

Wiswall, Matthew; Zafar, Basit

Working Paper

Belief updating among college students: Evidence from experimental variation in information

Staff Report, No. 516

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Wiswall, Matthew; Zafar, Basit (2011) : Belief updating among college students: Evidence from experimental variation in information, Staff Report, No. 516, Federal Reserve Bank of New York, New York, NY

This Version is available at:
<https://hdl.handle.net/10419/60900>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Belief Updating among College Students:
Evidence from Experimental Variation in Information

Matthew Wiswall
Basit Zafar

Staff Report no. 516
September 2011

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Belief Updating among College Students: Evidence from Experimental Variation in Information

Matthew Wiswall and Basit Zafar

Federal Reserve Bank of New York Staff Reports, no. 516

September 2011

JEL classification: D81, D83, D84, I21, I23, J10

Abstract

We investigate how college students form and update their beliefs about future earnings using a unique “information” experiment. We provide college students true information about the population distribution of earnings and observe how this information causes respondents to update their beliefs about their own future earnings. We show that college students are substantially misinformed about population earnings and logically revise their self-beliefs in response to the information we provide, with larger revisions when the information is more specific and is good news. We classify the updating behaviors observed and find that the majority of students are non-Bayesian updaters.

Key words: belief updating, college majors, information, uncertainty, subjective expectations, Bayesian updating

Wiswall: New York University, Department of Economics (e-mail: mwiswall@nyu.edu).

Zafar: Federal Reserve Bank of New York (e-mail: basit.zafar@ny.frb.org). The authors thank the NYU Center for Experimental Social Sciences for assistance conducting the information survey and experiment, and participants at a Federal Reserve Bank of New York seminar and the NYU Experimental Economics Working Group. Da Lin and Scott Nelson provided outstanding research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Schooling decisions, like most decisions, are made under uncertainty, in particular uncertainty about future realizations of schooling-related outcomes such as earnings (Manski, 1989; Altonji, 1993). For schooling decisions, such as choice of college major, one of the crucial elements of the decision making process is the student’s forecast of future earnings in each potential field. Standard economic theory assumes that individuals: (1) have perfect information and are rational forecasters, and (2) process new information about the various choice-specific outcomes as dispassionate Bayesians do. A recent and expanding literature has relaxed the first assumption and collected subjective expectations data.¹

This paper focuses on the second key assumption and studies the process by which college students update their beliefs regarding their future earnings. We conduct an experiment on undergraduate college students of New York University (NYU), where in successive rounds we ask respondents (1) their *self* beliefs about their own expected earnings if they were to major in different fields and (2) their beliefs about the population distribution of earnings. After the initial round in which the baseline beliefs are elicited, we provide students with accurate information on the population characteristics and then re-elicite their self beliefs. Hence, we observe how this new information causes respondents to update their self beliefs. We attempted to make our experimental design as realistic as possible and provided students with various kinds of public information, such as average earnings for US economics or business majors, which these students could encounter in mainstream media sources. Our experimental design creates a unique panel of subjective expectations data allowing us to study the process by which students update their own subjective beliefs in response to a series of known shocks to each student’s information set.

The experimental design we develop is motivated by studies that have found that individuals are not fully informed when making human capital decisions. Most relevant to our study, Betts (1996) finds that college students are misinformed about the population distribution of earnings of current graduates.² When provided with accurate information about the population distribution of earnings of current workers, this paper asks: (1) would students revise their self earnings beliefs in response to this information, and (2) how do they process such information?

In general we expect students to revise their self beliefs if they are misinformed about population earnings, and their self earnings beliefs are linked to their beliefs about population

¹See Manski (2004) for a review of the literature. In the context of schooling choices, studies that use subjective data on returns to schooling and other schooling-related outcomes include Smith and Powell (1990), Blau and Ferber (1991), Betts (1996), Dominitz and Manski (1996), Jacob and Wilder (2010), Kaufmann (2010), Stinebrickner and Stinebrickner (2010; 2011), Zafar (2010; 2011), Giustinelli (2011), Arcidiacono, Hotz, and Kang (2011), Attanasio and Kaufmann (2011), and Wiswall and Zafar (2011).

²Other studies in developing country contexts, such as Jensen (2010) and Nguyen (2010), also find that students have little idea about actual returns to schooling.

earnings. We find that students in our sample, despite belonging to a very high ability group, have biased beliefs about the population distribution of earnings. For example, they under-predict annual earnings of female workers with no college degree by \$15,000 and over-predict earnings of female graduates in Economics/Business by \$23,000. There is also considerable heterogeneity in errors in population earnings by individual characteristics, with more experienced students – those in their second or third year – having relatively more accurate beliefs about population earnings in some instances. We, however, do not find gender differences in accuracy of population earnings.

After providing students public information on population earnings, we find that the majority of respondents revise their self beliefs about their own future earnings at age 30. There is substantial variation in revisions across majors, from an average downward revision of 5% in self earnings in Economics/Business to an average upward revision of 54.5% in the no degree category. Thus, as in other studies that collect data on students' schooling choices and provide information about certain aspects of the choice, we find that students are not fully informed and that providing such information has an effect on their expectations.³

Our survey design with an embedded information experiment also allows us to address the second question and assess *how* students process such information and form expectations. The few studies that have analyzed how students from expectations use panel data on beliefs (Jacob and Wilder, 2011; Stinebrickner and Stinebrickner, 2010, 2011; Zafar, 2011). While these studies are able to study the evolution of expectations and changes in choices, they are limited in their ability to estimate the causal effect of information shocks on expectations. This is because in these previous panel datasets, where each wave is typically separated by several weeks, months, or years, it is extremely challenging to identify innovations in the agent's information set (Dominitz, 1998; Zafar, 2011). Other field experiments that disseminate information about different aspects of schooling choices get around this challenge since the researchers have control over what information is being provided to the respondents (e.g., Jensen, 2010; and Nguyen, 2010). While these studies analyze whether information affects choices, they are unable to shed light on the expectations formation process, largely because detailed data are needed to do so. Since we collect data not only on expected self earnings but also on the distribution of earnings, and on the respondents' priors about the information that we provide to them, we are able to examine directly the heterogeneity in belief updating.

