

An Empirical Test of How Events Turn the Cognitive Gears of Trust

D. Harrison McKnight¹ and Peng Liu²

¹ Accounting and Information Systems Department, Michigan State University
mcknight@bus.msu.edu

² Department of Information Systems & Decision Science, California State University, Fullerton
peliu@fullerton.edu

Abstract. This study empirically tests a social psychology-based Information Processing Model (IPM) that explains how events may change trust over time based on three cognitive mechanisms or “gears”: attention, attribution, and judgment. We briefly describe the IPM, and then empirically validate its assumptions and extensions. The IPM is contrasted with the incremental growth model (IGM) of trust change. We find more support for the IPM than the IGM.

Keywords: Trust, attention, attribution, threshold, risk, illusion, change.

1 Introduction

How does trust in another party change in response to a series of events over time? How do cognitive mechanisms produce trust changes over time? Few studies have addressed these questions (e.g., [3]). These questions are vital because trust in a technology changes over time. Further, the world is becoming scarier in terms of how adroitly online agents are attacking targets once considered safe, which changes trust.

McKnight et al. [6] addressed these questions by developing a psychology-based information processing model (IPM) of the major cognitive mechanisms involved in trust changes. Then they simulated the model to illustrate how it works under a set of plausible yet untested assumptions. But their assumptions may or may not be right. Incorrect assumptions would render the IPM useless. We test their assumptions.

The IPM can serve as a basic way to depict trust changes in the cyber world. By trust we mean a willingness to be vulnerable to another party [4]. See McKnight et al. [6] for the general developmental process theory approach behind this paper and the rationale for the theoretical model being tested. We next briefly describe the IPM. Then we explain the empirical methods used. We report the results, which validate much of the IPM and suggest refinements. This study enhances the model’s value.

2 The Baseline Information Processing Model

2.1 Model Overview

Model action initiates when an event takes place that may have trust ramifications (Figure 1, 1. Event). The event may be an interaction with the trustee, or it may be some signal or second-hand report about the trustee (e.g., media reports).

The model proceeds clockwise. Time between events may be short or long. The event passes through the perception and memory module, which provides a basis for noticing or ignoring the event. Gear 2 (attention gear) refers to whether the trustor notices the event. If the event is not noticed (2. “No” path), no subsequent gears are engaged, trust stays the same, and the process awaits the next event. Attention precedes attribution. The need for each subprocess and the way subprocesses hand off to others led us to think of them as cognitive gears that may each engage the next gear.

Gear 3 involves attribution to see if the event contradicts the current trust level. After attribution, the trustor assesses whether the attributed contradiction is serious enough to exceed the threshold cost of updating the trust level (gear 4). If not, the trust level is not updated and processing awaits the next event. If so, the update gear is initiated, which increases/decreases the trust level. The person then brings a new level of trust to the perceptual system that encounters the next event. Not only is the trust level updated, but the likelihood of attending to and attributing about an event is also updated. We next briefly present the assumptions of the model.

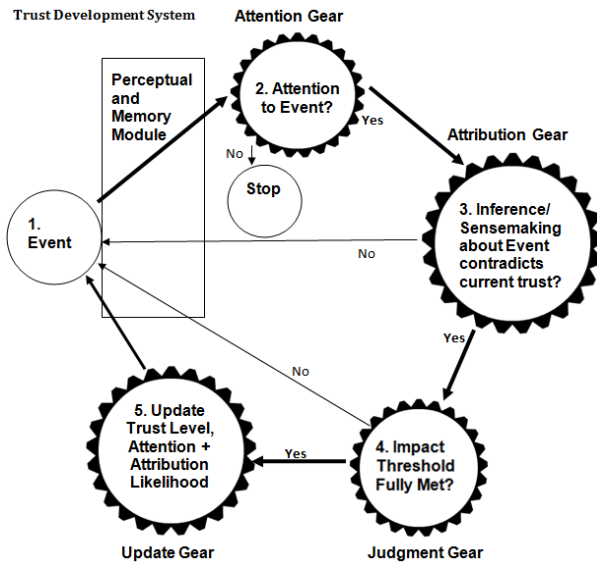


Fig. 1. Trust Development: A Baseline Information Processing Model (IPM)

The Attention Gear. By attention, we mean the person notices the event instead of ignoring it. Due to bounded rationality, people do not attend to everything. A change in trust due to an event cannot occur unless the trustor notices the event. Thus Figure 1 depicts attention as a necessary mental process that happens before trust changes. Although attention is necessary, it is not sufficient to change trust. The other gears are needed.

Assumption 1: Trustor attention to a behavioral event is a necessary but not sufficient condition for the change in a trustor’s level of trust.

The Attribution Gear. Attribution means how one makes sense of an event in a way that could change one's beliefs or attitudes. Attribution is a highly subjective, individual, and often irrational process, so its outcomes are difficult to predict. Hence, many events that are noticed are never fully or rationally appraised. Only when the cause of an event is decided, and further, only when the event's attribution clearly contradicts one's existing trust perceptions is a trust update very likely.

Assumption 2: Trustor attribution that a behavioral event contradicts current trust levels is a necessary but not sufficient condition for the change in the level of trust.

