The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself

David Eil
UC San Diego

Justin M. Rao*
UC San Diego

January 12, 2010

Abstract

This paper examines how individuals' beliefs respond to objective information about their ranking on a neutral quality – a meaningless number on a card – or on a quality that has a significant self-image component – intelligence or beauty. For favorable news in the image tasks, subjects respected signal strength and update as "optimistic Bayesians," but they heavily discounted and largely ignored signal strength in processing unfavorable news, leading to noisy posterior beliefs nearly uncorrelated with Bayesian inference. None of these patterns were observed in the control. Subjects did not display confirmatory bias, but the results can explain its root cause – disconfirming signals are treated as "bad news" due to the self-esteem loss of realizing one was initially wrong.

*deil@ucsd.edu, jmrao@ucsd.edu. UC San Diego, 9500 Gilman Dr., La Jolla, CA 92093-0508. We would like to thank our advisors Vincent Crawford and James Andreoni for continued guidance, the Russell Sage Foundation Behavioral Economics Small Grant Program for financial support and Nageeb Ali, Joel Sobel and Charles Sprenger for helpful comments.

1 Introduction

Many studies have shown that people have a "psychological defense mechanism" to help maintain positive self-image and confidence. Classic examples are selective memory, biased attributional judgment (good outcomes attributed to skill, bad outcomes attributed to luck) and cognitive dissonance. These strategies accentuate the positive even while, on some level, recognizing the negative. Three related economic concepts, confirmatory bias, self-serving bias and overconfidence, are all rooted in filtering out unfavorable information. A common thread is that all these behaviors hinge on new information being processed differently depending on whether it is favorable or unfavorable to existing beliefs.

However, the literature on how new information is incorporated into existing beliefs (signal processing) has generally shied away from using signals that are inherently good news or bad news. While prominent findings show that people systematically depart from the predictions Bayesian rationality – examples include the "law of small numbers" (Tversky and Kahneman 1974) and base rate neglect (Grether 1980) – the results are difficult to apply to an environment in which beliefs are directly related to self-image and esteem. This is because the experimental protocols typically use an underlying quantity far-removed from self-image, such as the color of balls in an urn.

This paper aims to bridge the gap between these two literatures. We study how people incorporate objective pieces of new information into their existing beliefs when the underlying quantity has direct psychological importance. In this environment signals have intrinsic valence – news is good or bad. The qualities used to generate signal valence were intelligence (*IQ*), as measured by score on an IQ test, and physical attractiveness (*Beauty*), as rated by subjects of the opposite sex during speed dating. As a control, subjects also updated their beliefs under the same information structure on the number (1-10) on a card given to them in a sealed envelope. All subjects participated in one image task and the control task in an experimentally balanced order. For each task, subjects first revealed the distribution of their prior beliefs through a computer interface. Then, they received three rounds of up/down signals. Each signal was an anonymous pair-wise comparison of their rank (out of 10 subjects in their session) to a randomly selected anonymous subject. For *Beauty* there were 10 subjects of each gender and all comparisons were made within gender.

Our main finding is that subjects incorporated favorable news into their existing beliefs in a fundamentally different manner than unfavorable news. For favorable news in the image tasks, subjects tended to respect signal strength and to adhere quite closely to the Bayesian benchmark (albeit with an optimistic bias). In contrast, subjects discounted or ignored signal strength in processing unfavorable news, which lead to noisy posterior beliefs that were nearly uncorrelated with Bayesian inference. These patterns were not observed in the control. For males the bias was strongest for *IQ*, while for females the bias was strongest for *Beauty*. We call this finding the good news-bad news effect.

The differential processing we observe indicates that bad news has an inherent "sting" that selective filtering helps ameliorate. For perfectly informative signals, it is not possible to apply the filtering ex-post, but it may be possible to avoid the signal ex-ante. As perfectly informative signals become more likely to be unfavorable, we expect higher degrees of avoidance. To examine signal avoidance we elicited subjects' willingness-to-pay (WTP) to learn their true rank through a price list equivalent to the Becker-DeGroot-Marchak mechanism. Consistent with our hypothesis, subjects who had received good news in the image tasks were willing to sacrifice some of their earnings to learn their true rank whereas subjects who had received bad news required a subsidy. Since the probability of learning one's true rank increased linearly with WTP, requiring a large subsidy made it highly unlikely or impossible that one's true rank would actually be revealed. In the control WTP did not vary across signal valence.

Our results help explain the underlying cause of confirmatory bias (CB). CB occurs when signals that agree with one's prior belief are incorporated into posterior beliefs in a different manner than signals that disagree with one's prior. Typically this involves underweighting disconfirming evidence and sometimes overweighting confirming evidence.[1] After viewing an equal number of confirming and disconfirming signals (neutral evidence) a CB-agent will view it as supporting the prior belief.

Evidence for CB comes mainly from the psychology literature. Lord et al. (1979) found that subjects, pre-screened for having strong views on the death penalty, who were given the same articles on the preventative value of capital punishment strengthened their conviction regardless their original stance. That is, both sides of the debate became more convinced they were right by the same evidence. Plouss (1991) used partisans in the nuclear safety debate and replicated Lord et al.'s finding. Mahoney (1977) sent a paper to journal referees with the results section altered to either agree or disagree with the referee's published results. He found that referees for which the result was disconfirming were far more critical of the methods section of the paper than referees for which the result was confirming, despite the fact that the method section was identical for both groups. Nickerson (1988) reviews the extensive evidence for CB and Rabin and Shrag (1999) highlights its economic importance .

Subjects in our experiment did not display CB. For instance, a subject who reported believing she had below average beauty tended to overreact to a signal that informed her she

[1] Rabin and Shrag (1999) presents a theoretical model of CB that uses this definition as a starting point.

was more attractive than a randomly selected comparison subject ("up") and underreact to the opposite signal ("down"). That is, "down," the signal which confirmed her prior, was given less weight than dictated by Bayesian inference. This is precisely the inverse of CB. Initially this result might appear surprising given past findings of CB. However, a careful survey of the literature reveals that CB presents itself when the underlying quantity is not directly tied to self-image. For example, whether or not the death penalty prevents violent crime does not make one feel good or bad intrinsically (unless, of course, one is on death row). Rather it has psychological importance because it is damaging to the ego to admit one was wrong on an issue they care about. Supporting this intuition, CB experiments generally use partisans' beliefs on hot-button debates. A feature of this protocol is that confirming signals are always good news – confirmation and valence are perfectly co-linear. In our design, confirming signals can be good or bad news. We find that what matters is valence, not agreement or disagreement with prior beliefs.

We do not dispute that CB is an empirical regularity. However, our results indicate that agreement with prior beliefs is not the underlying cause. We should only expect to observe CB when disconfirming signals are indeed interpreted as bad news — when being right is what matters most, CB is exhibited. This is exactly the structure of past experiments finding CB. An intuitive feature of our signal valence based explanation of CB is that CB is brought into the fold of the psychological defense mechanism, rather than just being a mechanical updating error.

The good news-bad news effect provides an explanation for self-serving bias. Self-serving bias occurs when self-interest shapes what one believes is fair. For instance, in laboratory bargaining games both sides tend to believe strongly their side is in the right even when they are given identical information (see Babcock and Loewenstein (1997) for a review). Information endorsing a social or moral norm that leads to material benefit has positive valence. Accordingly, it receives more weight in the updating process and self-serving bias naturally develops.

Another application of our results is to the empirical finding of widespread overconfidence in self-image related qualities (see Moore and Healy (2008) for a review).[2] Since strong favorable signals induce a large change in beliefs but strong unfavorable signals are heavily discounted, resultant beliefs are biased towards positive self-image and overconfidence develops. Furthermore, the pattern of information acquisition we observe can further fuel overconfidence as people likely to receive good news are the ones seeking out more precise signals.

[2] See also Alicke and Govorun (2005).

The question, however, remains: why would a bias such as the good news-bad news effect become part of our psychological constitution? Why is belief accuracy sacrificed in favor of irrational optimism? Many authors have grappled with this question when discussing the psychological defense mechanism more generally. They have argued that the economic and evolutionary value of self-confidence and self-esteem can potentially outweigh the associated costs.[3] Benabou and Tirole (2002) derive conditions on when so-called biases will emerge in a fully rational model. They show, for instance, that overconfidence can help agents commit to initially costly long-term projects that otherwise might not be undertaken. Compte and Postlewaite (2004) show that if self-confidence enhances performance, then rational agents will choose updating rules that are optimistically biased away from Bayesian inference. Brunnermeier and Parker (2005) modify preferences to include "anticipatory utility" and establish a similar result. These papers provide a solid foundation for the existence of valence-dependent updating in human psychology and our results support these models. The models can also be employed to help explain CB. Earlier we argued that CB can be explained by signal valence — these models show that under this explanation CB can result from a fully rational model, while previous models of CB, such as Rabin and Shrag (1998) had to assume it as a "behavioral bias."

The remainder of the paper proceeds as follows. Section 2 details the experimental design. Section 3 provides the inference results. Section 4 provides the information search results. A discussion follows in Section 5. Section 6 concludes.

2 Experimental Design

The experiment was conducted at the University of California San Diego Economics Laboratory. There were 7 *IQ* and 4 *Beauty* sessions. *IQ* sessions had 10 subjects each. *Beauty* sessions had 10 male subjects and 10 females subjects. Potential participants were solicited through an online subject database and were told only that the experiment would last 1.5 hours and earnings would be $25 on average. Given the sensitive nature of the experiment (receiving information on intelligence and physical attractiveness) we over-subscribed the sessions to account for people electing not to participate after reading the IRB consent form.[4] Across all sessions, only 2 subjects elected to leave after learning the nature of the ex-

[3]Inference mistakes can indeed have economic consequences. Examples include excessive trading in stock markets (Barber and Odean 2001) and overconfidence in physicians' diagnostic decisions (Christensen-Szalanski and Bushyhead 1988).