We begin our analysis of the updating process by first using a series of regressions to show that respondents update their beliefs in response to the information treatments, and do so in a logical way: Revisions in self beliefs are related to respondents' population errors (i.e., the

³For example, Hastings and Weinstein (2008) find that providing information to parents about school quality makes them more likely to choose high quality schools. Bettinger et al. (2011), and Dinkelman and Martinez (2011) find that providing information on financial aid improves certain educational outcomes.

gap between true population earnings and perceived population earnings – a measure of the informativeness of the revealed information for the respondents). However, the mean response of revisions in self beliefs to population errors is fairly inelastic: An error of a thousand dollars in population earnings results in a revision of \$34 in self earnings beliefs. This suggests that self beliefs about earnings are not entirely linked to the type of public population information we provide. There is, however, substantial heterogeneity in self earnings revisions in response to information. First, the response to population earnings is more pronounced the more relevant the information is– we find much stronger effects in treatments where respondents are provided with information on population earnings of graduates in specific majors than when they are provided with information about earnings of all workers. More importantly, as in Eil and Rao (2011) and Mobius et al. (2011), we find that the effect of information is asymmetric: There is significant updating when the information is good news for the respondent, i.e., when the respondent is informed that population earnings are higher than her prior beliefs, and no significant updating in instances where the respondent is informed that the population earnings are lower than her prior beliefs.

In the second part of the paper, we estimate a simple model of Bayesian belief-updating and ask how respondents’ observed revisions compare to the case if they were Bayesian. Our analysis shows there is substantial heterogeneity in the information-processing heuristics used by students, with the majority of respondents having non-Bayesian updating, either responding more (“Alarmist”) or less (“Conservative”) than the individual-specific Bayesian benchmark, and a sizable proportion of respondents (15-20%) updating illogically, i.e., revising their beliefs in a way that cannot be rationalized by our updating model. In analyzing the patterns of updating relative to the Bayesian benchmark, we document some important heterogeneity in belief-updating. First, we do not find gender differences in information processing heuristics. Second, relative to freshmen, experienced students are more likely to be non-updaters and less likely to react excessively to information (Alarmist updating). Third, we find evidence of valence-based updating: Respondents are twice as likely to be Conservative in their updating when the news is negative, i.e., when they are informed that population earnings are lower than their prior beliefs, than when the news is positive.

Finally, in the last section, we investigate the effect of our information treatments on future choices, and assess whether our intervention leads to welfare gains. We find that the information on earnings we provide causes nearly half of the students to revise their beliefs about graduating with the different majors. To get a sense of the impact of our information treatments on students’ choices, we compute the welfare change – defined as change in future expected earnings – for our sample. The mean welfare change in our sample is an increase of \$432 in age 30 earnings, and the welfare change is non-negative for three-quarters of our sample. We also

show that imposing Bayesian updating would severely underestimate the welfare gains from our experiment, implying that Bayesian updating is the sub-optimal heuristic for most students. This highlights the importance of using actual data on belief-updating rather than relying on a homogenous information-processing rule.

This paper adds to the large experimental literature on information processing. One strand of this literature explores the updating of ego-independent quantities such as which urn a ball is drawn from (Grether, 1980; El-Gamal and Grether, 1995). The second category studies information processing rules in settings that are more realistic and where beliefs have direct importance such as ability, performance, climate change, risk assessment, and effectiveness of contraceptives (see, for example, Viscusi and O'Connor, 1984; Cameron, 2005; Delavande, 2008; Eil and Rao, 2010; Grossman and Owens, 2011; Mobius et al., 2011). Our paper belongs to the second category: We consider the updating of earnings expectations in the context of college major choice—an important decision with significant economic consequences. In addition, most of the existing studies consider updating of binary outcomes, or have an information structure where the signal is binary. Our setting is a hybrid design that combines experimentally manipulated information as in laboratory experiments with a situation that is closer to real-world field experiments. As a result, our setup differs from the textbook case of Bayesian updating in two ways. First, information revealed to students may already be known to them. Second, while students are revising private beliefs about themselves, they receive public information. Both these differences have implications for the interpretation of our results. For example, our setup should be biased against the finding that respondents respond excessively to information. Yet, we find that nearly a quarter of our respondents fall in this category. We show that our classification of updating heuristics is robust to these features of the study design.

The next section describes the data and experimental setup. The following two sections explore the heterogeneity in population errors and analyzes the patters of revisions of self-earnings. Section 5 discusses the significance of the information experiment on choices and measures of student welfare. The last section concludes.

2 Data

2.1 Administration

Our data is from an original survey instrument administered to New York University (NYU) undergraduate students over a 3-week period, during May-June 2010. NYU is a large, selective, private university located in New York City. The students were recruited from the email list used by the Center for Experimental Social Sciences (CESS) at NYU. The study was limited to full time NYU students who were in their freshman, sophomore, or junior years, were at

least 18 years of age, and US citizens. Upon agreeing to participate in the survey, students were sent an online link to the survey (constructed using the SurveyMonkey software). The students could use any internet-connected computer to complete the survey, and were given 2-3 days to start the survey before the link became inactive. They were told to complete the survey in one sitting. The survey took approximately 90 minutes to complete, and consisted of several parts. Students were not allowed to revise answers to any prior questions after new information treatments was provided. Many of the questions had built-in logical checks (e.g., percent chances had to be between 0 and 100). Students were compensated \$30 for successfully completing the survey.