Note the IPM does not specify sufficient conditions for trust change. This is typical for process theories [7]. Process theories set up the necessary conditions and provide some assumptions to guide what makes change more or less likely. They assume non-deterministic and probabilistic change mechanisms.

The Threshold Gear. After making an attribution, a trustor judges whether it meets the threshold for revising trust. This is termed a judgment gear because it compares the benefits of updating trust with the costs of updating trust. We propose the threshold for willingness to change trust is met when the perceived benefits of changing trust significantly exceed the perceived costs of changing trust.

Assumption 3: Exceeding a judgment threshold is a necessary condition for the change in a trustor's level of trust.

The Trust Update Gear. Once the threshold is exceeded, the trust level is updated. The likelihood of attending to and attributing about the next event are also updated. If trust goes down, one will be more likely to notice and attribute about the next event because one's state worsens. If trust goes up, one will be less likely to notice.

2.2 Three Model Extensions—Risk, Negative Change, Illusion

Conditions that affect how those gears operate can be used to improve the model. We test McKnight et al.'s [6] three extensions, restating some for testing purposes.

Situational Risk Extension. Low risk situations garner less attention than do high risk situations. Risk justifies the mental energy needed for attribution. Because high risk involves higher attention and attribution, and because higher attention/attribution lead to more trust change, higher risk will lead to greater / more frequent trust change.

Extension 1: Higher perceived situational risk will significantly increase the likelihood of both attention and attribution to events. With higher perceived situational risk, trust changes will be significantly greater in magnitude and frequency than with lower perceived situational risk.

Negative Asymmetry Extension. Will trust decreases be greater in magnitude and frequency than trust increases? The literature suggests the answer is "yes" [7].

Extension 2: Negative events will be attended to and attributed about more than positive events. Negative trust changes will be significantly greater in magnitude and frequency than will positive trust changes.

Illusion Extension. In loyal close personal relationships people idealize their partner and use that idealized perception to dispel the effects of negative events. The same may occur in other relationships. Illusion means the extent to which a partner is more optimistic about their relationship's future than rationally merited. When one has positive illusions, a mix of positive and negative events will continue those illusions.

Extension 3: Illusionary (i.e., highly loyal) partners will more likely reinterpret negative events as positive events. Thus trust change for the highly loyal will be lower in magnitude and frequency than for those with low loyalty.

3 Methods

We collect online survey data to empirically test IPM Assumptions 1 and 2 (but not 3) and Extensions 1-3. The survey measures the trust respondents have in a technology product. Survey participants are students recruited from information systems and business courses at two U. S. universities. 69% of subjects were given extra credit (1% of course points); 31% received a 12% chance of receiving a \$25 Amazon.com gift certificate. The subjects were assigned quasi-randomly to give their opinions about one of four different technologies: a TomTom GPS, a Google driverless car, Ford's SYNC product, or Apple's Siri feature. These were chosen because student subjects should be interested in them. Also, we wanted the technologies to be relatively new to subjects so we could test the IPM in an initial trust, pre-adoption setting.

Subjects first provided their loyalty to each technology vendor. In real life, consumers often become loyal to a vendor like Apple. We wanted more loyal respondents than would be found in a totally random assignment to technology. So we assigned subjects with a loyalty score of 7/7 to that technology group. 62 were 7/7 loyal to Apple, 66 to Google, 10 to Ford and 4 to TomTom. When subjects had 7-level loyalty to Apple and another vendor, they were assigned to the Siri treatment. Subjects with no loyalty 7s were assigned randomly. Thus, we had more loyal subjects than a totally random distribution would provide. This also provides a greater number of subjects to whom the technology will seem relevant. Subjects provided their initial impressions of the technology and then saw a series of eight news briefs about that technology (four positive and four negative). Within each treatment group, we presented all subjects the same eight articles. The order of article presentation was done in eight alternating patterns nested within the technology randomization [+ - + - + - + -], [- + - + - + - +], [+ + - - + + - -], [- - + + - - + +], [+ + + + - - - -], [- - - - + + + +], [+ + - - - - + +], [+ + - + + - - -]. This design allows us to vary the way events are fed to them in systematic ways, enabling us to test theory not reported here. To select the events used we first had seventeen students rate eighty-five candidate news briefs for positivity/negativity on a 1-7 scale (from strongly negative to strongly positive). Each of the eighty-five news briefs was a shortened version of a recent months' newspaper

articles found using Nexis. The content meaning was carefully preserved during editing. We chose eight news briefs for each technology product based on the average of students' ratings. From the rated articles we chose two highly positive, two highly negative, one moderately positive, one moderately negative, one slightly positive and one slightly negative briefs for each technology (Table 1). The mean positivity rating differences among the four technology products are not statistically significant.