[4]Upon entering the lab subjects were told to carefully read the consent form which told them that they would receive feedback on either their intelligence or physical attractiveness. At that point they were given an opportunity to discreetly exit if they no longer wanted to participate.

periment. As such, it was necessary to randomly select subjects to leave to pare the sessions down the required number of subjects.

The experiment proceeded in three stages. The first stage of both session types collected the necessary information to rank the subjects on the intelligence or attractiveness. At this juncture subjects were not told why this information was being collected. In *IQ*, subjects took a 25 question IQ test. The questions were taken from a standard Wechsler Adult Intelligence Scale test and involved logic, spatial and verbal reasoning, and general knowledge. Subjects were told the "aims of the experiment depend on you making an honest effort on this IQ test." To incentivize them further 1 question was chosen at random (ex-post) and if they answered this "payment question" correctly $5 was added to earnings.

In *Beauty*, subjects engaged in a speed dating exercise.[5] Each person met 5 subjects of the opposite sex and engaged them in a 4 minute conversation. These meetings were face-to-face with a partition separating each pair. The only restriction placed on the conversations was that they could not reveal identifying information such as their full name or place of residence. To help break the ice, a conversation topics were suggested. During the meetings soft music was played to mimic a real speed dating environment. Most subjects engaged in lively conversation and seemed to enjoy the exercise.

After each meeting, subjects filled out a "speed dating questionnaire" rating their conversation partner (scale 1-10) on three dimensions: friendliness, attractiveness and ambition. These forms were kept in a manila envelope during the meetings and were filled out at separate partitioned work stations to maintain anonymity and encourage honest assessment. Ambition and friendliness were included to reduce the anxiety of filling out the questionnaire. The average of the 5 physical attractiveness ratings were used to give each subject a beauty score. After the speed dating exercise, these sessions proceeded as two separate 10 person sessions of each gender.

While stage 2 instructions were read, the scores for the image task in stage 1 were tabulated. The scores ranked the subjects 1 (highest score) through 10 (lowest). Rankings were within gender group for *Beauty*) and ties were broken by a flip of a coin. Subjects were informed of the ranking procedure and understood that the scores gave a strict ordering.

At the beginning of stage 2, subjects were passed a sealed envelope containing a card with a unique number 1-10 written on it. In *Beauty* sessions the numbers were unique within gender group. The number determined rank on the "card task" (*Control*) and again subjects were explained the ranks gave a strict ordering. They were also truthfully told that card rank was determined randomly.

[5]The design of the speed dating exercise is similar to Fisman et al. (2006) which used speed dating ratings to examine gender differences in mate selection.

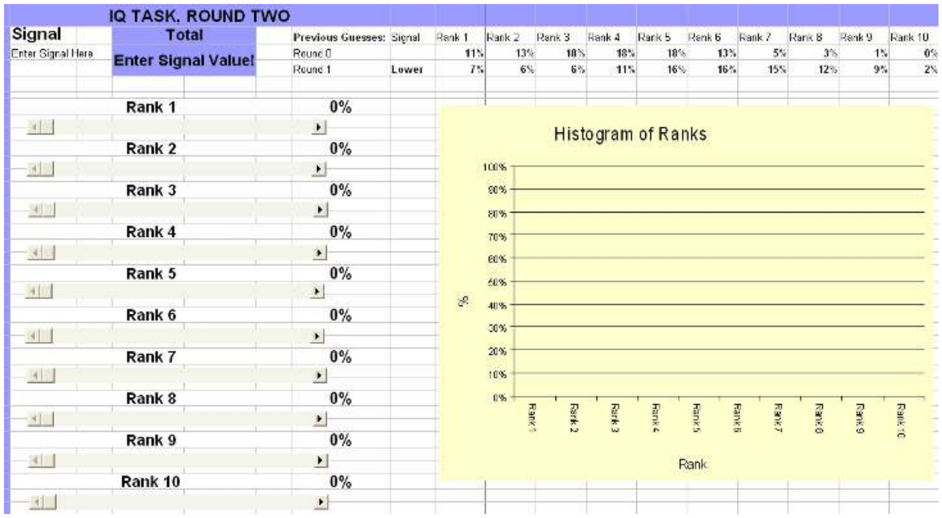

Figure 1: Instrument screen-shot.

Stage 2 consisted of 2 sets (image and control) of 4 rounds each, enumerated 0 through 3. In round 0, subjects entered their prior belief in percentage terms that they occupied each of the 10 ranks on the task. This was done through a computer interface as seen below in Figure 1.

The percentages always started at 0 for each rank and a graphical representation was provided through a auto-adjusting histogram. The total was given and had to add to 100% to be eligible for payment. Adjustments were made by clicking the arrows or moving the slider-boxes.

Following rounds 0, 1 and 2 each subject's rank was compared to a randomly selected comparison participant (again these were done within gender group for *Beauty* sessions). This comparison participant was anonymous and the comparisons were bi-lateral and unique each round (drawn without replacement). The result of the comparison was conveyed via a message that read either "You are ranked higher" or "You are ranked lower." The message cards were labeled by round and subject ID letter. As the messages were handed out, subjects were reminded each time that "higher" meant closer to the top rank of 1 and "lower" meant closer to the bottom rank of 10. After receiving the message, subjects entered it into their spreadsheet to ensure the message was accurately received. They then entered their new guesses over their rank on the task. The entire history of rank guesses and signals were provided at the top of the screen. The order of the image and control sets was randomized.

Payment for rank guesses was done through the incentive compatible quadratic scoring rule (Selten 1998). Following Moore and Heally (2008) subjects were shown the scoring rule formula and told, "Although this formula looks quite complicated, what it means to you is simple. You make the most money on average by honestly reporting your beliefs of the probability you occupy each rank. The formula rewards accuracy in a way that the way to maximize your average winnings is to report honestly." Eliciting beliefs in a laboratory setting indeed a challenge and many authors have highlighted the shortcomings of the quadratic scoring rule. For instance, risk aversion can induce median bias (Holt 1986). However, our protocol is designed to identify the processing differences between good and bad news. That is, we will be using comparisons that condition the elicitation mechanism used.

Following completion of stage 2, subjects were given an opportunity to learn their true rank on both the control and image tasks. This was done separately for each task through a price list equivalent to a discretized version of the Becker-DeGroot-Marschak (BDM) mechanism (Becker, DeGroot and Marschak 1964). This allowed us to capture subjects' willingness to pay for complete information. Each item of the price list had a choice of the following form:

> Do you prefer (A) Receiving $\$x$ and learning your true rank on the [IQ/Beauty or Card] task or (B) Receiving \$0 and not learning your true rank?

x ranged from -7.00 to 7.00 in \$1.00 increments. When $x < 0$ the language was changed to "paying $-x$" (e.g. if $x = -3$ it read "paying \$3.00"). There were 15 choices for each task. A random number was drawn for each task to determine which decision would be executed. Subjects were told they should answer each question as if it were going to be executed.[6] If the executed choice led to the subject learning their true rank on a given task, then they were given an "acknowledgment form" on which they initialed by their rank indicating they indeed learned the truth. In *Control* they also got to open the envelope. To ensure choice anonymity, if the executed choice did not lead to a subject learning their true rank then they received a blank acknowledgement form.

Each session lasted about 1.5 hours and subject earnings averaged \$23.33. Full instructions can be found in the appendix.

3 Inference Results

The purpose of the design was to have objective signals identical across the image and control treatments. This allows for an easy comparison of the patterns of observed in the neutral *Control* to *Beauty* and *IQ* for which subject rank is tied to self-image. Also, conditional on

[6]Note that this procedure is equivalent to the BDM but does not require the lengthy and confusing explanation of why subjects have an incentive to report their true willingness to pay.

priors collected in round 0, the design allows us to calculate closed form Bayesian posteriors for rounds 1-3, which are used as the normative benchmark. Objective signals also have the attractive property that the results of the experiment cannot be explained by subjects simply "inverting" (or misinterpreting) some unfavorable signals.

In the design, we tried to be as ecologically valid as possible. In the image treatments, priors were not imposed rather they were simply the actual prior probabilities the subjects walked into the lab with. The image qualities, physical attractiveness and intelligence, are both economically and socially relevant. The ego utility or status of these qualities was not artificially imposed. Rather it relied on the subjects' underlying preferences. While many signals in the real world are of a subjective nature (compliments, praise in letter of recommendation) objective up/down comparisons are quite common as well.[7]

This section has 3 parts. We first present the analysis of the posterior beliefs, then we examine the round-to-round changes in greater details. Finally, we argue that the gender differences observed support our claim that self-image association is driving the results. 7 of the 150 subjects were eliminated from the analysis because they clearly displayed a lack of understanding of the experimental protocol.[8]

3.1 Posteriors

We label "up" messages $s = 1$ and "down" messages $s = 0$. Recall, $s = 1$ indicates the subject's rank is closer to 1 than the comparison participant. In *Beauty* rank 1 represents the subject whose physical attractiveness rated highest, in *IQ* rank 1 represents the subject with the highest IQ test score and in *Control* rank 1 represents the subject whose random "card number" (1-10) is 1.

The purpose of this section is to isolate how favorable ($s = 1$) and unfavorable ($s = 0$) news is incorporated into posterior beliefs. Our hypothesis is that for the image tasks, subjects will respond more to favorable news and less to unfavorable news. For the control, no difference is expected. For each subject in each round we calculated what their posteriors should be according to Bayes rule, given the priors they reported in round 0. We then took the expected rank under this distribution and called this value the "mean Bayesian belief".

[7] Examining subjective signals is also possible with our experimental framework. Suppose we had elicited subjects' priors in *Beauty* before speed dating and then directly after. The feedback during the conversations is a subjective signal and is of interest to study, however we used objective signals to provide closed form solutions and leave the analysis of subjective signals to future work.