2.2 Survey Instrument

The survey instrument consisted of three stages (see Figure 1):

1. Initial Stage: Respondents were asked their *population* beliefs—beliefs about the earnings of current workers in the labor force, and *self* beliefs—beliefs about own earnings and other outcomes, conditional on completing various majors.
2. Intermediate Stage: Respondents were randomly selected to receive 1 of 4 possible information treatments. Each information treatment revealed statistics about the earnings and labor supply of a certain group of the US population. The information was reported on the screen and the respondents were asked to read this information before they continued. Respondents were then re-asked their population beliefs (on areas they were not provided information about) and self beliefs.
3. Final Stage: Respondents were given all of the information contained in each of the 4 possible information treatments. Respondents were then re-asked about their self beliefs.

The 4 information treatments consisted of statistics about the earnings and labor supply of the US population. Table 1 lists the 4 information treatments:

1. All Individuals Treatment: revealed earnings for the population of all US workers currently aged 30.
2. College Treatment: revealed earnings for the population of college graduates currently aged 30.
3. Female Major Specific Treatment: revealed earnings for female bachelor degree holders currently aged 30 by specific college major.

4. Male Major Specific Treatment: revealed earnings for male bachelor degree holders currently aged 30 by specific college major.

We often combine results from the treatments where we classify the All Individuals and College Treatments as *General* treatments, and the Female and Male Major Specific Treatments as *Major Specific* treatments.

The information treatments were calculated by the authors using the Current Population Survey (for earnings and employment for the general and college educated population) and the National Survey of College Graduates (for earnings and employment by college major). Details on the calculation of the statistics used in the information treatment are in the Appendix; this information was also provided to the survey respondents at the conclusion of the survey. Survey respondents were randomly provided with one of these information treatments in the intermediate stage. Before the population information was revealed, respondents were asked about their prior beliefs about these population statistics. After revelation of information, respondents were re-asked some of their self beliefs, including the major-specific earnings distribution at age 30.

The goal of this paper is to shed light on how students form earnings expectations. For that purpose, we focus on updating of self beliefs for earnings. Respondents were asked about earnings in their first job after college and for later periods at ages 30 and 45. Since the information about population earnings pertained to current 30 year olds, we focus on updating of earnings reported for age 30. In this paper, we use Initial Stage and Intermediate Stage beliefs in the analysis only.

We asked about earnings conditional on completing different college majors. Because of time constraints, we were forced to make difficult choices in the aggregation of college majors. We aggregate college majors to 5 groups: 1) Business and Economics, 2) Engineering and Computer Science, 3) Humanities, Arts, and Other Social Sciences (e.g. Sociology), 4) Natural Sciences and Math, and 5) Never Graduate/Drop Out. We provided the respondents a link where they could see a detailed listing of college majors (taken from various NYU sources), which described how each of the NYU college majors maps into our aggregate major categories. Before the official survey began, survey respondents were first required to answer a few simple practice questions in order to familiarize themselves with the format of the questions.

Expected earnings at age 30 were elicited as follows: "*If you received a Bachelor's degree in each of the following major categories and you were working FULL TIME when you are 30 years old what do you believe is the average amount that you would earn per year?*". We also provided definitions of working full time ("working at least 35 hours per week and 45 weeks per year"). Individuals were instructed to consider in their response the possibility they might receive an advanced/graduate degree by age 30. Therefore, the beliefs about earnings

we collected incorporated beliefs about the possibility of other degrees earned in the future and how these degrees would affect earnings. We also instructed respondents to ignore the effects of price inflation. The instructions emphasized to the respondents that their answers should reflect their own beliefs, and to not use any outside information.⁴

Our questions on earnings were intended to elicit beliefs about the distribution of future earnings. We asked three questions on earnings: beliefs about expected (average) earnings, beliefs about the percent chance earnings would exceed \$35,000, and percent change earnings would exceed \$85,000. The last two were elicited as follows: "*What do you believe is the percent chance that you would earn: (1) At least \$85,000 per year, (2) At least \$35,000 per year, when you are 30 years old if you worked full time and you received a Bachelor's degree in each of the following major categories?*"

We paid respondents a fixed compensation for completing the survey, and did not elicit respondents' beliefs using a financially incentivized instrument such as a scoring rule. This is because it is well known that proper scoring rules generate biases when respondents are not risk neutral (Winkler and Murphy, 1970). It should be pointed out that even if respondents are risk neutral, incentivized belief elicitation techniques are not incentive-compatible when the respondent has a stake in the event that they are predicting (the "no stake" condition in Karni and Safra, 1995), as is the case when reporting future earnings. In addition, Armantier and Treich (2011) show that beliefs are less biased (but noisier) in the absence of incentives. Finally, for self beliefs, we anyway do not have an objective measure against which their accuracy may be evaluated since we ask respondents for their individual self beliefs about future, unrealized, events.

2.3 Sample Selection and Descriptive Statistics

Our sample is constructed using the following steps. First, we drop 6 students who report that they are in the 4th year of school or higher, violating the recruitment criteria. Second, we drop 75 respondents (about 15 percent of the sample) whom we believe either made errors in filling out the survey or did not take the survey seriously. These include 21 students who report a change in graduation probabilities of greater than 0.5 in magnitude in any of the 5 major categories; 7 respondents who report full-time earnings below \$10,000 in any of the graduating majors, 2 students who report expected earnings of more than a million dollars; and 45 students who revise their self age 30 earnings by more than \$100,000 in any major category. This leaves us with a total of 420 respondents.