Table 1. Average Negative/Positive Rating of Events

Events (Negative-Positive levels)	Technology 1 Apple Siri	Technology 2 Google Car	Technology 3 TomTom GPS	Technology 4 Ford Sync
Event 1 (++)	5.7	5.6	5.6	5.1
Event 2 (-)	3.2	3.1	3.1	3.4
Event 3 (+)	4.9	4.9	5.1	4.3
Event 4 (---)	1.6	1.8	1.8	1.7
Event 5 (+++)	6.0	6.0	6.3	6.3
Event 6 (--)	2.2	2.4	2.5	2.2
Event 7 (+++)	6.0	6.2	6.3	5.9
Event 8 (---)	1.6	1.5	1.5	1.8
Average	3.9	3.9	4.0	3.8
Scale: 1 Strongly Negative, 2 Moderately Negative, 3 Slightly Negative, 4 Neutral, 5 Slightly Positive, 6 Moderately Positive, 7 Strongly Positive.				

Table 2. Respondent Experience with the Technology

	Average Across Technologies	Apple Siri	Google Car	TomTom GPS	Ford Sync
Time Length of Use (1)	1.5	1.9	1.0	1.8	1.3
Use Frequency (2)	1.6	2.2	1.0	1.6	1.5
Time length of use and use Frequency are measured on 1-7 pt. scale. (1) 1: Have not used at all; 2: < 1 year; 3: 1-2 years; 4: >2 but <3 years; 5: >3 but <4 years; 6: >4 but <5 years; 7: > 5 years (2) 1: Not at all; 2: Rarely; 3: 1-2 times/month; 4: 1-2 times/week; 5: 3-5 times/week; 6: Daily; 7: Multiple times/day					

The study proceeded as follows. Instructors introduced the survey to students with a 2 minute introduction. The introduction included a Powerpoint invitation page (also posted on the course website) with a Web address students could click on to start the survey. Students were given two-three weeks to do the survey on their own time, with two-three reminders. 60% were male and the average age was 21. They were asked their level of loyalty to Apple, Google, TomTom, and Ford: (e.g., "As long as Apple makes computers and phones, I doubt that I would buy these from anyone else."). Then they were assigned a technology and shown a brief Wikipedia summary of what the technology was/did. This helped all subjects know its features. Next, we asked subjects their use, enjoyment, trust, and intention to use the technology. Next they

were told to read carefully eight technology news briefs. After each they were asked their attention level and to briefly explain why; then their attribution level and why. They were asked their trust in the technology and their intent to use the technology. After the eight news brief cycles, subjects were asked their technology enjoyment (“I would enjoy using the [technology].”), perceived risk, and perceived relevance.

For the attention question, we related the technology to them personally: “If you saw the above news brief while you were considering [technology] for possible use, would you ignore the news brief or would you pay attention to it?” (1= “I would Definitely Ignore it”; 7= “I would Definitely Pay Attention to It”). The wording reflects that we were already asking them to read the news brief, garnering basic attention, but goes further to tap the depth-of-attention idea [2]. The online survey system records the time they spend reading the news brief and the time spent answering the attention question and its why explanation. We use both their numeric survey response and the time reading the brief and answering the attention questions to represent event attention. If respondents spend more time reading or mulling over this question, it means they pay more attention [2]. For attribution, we ask, “To what extent (if at all) does this news brief influence you to reconsider how much you trust the [technology]?” (1= Not at all; 7= A Huge Amount). As Figure 1 shows, this wording reflects an attribution thorough enough to possibly consider a trust change. We also use the time they spend on this question as an alternative measure of attribution about the news brief. If respondents spend more time, it means they make sense of it more completely.

At the end of the session, we tap their risk feelings: “How would you characterize the decision of whether to use this technology product?” (1= “significant opportunity” to 7= “significant risk” – [9]). Trust is measured by “For doing [actions related to the technology], I feel I can depend on the [technology].” [5]. Choices range from 0= “strongly disagree” to 10= “strongly agree” so it is comparable to the 0-1.0 scale used in [6]. Due to the length of a survey with eight event iterations, we had to use 1-item trust and perceived risk scales, each from a well-validated multi-item scale (e.g., [5]).

4 Results

We found the mean attention level (4.9/7.0) was higher than the attribution mean (3.9/7.0). The average trust level decreased from 6.7 to 6.3 (Table 3). Not every trust level went down; only 42.4% of subjects lowered their trust from T0 to T8 (Table 4). Table 4 also shows that 28.9% of the trust levels did not change at all, and another 28.7% went up. The average trust change magnitude over 8 periods was -0.52.

Assumptions 1 and 2 posit that attention and attribution are necessary to trust change. Developmental process studies use two key metrics: magnitude and frequency of change [6]. Magnitude is the size of the change in trust level after an event. Frequency means whether the trust level changes or not after an event, 1 means trust change, 0 means none. We test to see if high attention/attribution are related to higher magnitude and higher frequency of trust changes. For most tests, we use two-group means difference tests. This method was chosen both because it is the simplest test and enables the results to be displayed in a clear manner. The 422 respondents are

each presented eight news briefs which means eight possible trust changes (Trust(T)-Trust(T-1)). The sample size for some tests is 8 X 422 = 3376. We use absolute trust change value for assumptions 1 and 2 and extension 2, because we are interested first in the magnitude of trust change. We recoded the time index for attention and attribution into a 1 to 7 scale so it is equivalent to the measured item. Group 1 has attention/attribution levels from 1 to 3, and group 2 has attention/attribution levels 5 to 7. Assumption 1 and 2 say attention and attribution are necessary for trust change. If so, then those with higher attention and attribution should have a higher trust change. As Tables 5A and 5B show, the high attention group has a larger (5A) and more frequent (5B) change in trust than the low group. The same is true for attribution. Assumptions 1 and 2 are supported in terms of both magnitude and frequency of trust change.

