (15, 16). If people find it difficult to imagine the ways in which their traits, values, or preferences will change in the future, they may assume that such changes are unlikely. In short, people may confuse the difficulty of imagining personal change with the unlikelihood of change itself.

Although the magnitude of this end of history illusion in some of our studies was greater for younger people than for older people, it was nonetheless evident at every stage of adult life that we could analyze. Both teenagers and grandparents seem to believe that the pace of personal change has slowed to a crawl and that they have recently become the people they will remain. History, it seems, is always ending today.

References and Notes

1. S. D. Gosling, P. J. Rentfrow, W. B. Swann Jr., *J. Res. Pers.* **37**, 504 (2003).
2. B. W. Roberts, K. E. Walton, W. Viechtbauer, *Psychol. Bull.* **132**, 1 (2006).
3. M. E. Lachman, S. L. Weaver, *The Midlife Development Inventory (MIDI) Personality Scales: Scale Construction and Scoring* (Brandeis University, Waltham, MA, 1997).
4. D. K. Mroczek, C. M. Kolarz, *J. Pers. Soc. Psychol.* **75**, 1333 (1998).
5. K. M. Prenda, M. E. Lachman, *Psychol. Aging* **16**, 206 (2001).
6. K. O. McGraw, S. P. Wong, *Psychol. Methods* **1**, 30 (1996).
7. S. H. Schwartz, in *Advances in Experimental Social Psychology*, M. P. Zanna, Ed. (Academic Press, Orlando, FL, 1992), vol. 25, pp. 1–65.
8. G. Paolacci, J. Chandler, P. G. Ipeirotis, *Judgm. Decis. Mak.* **5**, 411 (2010).
9. M. Buhrmester, T. Kwang, S. D. Gosling, *Perspect. Psychol. Sci.* **6**, 3 (2011).
10. C. Sedikides, M. D. Alicke, in *The Oxford Handbook of Human Motivation*, R. M. Ryan, Ed. (Oxford Univ. Press, Oxford), pp. 303–322.
11. W. B. Swann Jr., in *Handbook of Theories of Social Psychology*, P. Van Lang, A. Kruglanski, E. T. Higgins, Eds. (Sage, London, 2012), pp. 23–42.
12. D. R. Addis, A. T. Wong, D. L. Schacter, *Neuropsychologia* **45**, 1363 (2007).
13. M. D. Robinson, G. L. Clore, *Psychol. Bull.* **128**, 934 (2002).
14. M. Ross, *Psychol. Rev.* **96**, 341 (1989).
15. N. Schwarz et al., *J. Pers. Soc. Psychol.* **61**, 195 (1991).
16. A. Tversky, D. Kahneman, *Cognit. Psychol.* **5**, 207 (1973).

Acknowledgments: We acknowledge the support of Research Grant BC5-0722132 from NSF to D.T.G. and T.D.W. Raw data from all studies are on deposit at the Inter-university Consortium for Political and Social Research (deposit no. 32668) and can be accessed at www.icpsr.umich.edu.

Supplementary Materials

www.sciencemag.org/cgi/content/full/339/6115/96/DC1
Supplementary Text
Table S1

24 August 2012; accepted 16 November 2012
10.1126/science.1229294



Supplementary Materials for

The End of History Illusion

Jordi Quoidbach, Daniel T. Gilbert,* Timothy D. Wilson

*To whom correspondence should be addressed. E-mail: gilbert@wjh.harvard.edu

Published 4 January 2013, *Science* **339**, 96 (2012)
DOI: 10.1126/science.1229294

This PDF file includes:

Supplementary Text
Table S1

Supplementary Materials

1. “Leurs Secrets du Bonheur” (“Their Secrets of Happiness”) is a French television show that aired on the channel *France 2* from October 2011 to January 2012. It invited viewers to participate in social science studies at the show’s website. We received permission to place a link to our studies on that website. Participants who clicked that link were assigned to one of our studies. Participants were given no financial compensation but were told before participating that they would receive feedback about their levels of wellbeing when the study was complete. Participants in Study 1, the follow-ups to Study 1, Study 2, and Study 3 were recruited via this method. Participants in Study 4 were recruited through the Amazon Mechanical Turk website.
2. During a first wave of data collection in November, 2011, participants who clicked our link were randomly assigned to participate in Study 1, a follow-up to Study 1, or Study 3. During a second wave of data collection in January, 2012, participants who clicked our link were randomly assigned to participate in a follow-up to Study 1 or Study 2.
3. In addition to the measures described in the manuscript, participants in Study 1, the follow-ups to Study 1, Study 2, and Study 3 completed numerous other questionnaires for other research projects (e.g., measures of satisfaction with life, depression, political orientation, income, etc.).
4. For clarity of presentation, we applied a Gaussian filter to smooth short-term fluctuations and highlight longer-term trends in Figure 1. A Gaussian filter replaces each value with the weighted average of neighboring values, and those weights are defined by a Gaussian function. We set the standard deviation of the Gaussian function to 4 years—with repetition of the values at both extremities to avoid edge effects—meaning that all low-frequency fluctuations within a four-year period were smoothed. Figure S1 shows the unfiltered data.