⁴We included these instructions: "*This survey asks YOUR BELIEFS about the earnings among different groups. Although you may not know the answer to a question with certainty, please answer each question as best you can. Please do not consult any outside references (internet or otherwise) or discuss these questions with any other people. This study is about YOUR BELIEFS, not the accuracy of information on the internet.*"

Table 2 shows the characteristics of our final sample. 36.5 percent of the sample (154 respondents) is male, 38.5 percent is white and 45.5 percent is Asian. The mean age of the respondents is about 20, with 40.5 percent of the respondents freshmen, 36 percent sophomores, and the remaining juniors. Three-fourths of the respondents completed the survey in under two hours, with 90% of all respondents completing the survey in three hours or less. The average grade point average of our sample is 3.5 (on a 4.0 scale), and the students have an average Scholastic Aptitude Test (SAT) math score of 701, and a verbal score of 685 (with a maximum score of 800). These correspond to the 93rd percentile of the corresponding SAT population score distributions. Therefore, our sample represents a high ability group of college students.

3 Earnings Beliefs and Belief Updating

In this section, we examine self beliefs about what each individual expects to earn in different majors, beliefs about population average earnings, and revisions in self beliefs following the information treatment.

3.1 Self Beliefs about Earnings

We first describe self beliefs about *own* earnings at age 30 if the respondent were to graduate in each major. The first column of Table 3 reports the average, median and standard deviation of the distribution of reported average self earnings in our sample at the Initial Stage. At the Initial Stage of the experiment all subjects were asked the same baseline set of questions. Looking across majors in column (1), we see that students expect the highest earnings (\$100,000) if they major in economics/business, and lowest if they do not graduate (\$37,500). Among the graduating majors, students expect the earnings to be lowest in humanities and arts (\$64,100). The median point forecast is lower than the mean self earnings for all majors, indicating that the distribution of point forecasts of future earnings is right-skewed. There is also considerable heterogeneity in self beliefs as indicated by the large standard deviations. The extent of heterogeneity can also be viewed in the top panel of Figure 2, which shows the belief distribution of our respondents if they were to graduate in economics or business. For example, in the economics and business category, the 5th percentile of the self belief distribution is \$50,000, the 50th percentile is \$90,000, and the 95th percentile is \$175,000. The second column of Table 3 reports self earnings for the subset of students who report to be either majoring or intending to major in that field. Compared to the beliefs for the full sample (column 1), this group of students has higher mean beliefs in most majors. This is consistent with observed sorting by ability and positive selection into majors based on expected earnings (Arcidiacono, 2004; Gemici and Wiswall, 2011).

As described above, we also collected data on the subjective distribution of future earnings. For this purpose, students were asked about the probability they would earn at least \$35,000 and at least \$85,000 at age 30 if they were to graduate in each major. Columns (3) and (4) of Table 3 present the average probabilities reported by students. While students believe that the likelihood of earning at least \$35,000 is fairly similar across the graduating majors (at least 0.75), the subjective likelihood of earning at least \$85,000 varies substantially across the majors, with students expecting the highest probability of that happening in the economics/business and engineering/computer science categories (mean probability exceeding 0.6 in both), and the lowest probability in humanities/arts (0.4) among the graduating majors. It is not surprising that students report very low probabilities for the occurrence of these outcomes in the no graduate major.

3.2 Population Beliefs about Earnings

As described above, at the beginning of the Intermediate Stage, we divided the subject pool into 4 randomly selected information treatment groups and asked corresponding baseline population beliefs questions before we provided the information treatment. We asked the following question for the randomly selected subset of respondents who were later assigned the Male Major Specific Treatment: "*Among all male college graduates currently aged 30 who work full time and received a Bachelor's degree in each of the following major categories, what is the average amount that you believe these workers currently earn per year?*" For another randomly selected group of respondents who were later assigned the Female Major Specific Treatment, we asked the corresponding question about female graduates.

Columns (5) and (6) of Table 3 report the mean, median and standard deviation of beliefs about US population earnings of men and women by the 5 major fields, reported by the two subsets of our sample who received the *Major Specific* (Male or Female) treatments. Self beliefs may differ from population beliefs for several reasons: Students might think that future earnings distributions will differ from the current ones, or students may have private information about themselves that justifies having different expectations. The difference between self and population beliefs therefore provides some suggestion of the student's belief of their own earnings advantage or disadvantage relative to the population average.

Looking across each of these columns, we see that population beliefs follow the same pattern as self beliefs (columns 1 and 2), with students believing population earnings to be highest in the economics/business and engineering/computer science categories, and lowest in humanities/arts and the not graduate categories. Compared to self earnings beliefs, students report similar population beliefs for all fields, except for economics/business and natural sciences, for which self beliefs are significantly higher. It is also interesting to note that students accurately perceive

a wage gap in favor of men in all fields, with average earnings for males exceeding those for females. As with the self beliefs, the distributions of population beliefs are skewed right and show substantial heterogeneity. For example, in the Female Major Specific treatment, the median of the population beliefs for average earnings of female graduates in economics/business is \$75,000, while the 5th percentile is \$50,000 and the 95th percentile is \$130,000.⁵

For the other, more general, information treatments, respondents randomly assigned to the All Individuals Treatment were asked the following question about their population beliefs: "*Among all individuals (college and non-college graduates) currently aged 30 who work full time, what is the average amount that you believe these workers currently earn per year?*" Those in the College Treatment were asked about earnings of all college graduates currently aged 30 and working full time. Column (7) reports the population beliefs of respondents in these *General* treatments. Mean population beliefs in the All Individuals Treatment are substantially lower than those for all majors, except the no graduate category. This demonstrates that, at least in the aggregate, respondents accurately believe that college graduates have higher average earnings than the full population. Moreover, compared to population beliefs in the major specific treatments, the standard deviation is quite low, reflecting much lower heterogeneity in population beliefs about average unconditional average earnings across all individuals. In the College Treatment, the mean belief reported for college graduates is higher than that reported for humanities/arts in the Major Specific treatments, accurately reflecting that the college graduate population includes individuals with higher earning majors. As with all of the population beliefs about college major specific beliefs, there is substantial heterogeneity in the population beliefs about college graduates.