Table 3. Variable Means

	Average	Apple Siri	Google Car	TomTom GPS	Ford SYNC
Trust T0	6.7	7.3	5.7	7.4	6.7
Trust T8	6.3	7.3	5.6	6.1	6.0
%Change	-6.7%	0%	-1.8%	-17.6%	-10.4%
Intent to Use T0	4.9	4.8	5.3	4.5	4.7
Intent to Use T8	4.4	4.8	4.4	3.9	4.3
%Change	-10.2%	0%	-17.0%	-13.3%	-8.5%
Enjoyment T0	4.6	4.8	4.9	3.9	4.8
Enjoyment T8	4.6	5.0	4.6	4.1	4.4
%Change	0%	+4.2%	-6.1%	+5.1%	-8.3%
Risk of Tech	3.8	3.3	4.3	4.0	3.8
Relevance of Tech	4.2	3.8	4.2	4.7	4.2
Loyalty to Vendor	3.6	3.9	4.8	2.7	2.4
Note: Trust is measured on 1-10 scale; others on 1-7 scale					

Table 4. Trust Change Frequency from Time 0 to Time 8

Trust Change	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0
# Changing	2	4	0	6	7	14	16	29	37	64	122
% of Total	0.5	0.9	0	1.4	1.7	3.3	3.8	6.9	8.8	15.2	28.9
Cum.% of Total	0.5	1.4	1.4	2.8	4.5	7.8	11.6	18.5	27.3	42.4	71.3
Trust Change	1	2	3	4	5	6	7	8	9	10	n
# Changing	50	28	20	11	8	2	0	2	0	0	422
% of Total	11.8	6.6	4.7	2.6	1.9	0.5	0	0.5	0	0	100
Cum.% of Total	83.2	89.8	94.5	97.2	99.1	99.5	99.5	100	100	100	

To test extension 1 (high situational risk → higher attention, attribution, trust change), we need to know the risk level of the four technologies. The average Apple Siri risk is 3.3/7.0 (n=133); for SYNC 3.8/7.0 (n=81); TomTom 4.0/7.0 (n=81); and Google 4.3/7.0 (n=127). We tested for risk differences and found Siri’s risk significantly differed from that of the other three (0.000<p<0.0012). However, the risk of Google, TomTom, and Ford were not significantly different from each other. So Apple Siri represents the low risk group and the others form the high risk group.

Table 5a. Assumptions 1 and 2 Two-group mean Trust Change (Δ) Magnitude difference tests

	n	Mean Trust Δ	Std. Dev.	t-value	p-value	Support
Using Attention measured by survey:						
Low Attention (1-3)	609	0.71	1.36	-4.32	0.000	Assumption 1 is supported.
High Attention (5-7)	2185	1.02	1.63	.	.	
Using Attention measured by Subject response time:						
Low Attention (1-3)	1449	0.80	1.49	-3.85	0.000	Assumption 1 is supported.
High Attention (5-7)	1447	1.03	1.67	.	.	
Using Attribution measured by survey:						
Low Attribution(1-3)	1252	0.73	1.47	-7.10	0.000	Assumption 2 is supported.
High Attribution(5-7)	1324	1.18	1.73	.	.	
Using Attribution measured by Subject response time:						
Low Attribution(1-3)	1440	0.80	1.51	-3.78	0.000	Assumption 2 is supported.
High Attribution(5-7)	1438	1.02	1.59	.	.	

Table 5b. Assumptions 1 and 2 Two-group mean Trust Change (Δ) Frequency difference tests

Groups	n	Mean Trust Δ	Std. Dev.	t-value	p-value	Support
Using Attention measured by survey:						
Low Attention (1-3)	609	0.36	0.48	-5.03	0.000	Assumption 1 is supported.
High Attention (5-7)	2185	0.47	0.45	.	.	
Using Attention measured by Subject response time:						
Low Attention (1-3)	1449	0.39	0.49	-4.14	0.000	Assumption 1 is supported.
High Attention (5-7)	1447	0.46	0.50	.	.	
Using Attribution measured by survey:						
Low Attribution(1-3)	1252	0.35	0.48	-9.26	0.000	Assumption 2 is supported.
High Attribution(5-7)	1324	0.53	0.50	.	.	
Using Attribution measured by Subject response time:						
Low Attribution(1-3)	1440	0.37	0.48	-5.87	0.000	Assumption 2 is supported.
High Attribution(5-7)	1438	0.48	0.50	.	.	

In Table 6a, low risk Apple Siri has lower average attention level than high risk Google Driverless car, Ford SYNC and in one case, TomTom GPS. This supports

Extension 1 in 5 of 6 cases. In Table 6b, Siri also has lower attribution levels (at least $p < .10$) in 4 of the 6 tests. Extension 1 also proposes that trust change will be higher in higher risk situations. Table 7 shows that this is consistently supported.