[8] The first reason for elimination was not having beliefs that added up to 100%. The second reason group of subjects eliminated appeared to interpret "up" as meaning closer to 10 (as opposed to closer to 1) but then at some point realized their mistake. This was a function of the somewhat confusing feature of the English language in that "higher" means better in terms of rank in in a distribution but means the opposite in terms of the absolute number. 2 subjects had a computer failure which lead to partial data loss.

Panels A, B and C of Figure 2 plot subjects' observed mean belief by the mean Bayesian belief for *Beauty*, *IQ* and *Control* respectively. The data is split by signal valence. It compares subjects who have received all good news to all bad news. This retains all observations from round 1 (when the subjects have only received 1 signal) and eliminates some observations from rounds 2 and 3; this is done to get the tightest comparison possible. The eliminated observations are analyzed in the following subsection and tell the same basic story.

The fitted lines are from the OLS regressions presented in Table 1. Optimistic updating occurs when the OLS fitted line is steeper for good news as compared to bad news. This demonstrates that subjects are reacting relatively more than the Bayesian for favorable signals and less than the Bayesian for unfavorable signals. A constant slope across signal valence indicates that bad news and good news are incorporated into beliefs symmetrically.

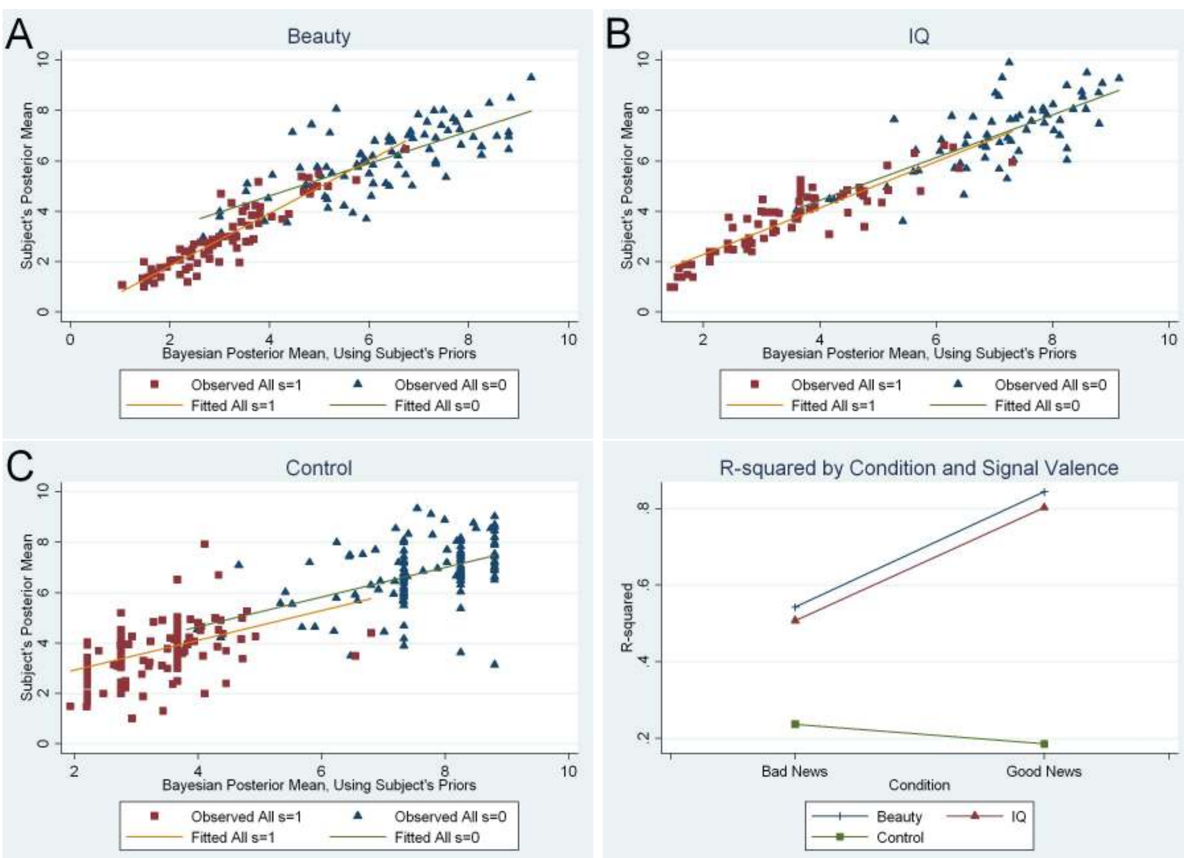

Figure 2: Posterior beliefs by signal valence and condition.

We see that the slope is indeed constant for in *Control* but that is steeper for good news in both image conditions. The effect is more pronounced in *Beauty* and table 1 shows that the difference is highly significant (as indicated by the coefficient on the interaction term). Subject's in *Beauty* engaged in optimistic updating as compared to the control. The difference in *IQ* is not statistically significant, however we see in panel D that subjects did differentially process information in this condition. In particular, subjects responded more predictably to good news, even if with less bias than in the *Beauty* treatment. Their response to bad news was noisier. The following subsection confirms this. The intuition, which is expanded upon in Section 5, is that the pattern of updating discussed in the introduction has the property that whether or not net optimism emerges depends on prior beliefs. In *Beauty* subjects showed more initial overconfidence, this means that on average down signals were more informative than up signals. Not respecting signal strength for unfavorable news leads to more bias in this case.

Table 1: Subject's Mean Belief
as a Function of Bayesian Mean

	Beauty	IQ	Control
Bayesian μ	0.641***	0.846***	0.589***
	(0.0620)	(0.0868)	(0.0862)
Bayesian μ*1{All $s = 1$}	0.406***	0.0714	-0.00118
	(0.0731)	(0.107)	(0.155)
1{All $s = 1$}	-2.311***	-0.580	-0.531
	(0.429)	(0.677)	(0.790)
Constant	2.054***	1.051	2.291***
	(0.412)	(0.642)	(0.666)
Observations	163	139	292
R^2	0.867	0.869	0.717

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panels A and B provide visual evidence that beliefs are far more noisy for bad news as compared to good news for the image treatments. No such patterns emerge in panel C (*Control*) – the intercept shift displayed in Panel C is well within 1 s.e. of 0. Panel D presents the R^2's from the regressions in Table 1 separated by signal valence (i.e. 2 per condition, 1 for s=0 and 1 for s=1). The differences in R-squared are statistically significant (as given by a variance ratio test) for *Beauty* ($p = 0.0000$) and *IQ* ($p = 0.0000$) but not for *Control* ($p = 0.6139$) The result of the R^2 comparison is striking and is visually apparent in Panels A-C. . R^2 was 58% higher for good news in *IQ* and 60% for good news in *Beauty*. For the image treatments (where the notion of good vs. bad news makes sense) subjects' beliefs adhered far more tightly to Bayesian rationality when the news was favorable than when it was unfavorable. This is seen through a slope coefficient closer to 1 and much higher R^2. The control shows that this was not simply an artifact of the experimental application of ranks.

3.2 Round-to-round Changes

The preceding subsection examined how posterior beliefs compared to their Bayesian counterparts. In this subsection we examine how beliefs change from one round to the next upon the arrival of new information. Again, the analysis is designed to isolate differences in information processing by signal valence. Recall a negative change is a changes towards 1, the highest rank.

Figure 3 presents linear fits of the round-to-round changes in mean beliefs as a function of the Bayesian change.[9] The associated regressions are in Table 2. The inclusion of higher order terms does not significantly improve fit and the linear analysis simplifies the comparisons between treatments. Across all conditions and signals, there is a pattern of over-responding to uninformative signals, i.e. the y-intercepts are not zero. The importance of this result will be discussed in the following section.

The 45 degree line is given for comparison to the Bayesian response. For *Beauty* and *IQ*, when the news is favorable (lower left quadrant) the fitted curve is much closer to the 45 degree line and has significantly steeper slope – signal strength is respected. The same does not hold true for unfavorable news. Table 2 shows that it is statistically indistinguishable from 0. Informative down signals are treated in a similar fashion to uninformative ones; beliefs are adjusted down in a noisily.[10] In the introduction we dubbed this pattern the good news-bad news effect. The effect shows up in both the posterior beliefs and the round-to-round changes. For the image tasks, signal valence fundamentally affects how new

[9]Both Figure 3 and Table 2 limit the analysis to changes in the correct direction. There were equal of number of changes in the wrong direction by signal valence. The inclusion of these outliers actually strengthens our argument as 2 were informative down signals in which the subject actually updated slightly updated towards 1. However, these points have such high leverage in the regression and represent less than 3% of the sample. As such they are excluded. Median regression is used in Table 2 to further reduce the influence of outliers.

[10]Our results would be considerably strengthened if we applied a kernel smoothing to prior beliefs. This is because Bayesian updating cannot put any weight on ranks that have prior probability equal to 0. For bad news, this tends to lead to an understatement of the optimistic updating because it was fairly common for a subject to have beliefs 50-40-10 on ranks 1, 2 and 3 respectively. The Bayesian posteriors have support of ranks 1-3, as such the mean does not move very much. However, it was exceedingly rare to see such concentrated (and likely mistaken) beliefs on the low end of the distribution. A smoothing procedure which transformed 50-40-10 to, for instance, 45-35-8-4-3-2-1-1-1 would lead to a much larger required downward change. In reporting beliefs, subjects generally left the possibility open that they were at the top of the distribution, so the smoothing procedure would not affect the analysis of up signals significantly.