To allow visual comparison of the results across studies, change scores in each study were transformed into percentages of change. So a score of 100% means the highest possible change score—that is, going from one extreme of the rating scale to the other for *all* the personality traits (Study 1), all the values (Study 2), or indicating that all of one's preferences will be different (Study 3). Scores of 0% indicates no change.

5. In Studies 2 and 3—but not in Study 1—the magnitude of the end of history illusion was larger among younger than older participants. Did the illusion merely diminish among older participants or did it actually disappear? In all three studies, the illusion was evident when we analyzed the data from our oldest participants as a group (i.e., predictors who were 50 years and older and reporters who were 60 years and older). Unfortunately, our samples did not contain a sufficient number of older participants to allow us to conduct meaningful analyses on participants at every age (see Table S1). More research will be needed to determine whether the illusion does or does not disappear at the very upper end of the age continuum.
6. In Study 3, the five preferences questions were originally scored on a 4-point scale from 1 (*Certainly the same*) to 4 (*Certainly different*). Although results using this continuous measure were significant (β condition = $-.06$, $p < .001$), we dichotomized the response scale for the sake of clarity. Also, in addition to asking participants about music, vacations, food, hobbies, and best friends, we also asked about their favorite movie. We eliminated this item from the analyses reported in the manuscript because more than 200 participants failed to complete it, suggesting that people do not find it easy to remember their favorite movie from a decade ago. In comparison, every participant completed every other item. Including this item in the analyses reported in the manuscript does not change the significance of the result (β condition = $-.12$, $p < .001$).
7. More than 80% of the participants in Study 1, Study 2, and Study 3 were women, so we also

performed regression analyses on men and women separately to ensure that the results were not limited to a single gender. These analyses revealed an end of history illusion for both genders. Specifically, analyses of men revealed an effect of condition in Study 1, Study 2, and Study 3 ($\beta = -.20, p < .001$, $\beta = -.39, p < .001$, and $\beta = -.14, p < .001$, respectively), and analyses of women revealed an effect of condition in Study 1, Study 2, and Study 3 ($\beta = -.12, p < .001$, $\beta = -.48, p < .001$, and $\beta = -.20, p < .001$, respectively).

Table S1. Number of participants (N) by age and condition.

Age	Study 1		Study 2		Study 3		Study 4	
	Reporters N	Predictor N	Reporters N	Predictor N	Reporters N	Predictor N	Reporters N	Predictor N
18		66		82		33		1
19		71		94		27		3
20		96		91		42		4
21		79		109		33		4
22		65		85		44		4
23		91		104		54		11
24		116		95		43		7
25		108		116		55		3
26		90		89		29		3
27		98		79		40		4
28	98	104	105	105	55	54	4	3
29	115	103	101	101	58	38	4	5
30	116	134	141	141	55	56	3	5
31	107	111	123	123	48	63	6	3
32	119	105	123	123	52	52	4	2
33	119	112	127	127	53	59	2	5
34	87	112	122	122	52	57	2	2
35	128	113	97	97	48	48	6	3
36	108	121	120	120	51	49	2	1
37	122	120	110	110	45	51	2	2
38	118	102	136	136	56	54	1	2
39	115	111	126	126	53	51	3	3
40	111	104	109	109	60	47	5	2
41	89	103	110	110	39	45	1	1
42	88	104	103	103	38	58	4	2
43	85	99	101	101	30	31	1	0
44	93	82	96	96	28	41	0	2
45	98	121	97	97	35	34	2	1
46	82	87	105	105	49	39	2	1
47	93	86	87	87	29	30	0	1
48	111	108	83	83	38	32	0	2
49	85	102	87	87	27	27	1	1
50	80	99	86	86	39	34	2	0
51	76	73	75	75	30	23	3	1
52	79	81	71	71	30	33	0	2
53	83	66	84	84	32	26	0	1
54	80	69	76	76	20	20	3	1
55	75	72	66	66	23	20	3	0

56	60	68	58	58	12	19	0	0
57	46	54	70	70	14	17	0	0
58	66	50	49	49	17	20	1	1
59	43		44		18		1	
60	67		66		16		0	
61	44		51		20		1	
62	33		55		15		1	
63	42		34		8		0	
64	36		52		6		1	
65	34		33		4		0	
66	16		20		4		0	
67	14		18		8		0	
68	13		15		2		0	

Reported

Predicted