Table 6a. Extension 1 Two-group mean difference tests for attention

Using Attention measured by survey:						
	n	Attn. Mean	Std. Dev.	t-value	p-value	Support
Apple Siri Group	1051	4.75	1.76	-2.64	0.008	Extension 1 is supported.
Google Car Grp.	1011	4.95	1.73	.	.	
Apple Siri Group	1051	4.75	1.76	-2.94	0.003	Extension 1 is supported.
TomTom Group	643	5.00	1.67	.	.	
Apple Siri Group	1051	4.75	1.76	-3.25	0.001	Extension 1 is supported.
Ford SYNC Grp.	644	5.04	1.75	.	.	
Using Attention measured by Subject response time:						
Apple Siri Group	1064	3.79	1.99	-4.51	0.000	Extension 1 is supported.
Google Car Grp.	1016	4.18	1.96	.	.	
Apple Siri Group	1064	3.79	1.99	-1.26	0.207	Extension 1 is Not supported.
TomTom Group	648	3.92	2.05	.	.	
Apple Siri Group	1064	3.79	1.99	-3.54	0.000	Extension 1 is supported.
Ford SYNC Grp.	648	4.14	2.01	.	.	

Table 6b. Extension 1 Two group mean difference tests for attribution

Using Attribution measured by survey:						
	n	Attr. Mean	Std. Dev.	t-value	p-value	Support
Apple Siri Group	1043	3.66	1.84	-1.50	0.133	Extension 1 is
Google Car Grp.	1002	3.78	1.87	.	.	Not supported.
Apple Siri Group	1043	3.66	1.84	-5.62	0.000	Extension 1 is supported.
TomTom Group	642	4.19	1.97	.	.	
Apple Siri Group	1043	3.66	1.84	-2.44	0.015	Extension 1 is supported.
Ford SYNC Grp.	645	3.89	1.94	.	.	
Using Attribution measured by Subject response time:						
Apple Siri Group	1052	3.89	2.06	-2.59	0.010	Extension 1 is supported.
Google Car Grp.	1012	4.12	1.93	.	.	
Apple Siri Group	1052	3.89	2.06	-0.19	0.852	Extension 1 is Not supported.
TomTom Group	644	3.91	2.00	.	.	
Apple Siri Group	1052	3.89	2.06	-1.92	0.055	Extension 1 is supported.
Ford SYNC Grp.	647	4.08	2.01	.	.	

Note: Low risk represented by Apple Siri; High risk represented by Google, TomTom, Ford

Table 7. Extension 1 Two-group mean difference tests for trust change Magnitude and Frequency

	n	Mean Magnit. Δ	Std. Dev.	t-value	p-value	Support
Apple Siri	1064	0.64	1.26	-4.58	0.000	Supported
Google Car	1016	0.91	1.40	.	.	
Apple Siri	1064	0.64	1.26	-9.74	0.000	Supported
TomTom GPS	648	1.46	2.21	.	.	
Apple Siri	1064	0.64	1.26	-2.89	0.004	Supported
Ford SYNC	648	0.83	1.38	.	.	
	n	Mean Frequ. Δ	Std. Dev.	t-value	p-value	Support
Apple Siri	1064	0.34	0.47	-6.65	0.000	Supported
Google Car	1016	0.48	0.50	.	.	
Apple Siri	1064	0.34	0.47	-7.03	0.000	Supported
TomTom GPS	648	0.51	0.50	.	.	
Apple Siri	1064	0.34	0.47	-3.27	0.001	Supported
Ford SYNC	648	0.42	0.49	.	.	

To test extension 2, we use Group 1 as the negative event group, and group 2 as the positive. Table 8a and 8b show that usually the negative event group does not produce higher average change in trust than the positive event group. Thus, extension 2 is mostly not supported. Table 9 shows the trust change results—again mixed. Overall, extension 2 is only supported in twelve of twenty-four tests (Tables 8-9). However, with extremely neg-/positive events, extension 2 is supported in five out of six tests.

Table 8a. Extension 2 Two-group mean difference tests for Attention

Events	n	Attn. Mean	Std. Dev.	t-value	p-value	Support
Using Attention measured by survey:						
Negative	1679	4.96	1.78	1.413	0.158	Not supported
Positive	1670	4.87	1.68	.	.	
Extremely Negative	837	5.10	1.86	1.153	0.249	Not supported
Extremely Positive	827	5.00	1.64	.	.	
Moderately Negative	420	4.97	1.63	-1.192	0.233	Not supported
Moderately Positive	422	5.10	1.55	.	.	
Slightly Negative	422	4.67	1.74	2.205	0.028	Supported
Slightly Positive	421	4.40	1.81	.	.	
Using Attention measured by Subject response time:						
Negative	1688	3.96	2.00	-1.04	0.298	Not supported
Positive	1688	4.03	2.01	.	.	
Extremely Negative	844	3.80	1.94	3.26	0.001	Supported
Extremely Positive	844	3.50	1.95	.	.	
Moderately Negative	422	3.68	2.05	-9.66	0.000	Not supported (wrong direction)
Moderately Positive	422	4.95	1.75	.	.	
Slightly Negative	422	4.56	1.95	2.65	0.008	Supported
Slightly Positive	422	4.20	2.00	.	.	