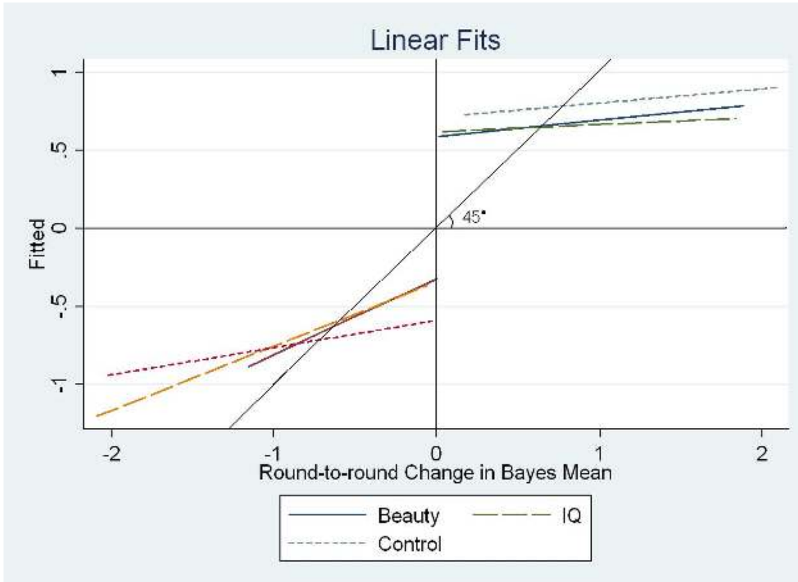

Figure 3: Round-to-round changes in mean belief by condition.

information is incorporated into beliefs. There do appear to be similar patterns in *Control* albeit of a much lower magnitude and lacking statistical significance.

Table 2: Mean Belief Changes by Signal Direction and Condition

Dependent Variable: $\Delta\mu$

Condition	**Beauty**	**IQ**	**Control**
$\Delta\mu_{Bayes}$	0.212**	0.0256	0.141
	(0.0833)	(0.112)	(0.153)
$\Delta\mu_{Bayes} \times 1\{s=1\}$	0.475***	0.540***	0.141
	(0.165)	(0.149)	(0.205)
$1\{s=1\}$	-0.642***	-0.772***	-1.119***
	(0.107)	(0.134)	(0.288)
Constant	0.437***	0.599***	0.625***
	(0.0826)	(0.102)	(0.222)
Observations	206	183	385
R^2	0.49	0.55	0.48

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

In the preceding subsection we showed that good news led to a much tighter adherence to Bayesian rationality than bad news. The following chart compares the standard deviation of the residuals for the regressions in Table 2, by signal valence. We can see the same pattern emerges, updating after down signals (s=0) is far more noisy in *Beauty* and *IQ*, but the essentially the same in *Control*. A variance ratio test gives $p = 0.0000$ for *Beauty*, $p = 0.085$ for *IQ* and $p = 0.347$ for *Control*. Once again the "image effects" in *Beauty* appear strong than *IQ* but are non-existent in *Control*.[11]

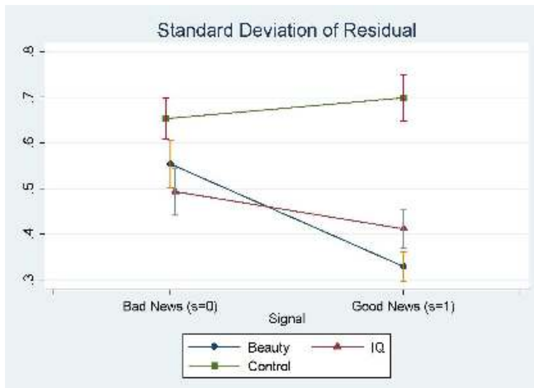

Figure 4: Noise in round-to-round updating by treatment and signal type

3.3 Image association and gender

We have argued that the chief difference between *Beauty*/*IQ* and *Control* is the degree to which the underlying quality is associated with self-image. *Control* uses random assignment of ranks in order to maintain image neutrality.[12] We have seen that differential processing of good versus bad new was most pronounced in *Beauty*. Under our working hypothesis, this would indicate physical attractiveness, as given by peer subjects in a speed dating exercise, has a greater impact on self-esteem for most subjects than intelligence, given by an IQ test.

There are many plausible reasons that this is the case. For one, formal feedback on physical attractiveness is not as common as formal feedback on intelligence. By the time they reach college, students receive many informative signals on intelligence such as SAT score, high school class rank and college GPA. There are no such analogs for appearance. [13] Rank out of 10 subjects on the appearance dimension is potentially more important to

[11]It is worth noting that since social comparisons were used in *Beauty* and absolute comparisons in *IQ* we cannot concluded that physical attractiveness has a stronger image association than intelligence. We elaborate on this point in Section 3.3.

[12]It is possible that some people consider themselves lucky and associate a high rank with having good luck. Also, even the artificial application of rank could have an effect due to an innate responsiveness to relative comparisons. Indeed, there does appear to be a slight effect consistent with this reasoning.

[13]Subjects were asked in the post-questionnaire whether they had ever posted a picture on HotorNot.com or similar website. Only 4% of subjects answered that they had.

self-image than their rank on the intelligence dimension as the rarity of formal feedback on appearance adds to its ego-impact. Another reason is that the "beauty scores" were formulated by subjects in the room while the IQ score used an external standard. Many experiments have shown that subjects care about the opinions of their peers in experimental settings (see for instance Andreoni and Bernheim, 2009).

There is evidence that possessing physical attractiveness is more important for females and possessing intelligence is more important for males in attracting potential mates (Fisman, Iyengar, Kamenica and Simonson 2006).[14] In Table 3 we look at how subjects beliefs changed relative to the Bayesian separated by signal history and gender. For both genders the values are mostly negative for *Beauty* as shown in Figure 2. We see that women were more optimistic in the face of the worst news (all $s = 0$, $t = -1.47$). In contrast, men were more optimistic in the face of the worst news in *IQ* ($t = 1.50$). The results indicate that men have a stronger aversion to believing themselves to be unintelligent, while women exhibit this aversion for physical unattractiveness. The effect only appears at the bottom the distribution. From a mate selection perspective this makes sense if finding mates becomes very difficult at a certain point on the left tail of the distribution.

Table 3: Total Change in Mean Beliefs Relative to Bayesian
$((\mu_t - \mu_0)$-$(\mu_t^B - \mu_0^B)$ by Gender

	Round 1		Round 2		Round 3		All
Signal History							
Signal History	F	M	F	M	F	M	Pooled t test
Beauty Condition							
0 $s = 1$	-0.13	0.14	-0.27	0.20	-0.47	0.02	-1.47
1 $s = 1$	0.04	-0.12	-0.15	-0.07	-0.16	0.01	0.05
2 $s = 1$			0.11	-0.19	-0.35	-0.31	0.90
3 $s = 1$					0.18	-0.17	0.86
IQ Condition							
0 $s = 1$	0.04	-0.28	0.04	-0.19	0.25	-0.44	1.50
1 $s = 1$	0.44	0.22	0.11	-0.28	-0.41	-0.06	0.48
2 $s = 1$			0.13	0.27	-0.09	0.01	-0.63
3 $s = 1$					0.45	0.13	0.68

[14]The Trivers-Willard hypothesis posits that if a couple has traits particularly beneficial to one gender then their offspring should more likely be that gender (Trivers and Willard 1973). Using a large scale survey, Kanazawa (2006) finds that beautiful parents are significantly more likely to have daughters. [15] Based on these findings, it is natural to expect more optimistic updating in *Beauty* for females and *IQ* for males.

4 Search Results

Our results in the last section showed that subjects dealt with the unwelcomeness of bad news by underweighting "down" signals. In this section we examine how people seek out new information as a function of their current beliefs. Here signals are perfectly informative, so they cannot be selectively filtered. Our hypothesis is that the "sting of bad news" leads subjects to avoid learning their true rank as the probability that the revelation is unfavorable increases. Theoretical models have shown how overconfidence can develop and persist if agents acquisition of new information is a function of their prior beliefs. We will briefly discuss two such models and then analyze how subjects' WTP to learn their true rank varied by their final round posterior beliefs.

4.1 Theoretical Models

Santos-Pinto and Sobel (2005) presents a model in which image qualities are multi-dimensional . Agents can augment their skill and do so on the dimensions they believe are most important to the overall quality. While the model does not deal with information search directly, the underlying mechanism is one of egocentric search. In the real-world, people may also receive informational feedback during skill augmentation.

In Köszegi (2006) information acquisition affects utility in two distinct ways: 1) it allows for better decision making 2) it changes beliefs which enter utility directly. Information search thus depends on the payoffs to accurate beliefs and the shape of the ego utility function. Köszegi assumes that ego utility equals 1 when mean beliefs are above the 50th percentile and 0 otherwise. Information can be damaging to agents who currently believe themselves to be better than average – these agents will exhibit rational ignorance. The type of agents exhibiting rational ignorance depends directly on the ego utility function. We will see that a function with high utility of occupying the top 2 ranks and very low utility of occupying the bottom 2 ranks has good explanatory power in our setting.

4.2 Results

As described in the methods section, after the 3 signal rounds subjects had an opportunity to learn their true rank. WTP is defined as the switching point on a price list offering a choice between receiving $\$x$ *and* learning their true rank on the task and receiving \$0 and *not* learning their rank. x ranged from -\$7.00 to \$7.00 in \$1 increments. As such subjects could require as much as \$7 subsidy to learn their true rank and could pay as much \$7. If the executed choice determined that a subject learn his true rank, then he had to initial by the rank on an "acknowledgment form."

Table 4 examines the relationship between WTP and the mean and standard deviation of their final round beliefs (i.e. after receiving all 3 signals). *Control* allows us to account for curiosity and other artifacts of the experimental setting unrelated to self-image.