3.3 Errors in Population Beliefs

3.3.1 Absolute Value of Errors

In the case of the groups receiving the Male and Female *Major Specific* treatments, the comparison of population beliefs (columns 5 and 6 of Table 3) in a given major with true population earnings (reported in Table 1) in the corresponding major shows that average student beliefs over-estimate the true average population earnings for all fields, except male earnings with the no-degree major. Columns (8) and (9) of Table 3 report the mean absolute error, defined as the absolute value of the difference between the true and perceived population earnings. We use the absolute value of the error here to assess the magnitude of the errors, without positive and negative errors cancelling out. The absolute errors are substantial, varying from a mean of \$15,000 for male no-degree workers to \$31,275 for male workers who graduated in eco-

⁵The distribution statistics for other majors and sub-populations is available on request.

nomics/business. Students also have considerable errors about the population average earnings for all workers and for college educated workers (Column 10). The absolute error in population beliefs is \$11,147 for all workers and \$21,000 for college educated workers.

3.3.2 Raw Errors

To provide a sense of the heterogeneity in population errors for at least one major category, the middle panel of Figure 2 shows the distribution of raw population errors regarding full-time females' earnings with an economics or business degree. Here we define raw errors as truth-belief, such that a negative error indicates over-estimation of the truth, and a positive error indicates under-estimation of the truth. Reflecting the dispersion in baseline beliefs, there is considerable heterogeneity in the level and sign of the errors, with non-trivial numbers of students making both positive and negative errors in all categories. While the median of this error distribution is -\$14,270 (i.e., over-estimation of population earnings by \$14,270), the 5th percentile is -\$69,270 and the 95th percentile is \$10,730 (under-estimation).

3.3.3 Heterogeneity in Population Errors

Are errors systemically related to observable respondent characteristics? Table 4 explores the heterogeneity in absolute population errors by treatment type and individual characteristics. The first column restricts the sample to respondents who received the General Treatments, and regresses the absolute error in population beliefs on a set of observable characteristics of individuals. The constant term indicates the mean absolute error in the All Individuals treatment. The absolute error in the College Treatment is substantially greater than the error in the All Individuals treatment. With regard to individual observables, we see that high ability respondents – defined as those with a score of at least 1450 out of 1600 on the SAT – make substantially smaller errors. Relative to freshmen, students in their sophomore or junior years also make significantly smaller absolute errors. These patterns of smaller errors for upperclassmen is consistent with the survey by Betts (1996). We also find that Asian respondents have substantially larger errors.

The estimates so far mask the heterogeneity in population errors by whether the error is positive (an underestimate of population earnings) or negative (over-estimate). For example, it is not clear whether the smaller absolute errors by high ability respondents are a consequence of less underestimation or overestimation of population earnings, or both. In columns (2) and (3), we restrict the sample to respondents with positive and negative errors, respectively, and regress the absolute error on the same set of controls. We see that the smaller absolute errors by high ability respondents and upperclassman (sophomores and juniors) are primarily driven by smaller negative errors, i.e., these groups make smaller over-predictions, on average, relative to

their counterparts. Similarly, the large absolute errors by Asians and in the College Treatment seem to be driven by large negative errors. As indicated in columns (1)-(3), females have larger absolute errors than male respondents, though these differences are not statistically significant.

Column 4 of Table 4 reports the OLS estimates of the same regression on the sample which received the Major Specific Treatments. Mean absolute errors are significantly larger in economics/business and engineering/computer science relative to the excluded major category (humanities/arts). Here, we see that female and high ability respondents make significantly larger absolute errors. Estimates in columns (5) and (6) suggest that these are driven by larger negative errors (i.e., larger overpredictions) by these groups, on average. None of the other individual characteristics are significantly different from zero at levels of significance of 95% or higher.

3.4 Revisions of Self Beliefs

We next explore how self beliefs are revised as the student respondents receive the information treatments. Recall that our experimental design has two rounds of information treatments in the intermediate and final stages.

The first column of Table 5 reports the mean and standard deviation of the distribution of percent revisions (intermediate-initial stage) in self beliefs about earnings. There is considerable heterogeneity in the updating of self beliefs across majors. The average of the percent revisions distribution varies from about -5 percent (downward revision) in economics/business to +55 percent (upward revision) in the no-degree category. As indicated by the standard deviations, within categories there is considerable heterogeneity. The bottom panel of Figure 2 shows the dispersion in students' revisions for earnings in economics/business in the Female Major Specific Treatment: the 5th percentile of the earnings revision is -50 percent, the 50th percentile is -15.48 percent, and the 95th percentile is +30 percent.

Columns (2) and (3) of Table 5 show the revisions of self beliefs in the combined Major Specific (Female and Male) and General (All Individuals and College) treatments, respectively.⁶ The revisions in the two treatment groups are statistically different for the no-degree category, with much larger upward revisions for respondents receiving the Major Specific treatment.