Table 8b. Extension 2 Two-group mean difference tests for Attribution

	n	Attr. Mean	Std. Dev.	t-value	p-value	Support
Using Attribution measured by survey:						
Negative	1672	3.96	1.92	3.56	0.000	Supported
Positive	1660	3.72	1.88	.	.	
Extremely Negative	833	4.18	2.00	1.87	0.061	Marginally supported
Extremely Positive	822	4.00	1.89	.	.	
Moderately Negative	419	4.00	1.79	3.77	0.000	Supported
Moderately Positive	419	3.55	1.71	.	.	
Slightly Negative	420	3.48	1.83	0.91	0.362	Not supported
Slightly Positive	419	3.36	1.93	.	.	
Using Attribution measured by Subject response time:						
Negative	1682	3.98	1.99	-0.44	0.660	Not supported
Positive	1673	4.01	2.01	.	.	
Extremely Negative	839	3.79	1.96	2.73	0.006	Supported
Extremely Positive	829	3.53	1.97	.	.	
Moderately Negative	421	3.78	2.07	-8.95	0.000	Not supported (wrong direction)
Moderately Positive	422	4.98	1.81	.	.	
Slightly Negative	422	4.56	1.85	4.33	0.000	Supported
Slightly Positive	422	4.00	1.94	.	.	

Table 9. Extension 2 Two-group mean difference tests for Trust Change

Events	n	Mean Magnit. Δ	Std. Dev.	t-value	p-value	Support
Negative	1688	0.95	1.53	1.24	0.217	Not supported
Positive	1688	0.89	1.62	.	.	
Extremely Negative	844	1.04	1.65	3.13	0.002	Supported
Extremely Positive	844	0.79	1.61	.	.	
Moderately Negative	422	0.83	1.38	-2.94	0.003	Not supported (wrong direction)
Moderately Positive	422	1.14	1.64	.	.	
Slightly Negative	422	0.90	1.41	0.73	0.463	Not supported
Slightly Positive	422	0.82	1.59	.	.	
	n	Mean Frequ. Δ	Std. Dev	t-value	p-value	
Negative	1688	0.45	0.50	3.10	0.002	Supported
Positive	1688	0.40	0.49	.	.	
Extremely Negative	844	0.47	0.50	5.09	0.000	Supported
Extremely Positive	844	0.35	0.48	.	.	
Moderately Negative	422	0.44	0.50	-2.83	0.005	Not supported (wrong direction)
Moderately Positive	422	0.53	0.50	.	.	
Slightly Negative	422	0.45	0.50	1.96	0.051	Marginally supported
Slightly Positive	422	0.38	0.49	.	.	

Table 10 shows the results for extension 3, that illusionary partners' trust change size and frequency will be lower. We operationalize illusionary partners as those who have high loyalty to a technology vendor because high loyalty is what illusionary effects involve [8]. This extension implies highly loyal subjects will have a more positive magnitude of trust change than will much less loyal subjects. The results for the total set of trust changes support assumption 3. Table 10's first four rows show that with the most extreme groups, those 1/7 loyal versus those 7/7 loyal, the low loyalty group had a mean trust change of -1.74, while the high loyalty group had a mean trust change of 0.03. The fact that average trust change for this high group was positive indicates the possibility of illusionary effects. The bottom three most loyal groups (scores of 1-3) also had more negative trust change means (-1.19) than the highest three (scores of 5-7) loyal groups (-0.05).

Table 10. Extension 3 Two-group mean difference tests for Trust Change

Group	n	Mean of Total Trust Change	Standard Deviation	t-value	p-value	Support
Low Loyalty(1)	69	-1.74	2.65	-4.61	0.000	Supported
High Loyalty (7)	116	0.03	2.46	.	.	
Low Loyalty(1-3)	159	-1.19	2.75	-4.22	0.000	Supported
High Loyalty (5-7)	216	-0.05	2.47	.	.	
+/- Breakdown:						
	n	Mean of Negative Trust Change	Standard Deviation	t-value	p-value	
Low Loyalty(1)	69	-5.71	5.42	-4.04	0.000	Supported
High Loyalty (7)	116	-3.09	3.41	.	.	
Low Loyalty (1-3)	159	-4.74	4.93	-3.18	0.002	Supported
High Loyalty (5-7)	216	-3.30	3.80	.	.	
	n	Mean of Positive Trust Change	Standard Deviation	t-value	p-value	
Low Loyalty (1)	69	3.97	4.70	1.53	0.127	Not supported
High Loyalty(7)	116	3.12	2.84	.	.	
Low Loyalty (1-3)	159	3.55	4.32	0.74	0.459	Not supported
High Loyalty (5-7)	216	3.25	3.32	.	.	
	n	Mean of Total Frequency	Standard Deviation	t-value	p-value	
Low Loyalty (1)	69	0.46	0.30	1.21	0.229	Not supported
High Loyalty (7)	116	0.41	0.29	.	.	
Low Loyalty (1-3)	159	0.43	0.30	0.38	0.704	Not supported
High Loyalty (5-7)	216	0.42	0.30	.	.	