Table 4: Willingness-to-pay as a Function of Final Round Beliefs

Condition	IQ	Beauty	Control
Final Round μ	-0.325***	-0.206**	0.0112
	(0.107)	(0.0855)	(0.0423)
Final Round σ	0.911**	0.917	0.237
	(0.382)	(0.575)	(0.174)
Constant	0.729	0.298	-0.339
	(0.523)	(0.750)	(0.336)
Observations	77	65	142
R^2	0.099	0.180	0.015

*** p<0.01, ** p<0.05, * p<0.1

Robust standard errors in parentheses

Subjects WTP to learn their true rank increased as their round 3 posterior mean moved towards 1 (highest) for both *IQ* and *Beauty* – for *Control* there was not any effect.[16] In fact, subjects in the image conditions subjects who believed they were below average (rank 6+) often required a subsidy to learn their true rank. In the image conditions there was an "informational value" as well – WTP increased with the standard deviation of the belief distribution.[17]

If most of the ego utility effects are driven by the tails of the distribution then small sample sizes within each condition become a problem. This is why a linear fit was used in Table 4 despite its lack of realism. In Figure 5 we pool *IQ* and *Beauty* and analyze the data using a flexible non-linear form, fractional polynomials.

Figure 5 shows that subjects in the middle of the distribution have a flat WTP function – they are generally not willing to pay but do not require a subsidy to learn their true rank. The function is convex for high ranks and concave for low ranks. The very best subjects had

[16] There were not gender differences across condition.

[17] One might suspect that the informational effect might differ by mean belief. Regressions available from the authors show that there are not significant differences if the interaction terms by mean belief quartile are added.

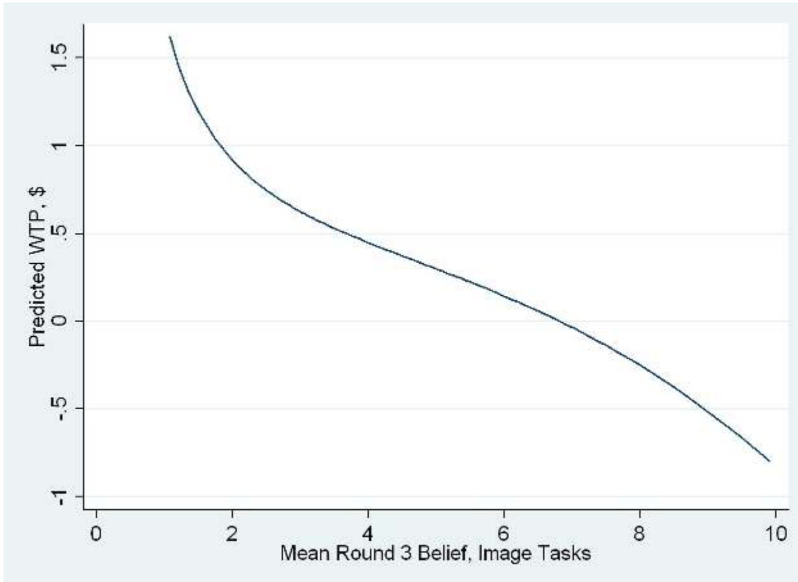

Figure 5: Fractional polynomial fit of WTP on mean final round belief in the image tasks.

a high WTP and the very worst required a subsidy.[18]

What type of ego utility function explains this data? Consider a subject with most of the probability mass of their beliefs in the top 3 ranks. In our data, on average this subject has positive WTP to learn their true rank. This is precisely the opposite behavior generated by Köszegi's functional form. In his model, the person would not want to risk receiving bad news. For this person to risk receiving bad news it must be that the potential ego gains outweigh the losses. This is the case if the ego utility function has increasing marginal utility as we move towards rank 1. There also must be increasing marginal disutility from median beliefs towards rank 10 so that subjects with the lowest round 3 beliefs require the largest subsidy to learn their true rank. Since agents in the center of the distribution have a flat WTP function, we infer ego utility is flat in this "average region." Essentially, subjects placed a high premium of knowing for certain they were the best in the room and a had a strong fear of learning they were the worst in the room.[19]

The type of search found has the property that beliefs will always be more precise as

[18]There is some evidence that male subjects have a higher WTP when they have beliefs near 1 as compared to females and females require a larger subsidy when beliefs are near 10. However, the differences are not significant. A larger sample size would be needed to examine gender differences at the tails of the distribution.

[19]An alternative plausible explanation, albeit one we view unlikely, is based in the psychological concept of "affirmation." Imagine an agent who is 99% sure she is rank 2. Suppose despite her near certainty, she has a high WTP to learn the truth. Positive (negative) feedback that has little informational value is known as affirmation (disaffirmation). Affirmation (disaffirmation) seeking (avoidance) has some explanatory power in our setting, however the significant effect of of round 3 belief variance shows that it is not the entire story.

we move towards the top of distribution. A consequence is that high types are likely to know with precision their percentile rank while low types will have more diffuse beliefs – nobody will believe with a high degree of confidence that they are in the bottom percentile ranks. This pattern of beliefs is not classic overconfidence (mean beliefs above average), but it can have economic consequences if people base decisions on the first moment of their distribution.[20] An example germane to our study is mate selection. Suppose nobody's mean belief places them in the bottom quartile of mate desirability. This can potentially disrupt the sorting necessary to reach equilibrium. Our second example involves a classic question in labor economics. Suppose the returns a college degree are positive only for those in the top 70% of the population. If subjects choose the mean of their beliefs to make decisions, then a lack of precision for below average types can lead them to enroll at sub-optimally high rates.

In round 0 subjects provided their beliefs with out an informational feedback. Our findings in this section are consistent with the prediction that below average subjects have a higher standard deviation in their beliefs than above average subjects. Across image tasks, the mean standard deviation of round 0 beliefs for above average types is 1.52, whereas it is 1.94 for below average types. The difference is highly significant. ($p = 0.00005, t = 4.263$). Overall, the precision heterogeneity in subjects' round 0 priors is consistent with the testable implication of our informational search findings.

5 Discussion

In this section we discuss how the good news-bad news effect can provide an explanation for observed confirmation bias and self-serving bias. We also illustrate how optimistic updating and differential search can compliment each other in generating overconfidence on image qualities, discuss the theoretical underpinnings of our findings and suggest real-world applications.

[20] The effect on mean overconfidence depends on the signal informativeness, whether priors exhibit over/underconfidence and whether agents update rationally. For instance, if agents are overconfident and signals are perfectly informative, then this search will tend to reduce net overconfidence. If agents are rational and signals are imperfectly informative then this ego utility function will actually induce underconfidence. This results follows from proposition 1 of Köszegi (2006). We discuss how this type of search combined with the inference pattern we observed can lead to overconfidence in Section 5.3.

5.1 The good news-bad news effect and confirmatory bias

In Section 3 we showed that subjects incorporated favorable and unfavorable news into their beliefs in fundamentally different ways – the tighter adherence to Bayesian rationality for good news was striking. Another well documented bias involving differential processing of information is the confirmatory bias (CB). The classic account of CB is that people overreact or react appropriately to new information that agrees with their prior belief and always underreact to information that disagrees with their prior belief.[21] The precise definitions of "agree" and "disagree" can be murky in some settings. We think in our experiment the most natural classification is as follows. For subjects whose mean belief places them in the top 5 ranks (above average) $s = 1$ is confirming and $s = 0$ is disconfirming. For subjects whose mean belief places them in the bottom 5 ranks (below average) $s = 1$ is disconfirming and $s = 0$ is confirming.

Results presented in Section 3 already indicated that our data show no evidence of CB. For instance, the lower responsiveness to bad news for subjects who believed they were below average, as shown in Figure 2 and Table 1, is precisely the opposite of the prediction of CB. In Table 5 we look for evidence of CB directly in comparing the subjects' round-to-round change in beliefs to the Bayesian change for confirming and disconfirming signals ($\Delta\mu - \Delta\mu_{Bayes}$). Negative values indicate optimistic updating towards rank 1.

Table 5: $\Delta\mu - \Delta\mu_{Bayes}$ as a Function of Prior Group and Signal Valence

	Good news		Bad news	
	Above avg	Below avg	Above avg	Below avg
Beauty	-0.14	-0.13	0.00	-0.13
	(-3.74)	(-1.20)	(-0.03)	(-1.32)
IQ	0.07	-0.08	-0.14	0.02
	(1.46)	(-0.70)	(-0.91)	(0.23)
Control	0.13	0.21	-0.36	-0.37
	(1.69)	(2.11)	(-5.16)	(-3.89)
CB prediction	$-$	$+$	$-$	$+$

t statistics ($H_0 = 0$) in parentheses

Examining *Beauty* we see above average subjects did tend to overreact to good news, consistent with CB. However, it also evident that below average subjects receiving $s = 1$ overreacted with commensurate magnitude (0.01 less). Specifically they did not ignore or

[21]Sometimes confirmatory bias is used to explain seeking information that is unlikely to be disconfirming (failing to test alternative hypotheses). This search-based CB finds support in our data.

under-respond to the disconfirming good news. For $s = 0$ is it is a similar story, CB predicts that below average subjects would over-respond to $s = 0$ but the opposite occurs — there is an under-response. For IQ, the patterns are similar but of a lower magnitude. What appears to matter most in the image conditions is the signal valence, not whether a signal confirmed the prior belief. The patterns in the *Control* do not exhibit CB either. The negative values for bad news are also seen in the lower constant in the regressions in Table 3.

The question remains: what is the underlying cause CB? If it were a hardwired feature of human updating machinery, then we would expect to see it in updating on image tasks. Yet we do not. A CB agent has an aversion to switching their beliefs and admitting they were wrong. One reason could be that they get ego utility (self-esteem) from believing they are right.[22] Disconfirming signals are bad news while confirming signals are good news. When the underlying quality has direct self-esteem consequences, we do expect CB because the ego utility of being right is outweighed by the direct image effects. This is precisely what we observe in our experiment. It also explains why we do not observe CB in *Control* – there is very low utility of being right about a random number. It is no surprise then that evidence of CB usually involves beliefs that do not have a direct image consequence but that are tied to identity (e.g. sides of a hot political debate (Plous 1991, Lord, Ross and Lepper 1979) or beliefs which drove costly effort in the past (Mahoney 1977)).