While the other revisions are statistically similar, it is interesting that the mean revision in engineering/business is larger in magnitude in the General treatment (downward revision of 8.02%) than in the Specific treatment (downward revision of 1.88%). Recall that in the

⁶For much of the remaining analysis, we pool the responses in the All Individuals and College treatments into the "General" treatment, and the Female and Male Major Specific treatments into the "Major Specific" treatment. This is because the results are qualitatively similar when we analyze the All Individuals and College treatments separately, and when we analyze the Female and Male Major Specific treatments separately. Pooling in this way keeps the tables simple.

General treatment, students receive information about earnings for either all individuals or for college graduates. This finding would seem to contradict the hypothesis that the General treatment is less relevant to individual self beliefs than the Major Specific treatment. However, if individuals respond to the overall level of the information relative to self beliefs and don't find the information provided in the General treatment irrelevant, the fact that the General treatment provides lower values for average earnings may cause a greater downward revision than the Major Specific treatment.

We next turn to the second round of information treatments. In the final stage, all respondents were provided with the information from all 4 treatments. At the start of the final stage, all students have the same information, although they have received this information in a different order. Column (4) of Table 5 shows that, as expected, revisions are generally larger in magnitude in between the initial and final stage than between the initial and intermediate stage. Moreover, being exposed to different information in the intermediate stage has an anchoring effect on respondents' revisions: Mean revisions are larger in magnitude for respondents who were assigned to the General treatment in the first stage, with the revisions being statistically different for three of the five major fields. The revision patterns in Table 5 suggest that students anchor their self beliefs to the statistics provided to them, even when, arguably, the information provided to them in the General treatment is less relevant. This is consistent with Tversky and Kahneman (1974), who find that the initial information provided to respondents has an anchoring effect on their choices, and that irrelevant information can affect behavior.

3.5 Self Belief Updating and Population Errors

We next examine whether errors in population beliefs regarding earnings relate to revisions of self beliefs in the intermediate stage. If students perceive a link between population earnings and self beliefs, then revealed errors in population beliefs should be systematically related to revisions of self beliefs. For example, if a respondent underestimates the population earnings (i.e., the error in population earnings is positive), the respondent should revise her self beliefs upwards upon receipt of information.

The updating patterns in Table 5 and population beliefs reported in Table 3 hint towards a logical positive relationship between the two. We see that students, on average, revise downward their self earnings beliefs the most in economics/business, which is the field with the highest average over-estimation in population earnings (compare population beliefs in columns (5) and (6) of Table 3 with true population earnings in Table 1). Similarly, self beliefs are revised upward the most for the not graduate category, which is the field with the largest under-estimation in population earnings.

To explore the link between revisions in beliefs and errors, we estimate a series of reduced-

form regressions using the intermediate - initial belief updating. Our dependent variable is the change (intermediate - initial) in age 30 self earnings reported by the respondent for each of the college majors. Our experiment provides respondents with various *packages* of information on earnings and labor supply for a given group. The randomly assigned group of students receiving the General treatments were provided with the population earnings for either the whole full-time working population or for college graduates. For the Major Specific treatments, students were provided with information for full-time workers with each of the various majors.

The first column of Table 6 shows that overall the error in the population earnings is positively related to self belief updating. This is evidence of logical updating in response to our information treatments. An error of \$1,000 in population earnings results in a revision of \$34 in self earnings. While the estimate is very precise (significantly different from 0 at the 1% level), the relatively "inelastic" response of revisions in self beliefs to population errors suggests that self beliefs about earnings are not entirely linked to the type of public population information we provide. In general, heterogeneous private information on the abilities and future earnings prospects of individuals may cause individuals to have an inelastic response to population information. Because our estimate is a combination of different treatments and individual responses, we next unpack this estimate and explore heterogeneity in updating by information type and individual characteristics.

3.5.1 Heterogeneity in Updating by Information Type

Column (2) in Panel A of Table 6 shows that, as one would have expected, it is the information revealed in the Specific treatments which leads to significant revisions, while no effect is found in the General treatments. The effect of the major specific information is more than twice as large as the overall pooled effect reported in column (1). This provides evidence that the quality or specificity of the information matters. Given the dependent variable is beliefs about earnings in each major, the major specific information evidently provides higher quality information with larger errors revealed by this information causing much larger belief updating.

Another dimension of information type is the direction of the errors revealed by the information. Column (3) shows that response to information is asymmetric. A positive error, i.e., under-estimation of population earnings, results in significant updating: An under-estimation of population earnings by \$1000 results in an upward revision in self earnings of \$181. On the other hand, we do not find a (statistically or economically) significant response in instances where the error is negative. Therefore, self beliefs are responsive to the information only when it is good news, i.e., when the respondent is informed that population earnings are higher than her prior beliefs. This pattern of asymmetric updating is consistent with Eil and Rao (2011), and Mobius et al. (2011), who find beliefs to be relatively more responsive to good news (where

good news is defined as feedback that improves one’s self-image). Table A1 in the Appendix shows an additional set of regressions in which the error is interacted with various treatment characteristics, and the effect of the error is allowed to vary depending on whether the information is positive or negative. We see that, in almost all the specifications, all the interactions with a positive error are statistically significant, while the coefficient on negative error is statistically different from zero in a few cases only. For example, the last column in Table A1 shows that respondents significantly revise their earnings beliefs in all major categories (excluding humanities/arts) in response to good news (i.e., when they are informed that population earnings are higher than their prior beliefs), and that the response to negative news is not statistically different from zero at the 95% level or higher for any major category.