To explore this more, we split the sample by the positive versus negative changes. The next four rows report that the same illusionary effect holds for negative changes. Trust change was significantly more negative (-5.71, -4.74) for the low loyalty groups than for the high loyalty groups (-3.09, -3.30). However, the next four rows show that this was not true for the positive trust changes; they were not significantly different, though they lie in the right direction. Hence, only the negative changes show the illusory effect. In this way, subjects process negative and positive differently. Frequency test results (Table 10, last 4 rows) did not support Extension 3.

We did one more analysis for Extension 3, which predicts loyal subjects will reinterpret negative events such that trust goes up. We compared the loyalty = 1 and 7 groups to see if the loyalty = 7 group, when seeing a negative news brief, changed their trust positively more than did the loyalty = 1 group. We found some evidence of this, per Table 11. Subjects in the high loyalty group were 15% likely to change their beliefs positively, versus only 9% in the low loyalty group. High loyalty subjects also were less likely to have negative change (28%) than those with low loyalty (37%).

The frequency of change is also part of the proposed IPM. Frequency here is measured by the number of changes in trust that have occurred divided by the number of events that have occurred [6]. Trust change frequency should be lower than the incremental growth model (IGM) suggests (see [6]). The IGM posits trust changes in small, frequent steps. Table 12 shows the trust change magnitude/frequency in our data. Note: 80% of changes were < 2, but a “long tail” of large changes also resulted. Finally, in keeping with the overall intent of the Information Processing Model of trust progression [6], we compare the average magnitude and frequency of trust change to those they simulated. Table 13’s first line shows the magnitude and frequency based on this study’s data, and the second and third rows show the IGM and Baseline IPM simulation figures. Although our data is best compared with the Baseline model figures, we also show for comparison purposes the other simulated results in McKnight et al. [6]. As expected, our data’s average magnitude of trust change (0.092) is greater than that that of the IGM (0.023). However, our data’s trust change magnitude is also higher than that of McKnight et al.’s [6] simulated baseline model (0.049) or even the magnitude in the high risk model (0.056). The frequency of trust change column shows that our data’s frequency (0.429) is less than that of the incremental model (0.835). However, it is much higher than that for the baseline model (0.097). Rather, it is closer to the frequency rate for the high risk model (0.526). This may be because we gave all subjects a chance to change their trust even when they had very low attention/attribution.

Table 11. Extension 3 Trust Change After A Negative Event

Trust Change	Low Loyalty Group (Loyalty =1)		High Loyalty Group (Loyalty =7)	
	Frequency	Percentage	Frequency	Percentage
Negative	102	37 %	132	28 %
No Change	150	54 %	264	57 %
Positive	24	09 %	68	15 %
Total	276	100 %	464	100 %

Table 12. Absolute Trust Change after a News Brief by Change Amount

Change Amount	0	1	2	3	4	5	6	7	8	9	10	n
Frequency	1929	769	305	138	79	70	32	17	13	13	11	3376
Percent	.57	.23	.09	.04	.02	.02	.01	.005	.004	.004	.003	

Table 13. Trust Change Averages by Model

Sample size = 1000	Magnitude	Frequency
Empirical Data	0.092*	0.429
Incremental Growth Model (IGM)	0.023	0.835
Baseline IPM Model (Assumes Low Risk)	0.049	0.097
For further comparison:		
Negativity Asymmetry Model	0.044	0.124
Medium Risk Model	0.055	0.146
High Risk Model	0.056	0.526
Illusion Model	0.054	0.096
* Our data's trust change magnitude was 0.92, but this was rescaled to match the 0.1-1.0 scale of the simulation.		

5 Discussion

5.1 Key Findings

The findings of this paper help us understand trust progression better. But they only constitute a start to this research stream. First, we find (Table 13) that, as the IPM posits, trust changes in larger increments than the IGM suggests it should. We also find that trust changes less frequently than the IGM suggests. Thus, trust researchers and practitioners should be careful not to adopt the idea (e.g., [1]) that person-to-technology trust always increases/decreases slowly, in small steps. In fact, what we find is that trust in little-known technologies can start at moderate-to-high levels, as our average starting trust level (6.7/10) shows. From there, not every event causes trust to change. Indeed, we found 57% of events caused no trust change at all. When trust did change, 53% of changes were one unit, while 47% were two or more units. The weighted average trust change of the 47% was 2.14. This calls into question the tiny but continuous changes the IGM predicts. Rather, trust changes are larger and less frequent than IGM predicts. It also shows that the Baseline IPM is too low in predicting trust change frequency. But this may depend on technology newness.

Second, we find both attention and attribution are involved in trust change. We find (Table 5A) when attention is low, the size of trust change is low (0.71), and that high attention produces larger trust changes on average (1.02). The same holds for low and high attribution. We see (Table 5B) that when attention is low, the frequency of change is low (0.36), while high attribution produces larger trust change frequency

(0.47). Thus, assumptions 1 and 2 help us understand how trust change operates. However, the IPM suggests more is involved in trust change: the threshold decision.