In past experiments showing CB, confirming signals were always good news. The perfect co-linearity of confirmation and valence rendered it impossible to distinguish between them. In our experiment, confirming or disconfirming news could be intrinsically favorable or unfavorable. In the image conditions what mattered was signal valence, not whether the signal confirmed the prior belief. In the context of the "psychological defense mechanism" discussed in the introduction, this result is not surprising. There is extensive evidence that people are more likely to remember good as opposed to bad events (Mischel, 1976), attribute good outcomes resulting from their actions to skill and bad outcomes to luck (attribution theory Heider, 1958) and shield themselves from uncomfortable truths such as the danger of their jobs (cognitive dissonance Feisteiger, 1957 and Akerlof and Dickens, 1982). The good news-bad news effect places CB within the realm of these self-protecting defenses.

5.2 Self-serving bias and signal valence

Self-serving bias is the tendency to conflate self-interest with moral judgment. It can lead to bargaining impasses since both parties strongly believe they are right (Babcock and

[22] It could also be anticipation utility. Suppose an agent made decisions today which pay out in the future based on the accuracy of their beliefs. If they change their mind, it means these actions will have a low expected payout and this lowers "anticipatory utility" today.

Loewenstein 1997). There is strong experimental evidence for self-serving bias (Messick and Sentis 1979, Babcock, Loewenstein, Issacharoff and Camerer 1995). In a typical experiment subjects are randomly assigned a role in bargaining game. The role involves a detailed back story and the subject's aim is to negotiate the best deal possible for themselves. Despite the fact that roles are randomly assigned, both groups tend to report that they deserve a greater share of the pie *based on the facts of the case*. Self-serving bias has also been shown to have economic and health consequences in the field (Larwood 1978, Babcock, Wang and Loewenstein 1996).

In bargaining, signals that one will win the case are have positive valence – they indicate that a financial windfall is forthcoming. They might also mean that one did not act wrongly in the past as in, for instance, a divorce hearing or patent infringement case. The good news-bad news effect predicts that these positive signals will receive disproportionate weight as compared to negative signals, and consequently both sides will exhibit excessive confidence that they are in the right. Notice that CB has a tough time explaining self-serving bias, especially in controlled experimental settings. With random assignment, it is unreasonable to posit that ex-ante both groups of subjects believe they will be placed in the role that "should" win. Signals cannot be confirming or disconfirming in this setting. Conversely, since winning means higher earnings, news has intrinsic valence.

5.3 The joint effect of the search and inference results

In Section 4 we showed that WTP to learn true rank increased as beliefs moved toward rank 1. Our design used revelation of the truth (instead of just another up/down comparison) to amplify the differences in WTP and increase the power of our statistical tests. It useful to consider how the results would apply to situations in which signals are not perfectly informative. Figure 3 shows that people tend to over-respond to relatively uninformative signals and do not respect the strength of informative bad news. When one believes already believes they are high in the distribution good news is uninformative (it should be expected) while bad news is very informative. This means that if high types are seeking signals and generally receiving positive feedback, then they will develop very optimistic beliefs.

Biased search and processing can also lead to path dependence in beliefs. Path dependence in beliefs is stressed in Rabin and Shrag (1999)'s model of CB. Biased search and inference has a similar effect . The bias induced by the good news-bad news effect alone increases with initial overconfidence because bad news becomes increasingly informative. Our search findings suggest these are the signals likely to be actually received. Overconfidence on image tasks will develop when priors are accurate or optimistic, but it might not develop if beliefs evolve to (or begin at) a sufficiently underconfident state. In the latter case, underconfidence can get locked in as people avoid new information.

One application of this result is to the finding that people tend to be overconfident on easy tasks and underconfident on hard tasks (for a thorough review see Moore and Healy, 2008). If all comparisons are relative, then the distribution of good news and bad news will be independent of task difficulty. However, for absolute feedback (e.g. test score as opposed to percentile rank) bad news is more likely for a difficult task and good news for an easy task. If people over-respond to relatively uninformative success on an easy tasks, overconfidence will develop. This observation suggests that it would be nice to look at information search when people can choose between absolute and relative information. Our results suggest that they would choose relative comparisons for difficult tasks and absolute comparisons for easy tasks.

5.4 Explaining the good news-bad news effect

We have argued thus far that a direct concern for self-esteem and self-image endows signals with intrinsic valence, which leads to differential processing. The notion of a direct utility to beliefs is gaining purchase among economic theorists. Akerlof and Dickens (1982) first used belief-based utility in their model of cognitive dissonance.[23] In their model, agents rationally maintain optimistic views on the danger of their jobs in order to avoid "fear utility" in the interim. Benabou and Tirole (2002) and Brunnermeier and Parker (2005) generalize this notion calling it "anticipation utility" and show that a rational agent will always "choose" an updating rule that is biased towards optimism. Compte and Postlewaite (2004) modify preferences so that self-confidence enhances performance and establish a similar result. The intuition behind these models is that using Bayes rule maximizes belief accuracy, so moving away from Bayes rule toward optimistic updating has only a second order cost (through the loss in accuracy) but a first order benefit through either anticipatory utility or enhanced performance. As such a rational agent always prefers to induce some optimistic bias.

When beliefs affect utility in these "non-standard" ways, agents have an incentive to sacrifice belief accuracy in order to increase the belief based utility, blurring the distinction between hopes and beliefs. Normatively speaking, the sacrifice in accuracy should depend on the marginal returns to belief accuracy as compared to the marginal returns to belief utility. Whether the good news-bad news effect is rational or not depends on whether agents can overcome the tendency to discount bad news when the returns to belief accuracy increase.

[23]The notion of cognitive dissonance was introduced formally by Festinger (1957). For a review see Simons et al. (1970).

An experimental design which increases the costs to inaccurate beliefs while holding belief utility constant would be able to answer this interesting question. We leave this to future work.

5.5 Real-world applications

In our experiment signal valence was generated by ranking 10 strangers after a short task measuring their intelligence or beauty. In an experimental setting, ethical considerations rightly limit "how bad" news can be. Rank on the image task, as compared to 9 strangers, endowed the signals with enough psychological valence to significantly change their updating behavior. In the real-world, news can range from the overwhelmingly positive to the jarringly negative. In these situations we would expect differential processing by signal valence to play an even larger role than it did in the experimental setting. We have already discussed real-world applications in which the effect of valence is mediated through CB, overconfidence and self-serving bias. However, there are applications that fall outside the scope of these previously identified inference mistakes. We discuss a few such examples below.

In finance, the good news-bad news effect predicts that an agent will respond differently to news on a security she owns as compared to news on a security that she does not own. If people have a difficult time responding objectively to news on securities they own, and on some level recognize this, then it make sense to hire financial professionals to do the job for them. They are willing to pay a financial manager for doing a job that an educated person could quite easily accomplish, at a far lower cost, by purchasing mutual funds through a discount online trading account.

Further evidence from finance come from the literature on post-earnings announcement drift (PEAD). Predictable drift in stock prices should not occur in an efficient market, but many authors have found that stocks experiencing earnings surprises do predictably drift over a 120 day period following earnings surprises (Foster, Olsen and Shevlin 1984). Bernard and Thomas (1989) found that the drift is significantly more pronounced for bad news stocks as compared to good news stocks. That is, the initial price movement accounts for less of the available information for bad news as compared to good news, which is precisely our primary finding. The authors show that these patterns are driven by small-cap stocks (those traded by non-institutional investors). For large-cap stocks good news and bad news are treated the same. The authors also rule out risk-premium-based and other explanations for the observed drift.

An application in medicine is a hospital patient's prediction of their prognosis – here news is literally life-or-death. Studies have found that cancer and HIV+ patients optimistically assess their chances of survival (Weeks, Cook, O'Day, Peterson, Wenger, Reding, Harrell, Kussin, Dawson and Connors 1998, Eidinger and Schapira 1984, Bhattacharya, Goldman and Sood 2009). It is not, however, that their physicians are misleading or giving them intentionally optimistic forecasts; the physicians' forecasts are shown to be accurate and unbiased. In fact, the phenomenon has attracted attention in the medical community because it complicates physicians' ability apply treatment they find most appropriate given the circumstances. An optimistic patient is liable to make decisions that are not in her best interest and physicians have adopted strategies to try to alleviate this problem (Paling 2003). A related application is to preventative care. Studies show that Americans tend not to heed early warnings signs of disease, which drives up health care costs in the long-run (Schuster, McGlynn and Brook 2005). Early signs of illness are unwelcome bad news and our prediction is that this news is discounted.

6 Conclusion

We have shown that image effects have an important implication for Bayesian inference. When updating on a attributes closely tied to self-esteem, new information is not only confirming or disconfirming but also intrinsically favorable or unfavorable. In our image treatments, good news and bad news were incorporated into prior beliefs in entirely different ways. For good news, the inference is less noisy, respects signal strength and conforms much more closely to Bayesian rationality. Conversely for bad news, beliefs are more noisy and scarcely resemble the normative Bayesian standard.

The image treatments used attributes designed to have a strong self-esteem association: rank out of 10 subjects in the session on intelligence as given by an IQ test and physical attractiveness (rankings within gender group) as given by peer ratings in a speed dating exercise. In our control condition we used randomly generated subject ranks. Here mistakes were made, but they were largely symmetric with respect to signal direction. In further support of our main finding, gender differences in image effects were consistent with studies in the mate selection literature that show physical attractiveness has a higher importance for women and intelligence a higher importance for men.

In Section 4 we further examined the "sting of bad news" by providing subjects with the opportunity to learn their true rank on both the image task and control. Subjects who believed they were near the top of the distribution had a high WTP to learn their true rank. Conversely, subjects who believed they were near the bottom of the distribution required a subsidy to learn the truth.