3.5.2 Heterogeneity in Belief Updating by Individual Characteristics

We next explore the extent of heterogeneity in the relationship between population errors and earnings beliefs using a set of observable characteristics for respondents. In column (4) of Panel B of Table 6, we include an indicator for female gender, and interactions of the error with female and male indicators. Estimates for the interaction terms indicates that men are more responsive to their errors about population earnings than women. The response by both men and women to errors is positive (indicating logical updating for both groups), but we estimate the response is 3 times larger for men than women, although given the precision level we cannot reject a hypothesis that these responses are closer.

The second regression in Panel B investigates whether responses to the information treatment differ by the grade level of the student. The interaction terms indicate that freshman, sophomores, and juniors all update logically to errors. Although each coefficient is statistically significant at least at the 5 percent level, we cannot reject the hypothesis that the responsiveness to errors is the same across the groups. Finally in the third column of Panel B we investigate whether there is heterogeneity in updating by the ability of the student, where we classify students as high ability if they have an SAT score greater than 1450 (30% of our sample respondents fall in the high ability group). The interaction terms reveal that high ability students are more responsive to errors they make, where the error coefficient response is nearly twice as large as that for low ability students. These estimates are consistent with either a hypothesis that high ability students are simply paying more attention to the information treatment we provide and are able to process the information better, and/or that high ability students perceive a stronger link between population earnings and their own earnings.

3.5.3 Non-Parametric Analysis

To further explore the relationship between errors and belief updating, we turn to a non-parametric analysis using a local linear regression. Figure 3 shows the local linear regression of self earnings revisions on population errors, for the General and Specific treatments separately. Several points are of note. First, except for very negative errors, the response in the General treatments is not statistically different from zero. Second, the response of revisions to errors is asymmetric, with a steeper slope for positive errors, in particular in the Specific treatments. Third, even conditioning on direction of error, the relationship does not seem to be linear. In the next section, we explore the heterogeneity in updating in more detail.

4 Model of Belief Updating

4.1 Bayesian Benchmark

In this section, we examine a formal model of belief updating. In particular, we explore how a respondent’s belief about expected self earnings reported in the intermediate stage—the *Observed posterior*—compares to a posterior if the updating process were approximately Bayesian, i.e., the *Bayesian posterior*. Our objective is to use our information experiment to construct a Bayesian benchmark level of updating for each respondent and then compare the actual observed updating for each individual to this benchmark. If the updating process were Bayesian, the posterior would be given by:

$$\text{Post}_{im} = \frac{\alpha_i}{\alpha_i + \beta} \text{Prior}_{im} + \frac{\beta}{\alpha_i + \beta} \text{Info}_m, \quad (1)$$

where Post_{im} is respondent i ’s belief in the Intermediate stage about expected self earnings in major m ; Prior_{im} is the belief reported in the Initial stage about expected self earnings in major m ; Info_m is the information treatment that i is provided about earnings in major m between the Intermediate and Initial stage; α_i is the individual specific precision of the prior; and β is the precision of the revealed public information. In our setup, Info_m depends on the treatment the respondent is assigned. But since the information is public, the precision associated with this information is homogenous.⁷

We form the precision of the prior on average earnings, α_i , using the self reported uncertainty in future earnings: $\alpha_i = \frac{1}{\text{Var}(\text{Prior}_{im})}$. The precision of the information is similarly formed using

⁷In the General treatments, since the respondent is provided with population earnings of either all workers or college graduates, there is only one piece of new information that is observed, and hence $\text{Info}_m \equiv \text{Info}$. In the Major Specific treatments, we assign the respondent the information about population earnings in the major corresponding to the self beliefs about earnings in each major. Therefore, in this case, Info_m varies by major.

the revealed distribution of population earnings: $\beta = \frac{1}{\text{Var}(\text{Info}_m)}$. With data on each of the components in (1), one can compute the Bayesian posterior, $\text{Posterior}_{im}^{Bayes}$, as follows:

$$\text{Post}_{im}^{Bayes} = \frac{\frac{1}{\text{Var}(\text{Prior}_{im})}}{\frac{1}{\text{Var}(\text{Prior}_{im})} + \frac{1}{\text{Var}(\text{Info}_m)}} \text{Prior}_{im} + \frac{\frac{1}{\text{Var}(\text{Info}_m)}}{\frac{1}{\text{Var}(\text{Prior}_{im})} + \frac{1}{\text{Var}(\text{Info}_m)}} \text{Info}_m. \quad (2)$$

To compute the variance of future earnings, recall that students were asked about the probability of earning at least \$35,000 and \$85,000 at age 30 if they were to graduate in each major, and they were also provided with information about the distribution of population earnings. We fit the responses of the respondent to the questions about the chance of earning more than \$35,000 and more than \$85,000 per year to a log-normal distribution, and obtain an estimate of $\text{Var}(\text{Prior}_{im})$ for each major and individual. Similarly, we use the empirical likelihood of earning more than \$35,000 and \$85,000 in the population – information that students were provided with in the treatments – to obtain an estimate of $\text{Var}(\text{Info}_m)$ for each major. Note that the latter variance is the same for each respondent since everyone within a treatment group receives the same information. After computing the Bayesian posterior, we can then investigate how the observed posterior, $\text{Post}_{im}^{Observ}$, i.e., beliefs reported in the Intermediate stage, compare with the Bayesian benchmark.

There are two important differences between our experimental design and the textbook case of Bayesian updating. First, the information we reveal may already be known by some respondents. As shown in Section 3, this is not the strictly the case for all of our respondents since all individuals had some errors in their beliefs about the population earnings distribution. However, the distribution of errors in population beliefs, discussed above, shows that there is substantial heterogeneity in how informative the information provided to respondents was. A second key difference in our experimental design from the textbook case is that we reveal *public* information but ask individuals about their *private* beliefs about themselves. Individuals can differ in how relevant they believe the population distribution of earnings is to their own future earnings. For example, if we observe that a respondent does not revise her beliefs in response to the information, even after controlling for her priors about the information, this could either imply biased, non-Bayesian, updating, or that the respondent simply did not find information on population beliefs relevant for self beliefs.