Third, we learn risk is important for modeling how trust changes. In fifteen out of eighteen tests, the level of risk mattered. Higher situational risk consistently caused larger and more frequent trust change. Usually, it also caused higher attribution and attention to the events. Fourth, our results were mixed regarding event negativity. We need to study our patterns of positive/negative events to see if that makes a difference. We also need to look for theory that could explain our results. One exception: for highly negative events, the results support the model more consistently. This may mean that negative events make a difference only when they are highly negative.

Fifth, we found solid support for a higher magnitude of trust change when subjects have high loyalty to a vendor. This implies tech vendors who build loyal customer bases will make their customers more resilient to large negative changes in trust level.

An overall pattern emerges from our study. In general, the results support the findings and resulting trust theories of social psychologists and sociologists like Holmes and Luhmann who suggest human trust change is subject to human attention/attribution imperfections. The IPM helps explain why many empirical results have not followed the simple pattern of the IGM.

5.2 Future Research Implications

Much more empirical work is needed to test the IPM. First, future research could test the judgment gear and the trend aspect of the model. Also, empirical longitudinal researchers could compare their trust progression patterns against the patterns McKnight et al.'s [6] Table 2 depicts. Second, additional factors (beyond our three extensions) that affect the trust mechanisms should be theorized and examined. That is, the IPM enables trust theorists to develop new theory because our approach allows them to concentrate—not on what factors lead to trust—but on what factors affect the cognitive mechanisms behind trust changes. This should become a fruitful domain.

Third, IPM opens ways to research trust stability, which is important in outsourcing and e-commerce relationships. Although the IPM discussion has recognized a difference between how trust develops at first and how it may become more stable later, the model may or may not have captured all the mechanisms that affect when trust becomes stable and how stable trust may become over time. Fourth, mathematical models could be developed to more elegantly represent variants of the models we simulate here. The model assumes events valued within normal distribution parameters. Future research could take into account the fact that some events can be extremely positive and others extremely negative. Fifth, the differences in impact between first-hand experiential events and second-hand or signaled events should be studied. Therefore, it does not accurately gauge how much effect outlier events may have on trust levels, on trust stickiness, and so forth. Future research should address this. Finally, the model should incorporate the structurally supporting effects of levels of institution-based trust.

5.3 Limitations and Future Research

One limitation is the use of student subjects evaluating unfamiliar technology, which limits the results' generalizability. Future research should use other sample types. More data needs to be collected to test the robustness of IPM. New attribution and attention measures should be developed to test the results' robustness. Another limitation is the use of a 50-50 mix of positive/negative events. This may not reflect reality and should be adjusted. The model is a general one and will not apply to a specific situation. However, empirical analysis of the trust situation will yield information that enables researchers to tailor the model to the context. Also, the sample is not random.

6 Conclusion

This study provides one data point showing how trust in a technology changes over time. It shows reasonable support for most of the IPM's assumptions and extensions. We find that attention, attribution, situational risk, and vendor loyalty influence trust change. Thus, human imperfections influence trust change. For change magnitude, we find it closer to the IPM model's predictions [6] than to the incremental growth model (IGM). For frequency, it lies almost equi-distance between the IGM and the IPM. More work is needed to clarify how attention and attribution operate to change trust.

References

1. Blau, P.M.: *Exchange and Power in Social Life*. John Wiley & Sons, New York (1964)
2. Kramer, R.M.: Divergent Realities and Convergent Disappointments in the Hierarchic Relation: Trust and the Intuitive Auditor at Work. In: Kramer, R.M., Tyler, T.R. (eds.) *Trust in Organizations: Frontiers of Theory and Research*, pp. 216–245. Sage, Thousand Oaks (1996)
3. Lee, J., Moray, N.: Trust, Control Strategies and Allocation of Function in Human-Machine Systems. *Ergonomics* 35, 1243–1270 (1992)
4. Mayer, R.C., Davis, J.H., Schoorman, F.D.: An Integrative Model of Organizational Trust. *Academy of Management Review* 20, 709–734 (1995)
5. McKnight, D.H., Choudhury, V., Kacmar, C.: Developing and Validating Trust Measures for e-Commerce: An Integrative Typology. *Information Systems Research* 13, 334–359 (2002)
6. McKnight, D.H., Liu, P., Pentland, B.T.: How Events Affect Trust: A Baseline Information Processing Model with Three Extensions. In: Dimitrakos, T., Moona, R., Patel, D., McKnight, D.H. (eds.) *IFIPTM 2012. IFIP AICT*, vol. 374, pp. 217–224. Springer, Heidelberg (2012)
7. Mohr, L.B.: *Approaches to Explanation: Variance theory and Process Theory*. *Explaining Organizational Behavior*, ch. 2, pp. 35–70. Jossey-Bass, San Francisco (1982)
8. Murray, S.L., Holmes, J.G.: A Leap of Faith? Positive Illusions in Romantic Relationships. *Personal and Social Psychology Bulletin* 23, 586–604 (1997)
9. Sitkin, S.B., Weingart, L.R.: Determinants of Risky Decision-Making Behavior: A Test of the Mediating Role of Risk Perceptions and Propensity. *Academy of Management Journal* 38, 1573–1592 (1995)