Past work has shown that humans have a robust "psychological defense mechanism" to maintain positive self-esteem. A separate literature has examined systematic departures from Bayesian rationality in signal inference. This paper draws from both literatures to study how people incorporate objective signals into their existing beliefs as the degree of self-esteem impact is varied. The results point to differential processing by signal valence as the underlying cause of confirmatory bias, self-serving bias and overconfidence. Furthermore they place these biases within the realm of the psychological defense mechanism and confirm the predictions of theoretical models that predict valence-dependent updating because agents have a direct utility to beliefs. The paper suggests that future research could benefit from further synthesis of these two fields.

References

Alicke, M.D. and O. Govorun, *The Better-than-Average Effect*, Psychology Pr, 2005.

Andreoni, J. and B.D. Bernheim, "Social Image and the 50-50 norm: A Theoretical and Experimental Analysis of Audience Effects," *Econometrica*, 2009, *77*, 1607–1636.

Babcock, L. and G. Loewenstein, "Explaining Bargaining Impasse: The Role of Self-serving Biases," *The Journal of Economic Perspectives*, 1997, pp. 109–126.

____ , ____ , S. Issacharoff, and C. Camerer, "Biased Judgments of Fairness in Bargaining," *The American Economic Review*, 1995, pp. 1337–1343.

____ , X. Wang, and G. Loewenstein, "Choosing the Wrong Pond: Social Comparisons in Negotiations that Reflect a Self-serving Bias," *The Quarterly Journal of Economics*, 1996, *111* (1), 1–19.

Barber, B.M. and T. Odean, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 2001, *116* (1), 261–292.

Becker, G.M., M.H. DeGroot, and J. Marschak, "Measuring Utility by a Single-response Sequential Method," *Behavioral Science*, 1964, *9* (3), 226–232.

Bernard, V.L. and J.K. Thomas, "Post-earnings-announcement Drift: Delayed Price Response or Risk Premium?," *Journal of Accounting research*, 1989, *27*, 1–36.

Bhattacharya, J., D. Goldman, and N. Sood, "Market Evidence of Misperceived Mortality Risk," *Journal of Economic Behavior and Organization*, 2009.

Brunnermeier, M.K. and J.A. Parker, "Optimal Expectations," *The American Economic Review*, 2005, *95* (4), 1092–1118.

Carrillo, J.D. and T. Mariotti, "Strategic Ignorance as a Self-disciplining Device," *Review of economic studies*, 2000, pp. 529–544.

Christensen-Szalanski, JJ and J.B. Bushyhead, "Physicians' Use of Probabilistic Information in a Real Clinical Setting," *Professional Judgment: A Reader in Clinical Decision Making*, 1988.

Compte, O. and A. Postlewaite, "Confidence-enhanced Performance," *American Economic Review*, 2004, pp. 1536–1557.

Darley, J. and P. Gross, "A Hypothesis Confirming Bias in Labeling Effects," *Journal of Personality and Social Psychology*, 1983, *XLIV.*

Eidinger, R.N. and D.V. Schapira, "Cancer Patients' Insight Into Their Treatment, Prognosis, and Unconventional Therapies," *Cancer*, 1984, *53* (12).

Festinger, L., *A Theory of Cognitive Dissonance*, Stanford University Press: Stanford, CA., 1957.

Fisman, R., S.S. Iyengar, E. Kamenica, and I. Simonson, "Gender Differences in Mate Selection: Evidence from a Speed Dating Experiment," *Quarterly Journal of Economics*, 2006, *121* (2), 673–697.

Foster, G., C. Olsen, and T. Shevlin, "Earnings Releases, Anomalies, and the Behavior of Security Returns," *The Accounting Review*, 1984, *59* (4), 574–603.

Frey, D., "The Effect of Negative Feedback About Oneself and Cost of Information on Preferences for Information About the Source of this Feedback," *Journal of Experimental Social Psychology*, 1981, *17* (1), 42–50.

Gelman, A., "Letter to the editors regarding some papers of Dr. Satoshi Kanazawa," *Journal of Theoretical Biology*, 2007, *245* (3), 597–599.

Grether, D.M., "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 1980, *95*, 537–557.

Heider, F., *The Psychology of Interpersonal Relations*, Wiley: New York., 1958.

Holt, C.A., "Scoring-rule Procedures for Eliciting Subjective Probability and Utility Functions," *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti. North Holland Press, Amsterdam*, 1986.

Kanazawa, S., "Beautiful Parents Have More Daughters: A Further Implication of the Generalized Trivers–Willard Hypothesis," *Journal of theoretical biology*, 2007, *244* (1), 133–140.

Kelley, HH and JL Michela, "Attribution Theory and Research," *Annual review of psychology*, 1980, *31* (1), 457–501.

Kőszegi, B., "Ego utility, Overconfidence, and Task Choice," *Journal of the European Economic Association*, 2006, *4* (4), 673–707.

Larwood, L., "Swine flu: A Field Study of Self-serving Biases," *Journal of Applied Social Psychology*, 1978, *8* (3), 283–289.

Lord, C., G. Ross, and L. Lepper, "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered evidence," *Personality and Social Psychology*, 1979, *37* (11).

Mahoney, Michael J., "Publication Prejudices: An Experimental Study of Confirmatory Bias in the Peer Review System," *Cognitive Therapy and Research*, 1977, *1* (2).

Messick, D.M. and K.P. Sentis, "Fairness and Preference," *Journal of Experimental Social Psychology*, 1979, *15* (4), 418–434.

Mischel, W., "Determinants of Selective Memory About the Self," *Journal of Consulting and Clinical Psychology*, 1976, *44* (1), 92–102.

Moore, D.A. and P.J. Healy, "The Trouble with Overconfidence," *Psychological Review*, 2008, *115* (2), 502.

Nickerson, Raymond, "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises," *Review of General Psychology*, 1998, *2* (2), 175–220.

Paling, J., "Strategies to Help Patients Understand Risks," *British Medical Journal*, 2003, *327* (7417), 745–748.

Plous, S., "Biases in the Assimilation of Technological Breakdowns: Do Accidents Make Us Safer?," *Journal of Applied Social Psychology*, 1991, *21.*

Rabin, M. and J.L. Schrag, "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, 1999, *114* (1), 37–82.

Santos-Pinto, L. and J. Sobel, "A model of Positive Self-image in Subjective Assessments," *American Economic Review*, 2005, *95* (5), 1386–1402.

Schuster, M.A., E.A. McGlynn, and R.H. Brook, "How Good is the Quality of Health Care in the United States?," *The Milbank Quarterly*, 2005, *83* (4), 843–895.

Selten, R., "Axiomatic Characterization of the Quadratic Scoring Rule," *Experimental Economics*, 1998, *1* (1), 43–62.

Simons, H.W., N.N. Berkowitz, and R.J. Moyer, "Similarity, Credibility, and Attitude Change: A Review and a Theory," *Psychological Bulletin*, 1970, *73* (1), 1–16.

Taylor, S.E. and J.D. Brown, "Positive Illusions and Well-being Revisited: Separating Fact from Fiction," *Psychological Bulletin*, 1994, *116*, 21–21.

Trivers, R.L. and D.E. Willard, "Natural Selection of Parental Ability to Vary the Sex Ratio of Offspring," *Science*, 1973, *179* (4068), 90.

Tversky, A. and D. Kahneman, "Judgment Under Uncertainty: Heuristics and Biases," *Science*, 1974, *185* (4157), 1124–1131.

Weeks, J.C., E.F. Cook, S.J. O'Day, L.M. Peterson, N. Wenger, D. Reding, F.E. Harrell, P. Kussin, N.V. Dawson, and A.F. Connors Jr, "Relationship Between Cancer Patients' Predictions of Prognosis and Their Treatment Preferences," *Journal of the American Medical Association*, 1998, *279* (21), 1709–1714.

7 Appendix

7.1 Experimental Instructions

Instructions taken from a *Beauty* session.

Welcome

Thank your for participating in our experiment. We will begin shortly. Today's experiment will last under an hour and half.

Informed Consent

Placed in front of you is an informed consent form to protect your rights as a subject. Please read and sign it. If you would like to choose not to participate in the study you are free to leave at this point. If you have any questions, we can address those now. We will now pass through the aisle to pick up the forms.

Anonymity

Your anonymity in this study is assured. Your name will never be collected or connected to any decision you make here today. Your email was collected for invitation purposes only, it will never be connected to your performance in the study. Furthermore, your earnings will be paid in a sealed envelope with your subject letter ID so that even those running the study will not know your earnings. No other subject in the study will know any of your choices or performance in the study.

Rules

- Please turn your cell phones off.
- If you have a question at any point, just raise your hand.
- Please put away any books, papers, computers, etc that you have brought with you.

Stages of Today's Experiment

Today's experiment will have two stages. The instructions for each stage will be read at the beginning of that stage. At the conclusion of stage 2, your payments will be prepared and the experiment will end.

Your Earnings

Some of the decisions you make today will impact your earnings. We will explain exactly which decisions and how earnings will be calculated at the appropriate time. Your earnings will be paid in cash, placed in a sealed envelope with your subject number. Although earnings will vary by subject, most subjects will earn between $20 and $30.

Stage 1 Instructions, Part 1

You will notice that this experiment is exactly half males (Subject ID M-A through M-J) and half females (W-A through W-J). Shortly we will begin a "speed dating" exercise in which you will meet 5 members of the opposite sex for 4 minutes a piece. Suggested conversation topics have been provided for you. Try to stay on these topics. These are:

- What are your plans following college?
- What is your favorite book/movie/album?

Following each 4 minute "meeting" you will fill out a brief meeting questionnaire. On this you will rate the person you just spoke with on three categories: ambition, attractiveness, friendliness. This will be done at a separate work station so that no other subject will ever see your ratings.

In your subject packet are is meeting questionnaires with your subject ID and the subject ID letters of the 5 participants you will meet. Males will move clockwise around the room on the switch. You will fill out the questionnaires privately at a different station. Under absolutely no circumstances will any other subject see the ratings you record. Your honesty is paramount to the aims of this study. After you fill out the form, place it in the manila envelope with your subject letter and proceed to the next station. We will now ask all the males on the inner column to switch with the females across the aisle and show you exactly how this stage will proceed.

Stage 1 Instructions, cont.

Please return to your original seat and we will proceed with the study.

We are now passing you a sealed envelope with your subject letter ID written on it. Within each is a note card with a number on it. For the males, these numbers range from 1-10, with each male having 1 number. The same applies to the females. For the remainder of the study, we will refer to the left side of the room as the "male group" and right side as the "female group." Within each group, each one of you has different number. The numbers have been assigned randomly by rolling a ten sided die prior to the study [subjects are shown the die]. The role of this envelope in the study will be explained at the appropriate time. You will not be allowed to open the envelope until the experiment has concluded. [stop] Are there any questions?

Things To Remember

The numbers in the envelope labeled "Card Number" range from 1 to 10, each subject in a given group has a different number. These numbers were assigned randomly. *Do not open the envelope.*

This Concludes Stage 1.

Stage 2 Instructions , Set 1 (Appearance Task)

Stage 2 consists of 2 "sets" each consisting of 4 "rounds."

In stage 2 we will ask you tell us the probability you think you occupy each of 10 ranks for the "task." The reason that there are only 10 ranks is that you will be making comparisons only to members of your gender group. So since there are 10 subjects in each group, this means that for each task there exists a ranking 1-10.

For this task, the appearance task, your rank is given by your "appearance score" which is the average of the 5 ratings on "appearance" in stage 1 on the experiment. The subject in your gender group with the highest average score is ranked 1, the one with the lowest, with your group, is ranked 10. A sample decision screen is now being displayed. You can control each "rank box" by clicking the slider and/or the arrows adjacent to the box. (Short description and demonstration).

In each round, you will be asked to fill in these boxes with what you think are the probabilities that your appearance rank is the number beside the box. You can do so by moving the slider or clicking the arrows. Notice that the card numbers give ranks 1-10, just as the boxes. These probabilities must add up to 100%. If they do not add up to 100%, you will be paid nothing for the round. The total is displayed to you on the screen. If they do add up to 100%, you will be paid according to the following formula:

$$2 - \sum_{i=1}^{1} 0(1\{rank = i\} - p_i)^2 \tag{1}$$

While this payoff formula may look complicated, what it means for you is simple: **you get paid the most on average when you honestly report your best guesses of the probability for each rank.** The range of payoff is 0-2 dollars for each round of guesses.

Each "decision form" is labeled by round and task number. These labels are on the tabs of the spreadsheet. (short demonstration moving between tabs). Once you complete a decision form, please click forward to the next tab. You may not go back and change your previous round guesses. However the guesses from your all previous rounds of that set will be displayed at the top of the screen. The spreadsheet records the time of entry and will check to make sure that all entries were down at the right time. You will not be eligible for payments if you have an entry that is out of order (i.e. if you went back and changed a previous entry).

Before rounds 1, 2, and 3, your rank will be compared to a randomly chosen "comparison participant" FROM YOUR GENDER GROUP. You will never know the identity of these participants. In each round, it will be a different participant. To convey the result of the comparison of ranks, we will pass you a "signal card" with the words "You are ranked higher" or "You are ranked lower" written on it. If you get the message "You are ranked higher" in a given round, it means that your appearance score was higher (equivalently your rank was closer to 1) than your comparison participant for that round. If you get the message "You are ranked lower" in a given round, it means that your appearance score was lower than your comparison participant for that round. For example, if your appearance rank was 5 and we compared you to a subject whose score ranked them 7, then you would receive the message "You are ranked higher." At the beginning of the next round, you will make a new set of guesses.

After receiving your message in rounds 1-3, click the box "Enter message here" and change it to reflect the message you actually received. This way we know you got the message.

In rounds 0-2, we will ask you to the probability you expect to get a "higher" message in the next round. Your payment for that guess will be made in the same fashion as the rest, so once again you have the incentive to truthfully tell us the probability you expect to to receive a "higher" signal in the subsequent round.

Things to Remember

- In round 0 you will make your guesses before receiving any messages.
- All rank comparisons and guesses are WITHIN GENDER GROUP.
- Your appearance score gives your rank for this task, highest score is rank 1 and so forth. The "signal cards" compare your rank to a randomly chosen "comparison participant" from your gender group.
- Before rounds 1, 2 and 3 we will compare your rank to a randomly chosen "comparison participant," which will differ by round.
- You will receive a message that says either "You are ranked higher" or "You are ranked lower"
- After receiving the message, you will make new guesses for the next round.
- In rounds 0-2, you will also make a guess at the probability of receiving a "higher" message in the subsequent round.
- You will make the most money on average by making your probabilistic guesses honestly.
- You will make your guesses by clicking the arrows or moving the slider adjacent to the "rank boxes" on the spreadsheet in front of you. The graph will change as you click the arrows or move the slider and the chosen percent for that rank will be displayed.
- Your guesses must total 100% in each round to be eligible for payment.
- Once you complete the decision form for a round, please click forward to the next round using the tabs at the bottom of the screen. You may not click back to change your answers as this will violate the aims of the experiment and make you ineligible for payments.
- Once everyone is done with the round, we will proceed to the next, so take your time when making decisions.

Stage 2, Set 2 (Card Number Task)
We will now begin set 2.
In the enveloped with your subject ID letter is a note card with a number from 1-10 on it. Within each gender group, one person has each one of these numbers.

This randomly assigned number gives your rank for this task, 1-10. Each of the 10 ranks is filled by exactly 1 subject within your group. Set 2 will proceed exactly like set 1, except that your rank is determined by your card number.

Things to Remember

- In round 0 you will make your guesses before receiving any messages.
- All rank comparisons and guesses are WITHIN YOUR GENDER GROUP.
- The subject whose card number is "1" is ranked 1. The "signal cards" compare your rank to a randomly chosen "comparison participant" from your gender group.
- Before rounds 1, 2 and 3 we will compare your rank to a randomly chosen "comparison participant," which will differ by round.
- You will receive a message that says either "You are ranked higher" or "You are ranked lower."
- After receiving the message, you will make new guesses for the next round.
- In rounds 0-2, you will also make a guess at the probability of receiving a "higher" message in the subsequent round.
- You will make the most money on average by making your probabilistic guesses honestly.

- You will make your guesses by clicking the arrows or moving the slider adjacent to the "rank boxes" on the spreadsheet in front of you. The graph will change as

- you click the arrows or move the slider and the chosen percent for that rank will be displayed.

- Your guesses must total 100% in each round to be eligible for payment.

- Once you complete the decision form for a round, please click forward to the next round using the tabs at the bottom of the screen. You may not click back to change your answers as this will violate the aims of the experiment and make you ineligible for payments.

- Once everyone is done with the round, we will proceed to the next, so take your time when making decisions.

Stage 3, Final Decisions
You will now have the opportunity to learn your true rank on the Appearance and Card Number task. Whether we reveal your true rank on the Appearance and/or Card Number task will be determined by your choices in the following decisions. The decision will proceed as follows. We are now passing out two "final decision forms." Each form is labeled either "Appearance task" or "Card Number task." Each form asks you to make 15 decisions, indicating whether your prefer option A or option B. Here is a sample of the decision form for the Appearance task: [sample shown]

You may only switch once from column A to B, if you switch at all. Why should you tell the truth? Once we collect the forms we will randomly choose a number (one for each form) between 1 and 15 inclusive and execute your decision for that number. Your final earnings will reflect your choices.

Example 1: Suppose you switch from A to B at decision (3). You are indicating that for a payment of $5.00 you'd be willing to learn the true rank on the task, but for $4.00 you would rather not learn the rank and receive $0. Suppose in this case we randomly draw 15. In this case you would not learn your true rank and receive $0. Suppose instead we draw 2. In this case you would receive $6.00 and learn your rank on the task.

Example 2: Suppose you do not care if you learn your true rank or not. In this case you are best served switching at 8. The reason is this maximizes your chance of receiving money (if 1-7 are drawn) and ensures you will never pay money.

What happens if it is determined you will learn your true rank?
After we collect your decision forms and determine the decisions that will be executed, we will know which subjects have chosen to learn their true rank. If it is determined that you will learn your true rank on either (or both) task(s), we will pass you a piece of paper with the rank(s) written on them. If it was determined that you will not learn your true rank, then this space will be left blank for you. For each rank that was revealed to you (if applicable) you must write the phrase "acknowledged, rank #" right next to the rank, where # is your rank on that task. We will then pick up the forms, making sure that everyone whose rank was revealed to them has acknowledged. We will pass out these forms to every subject, regardless of whether you chose to learn your true rank. As such, no other subject will know whether or not you learned your true rank. Furthermore, no other subject will ever learn your true rank. This process ensures your anonymity. We will pick up your envelope at this time as well.

How are any payments handled?
Your final earnings will reflect the monetary component of the executed decisions.

Things to Remember

- Your choices in this decision (and chance) will determine whether or not you learn you true rank on the Card Number and/or Appearance task.

- You can ensure that you always learn your true rank on a given task by always checking column A. You can ensure you never learn your true rank by always checking column B.

- If you do switch between column A and B, you can only do so once.

- Following these choices, the experiment will conclude.
- Two random numbers (1-15) will be selected to determine which decision will be executed for each form.
- If your choice indicates you learn your true rank you will be notified by an acknowledgment form, which you will initial to indicate you saw the rank.
- You may fill out the two forms in order you wish.
- You have a financial incentive to answer truthfully.