The difference between the interpretation of the Bayesian updating we analyze here and the textbook case is a consequence of our experimental setup. In typical studies of belief updating (Tversky and Kahneman, 1974; Grether, 1980; Viscusi and O’Connor, 1984; and Viscusi, 1997; Cameron, 2005; El-Gamal and Grether, 1995; Eil and Rao, 2011; Mobius et al., 2011), respondents are provided with signals about the same quantity over which revision of beliefs are being analyzed. For example, in the frameworks used by Eil and Rao (2011), Mobius

et al. (2011), and Grossman and Owens (2011), respondents are revising their beliefs about either their own intelligence or beauty, and receiving feedback about the same underlying entity for which beliefs are being reported. That is not the case in the design used in our study: We observe belief updating about future self earnings, formed from both past private and public signals, but the signals that students receive in our experiment are about population beliefs. Our study design was motivated by the kinds of information that are typically available to students when making real world schooling choices. The kind of information that we provided to respondents is precisely the kind that are available in mainstream sources.⁸ Information along similar lines has been provided to students in other contexts, and it has been shown to have an impact on actual schooling choices (Jensen, 2010; Nguyen, 2010).

4.2 Are Students Bayesian?

First, we use our Bayesian benchmark as a device to characterize the heterogeneity in belief updating. Figure 4 plots the observed updating in average self earnings ($\text{Post}_{im}^{\text{Observ}} - \text{Prior}_{im}$) and the Bayesian revision ($\text{Post}_{im}^{\text{Bayes}} - \text{Prior}_{im}$) by major category. If students are Bayesian, all the points would be along the 45-degree line. That is clearly not the case. We split the data by General and Specific treatments. Also shown are the fitted lines from an OLS regression of observed revision on Bayesian revision. The fitted lines are flatter than the 45-degree line for both treatments in all major categories. This indicates that, on average, students respond less to the information than the Bayesian benchmark. The figures also show less sensitivity to the information in the General treatments, at least for some of the majors.

4.2.1 Characterizing Belief Updating Heuristics

As indicated by the scatterplot in each of the panels of Figure 4, there is substantial heterogeneity in students' response to information, and it appears that some of this updating is non-Bayesian according to our benchmark. We next characterize the updating heuristics used by our respondents. We classify each respondent to an updating type, depending on how her observed posterior compares with our Bayesian benchmark posterior. We use five possible heuristics to classify a respondent's updating. A respondent's type is: (1) Bayesian if her posterior belief is within a band around the Bayesian posterior; (2) Alarmist if, relative to the Bayesian benchmark, the response is more exaggerated; (3) Conservative if she updates in the right direction but less than a Bayesian; (4) Confused if the updating is in the wrong direction,

⁸For example, the Chronicle of Higher Education lists earnings by major and subject area: <http://chronicle.com/article/Median-Earnings-by-Major-and/127604/> (accessed September 10, 2011). Similarly, the BLS publishes a yearly handbook with information on earnings, job prospects, and working conditions etc. at hundreds of different types of jobs in the Occupational Outlook Handbook (<http://www.bls.gov/oco/>).

i.e., inconsistent with the direction prescribed by Bayesian updating; and (5) Non-Updater if there is no response to the information. We borrow this nomenclature from the previous psychological and experimental economics literature on belief updating (Grether, 1980; Kahneman and Tversky, 1982; El-Gamal and Grether, 1995). This literature classifies individuals as using the Conservative heuristic if they fail to sufficiently adjust their beliefs in light of new information, and classifies individuals as using the "Representative" heuristic if they rely too heavily on recent information; we instead use the term "Alarmist" to refer to such updating.

In the case where $\text{Posterior}_{im}^{Bayes} > \text{Prior}_{im}$, i.e., a respondent should revise beliefs upward on receipt of information, we classify the respondent's type, Type_i , as follows:

$$\text{Type}_i = \begin{cases} \text{Bayesian} & \text{if } |\text{Post}_{im}^{Bayes} - \text{Post}_{im}^{Observ}| \leq \text{Band}_{im} \\ \text{Alarmist} & \text{if } \text{Post}_{im}^{Observ} > \text{Post}_{im}^{Bayes} + \text{Band}_m^+ \\ \text{Conservative} & \text{if } (\text{Post}_{im}^{Observ} \geq \text{Prior}_{im}) \ \& \ (\text{Post}_{im}^{Observ} < \text{Post}_{im}^{Bayes} - \text{Band}_{im}^-) \\ \text{Confused} & \text{if } \text{Post}_{im}^{Observ} < \text{Prior}_{im} \\ \text{Non-Updater} & \text{if } \text{Post}_{im}^{Observ} = \text{Prior}_{im}, \end{cases} \quad (3)$$

where Band_{im} is a band around the Bayesian posterior within which the respondent is considered to be Bayesian. The upper end of the interval, Band_m^+ , is 10% of the sample standard deviation in beliefs reported at the baseline, $\overline{\text{std}(\text{Prior}_m)}$. The lower end of the band, $\text{Band}_{im}^- = \min\{0.10*\overline{\text{std}(\text{Prior}_m)}, 0.5*|\text{Posterior}_{im}^{Bayes} - \text{Prior}_{im}|\}$. We choose a non-symmetric band with a tighter lower bound because, in cases where the sample standard deviation in beliefs is very large, we may be left with no conservative types. This criteria ensures that there are always some respondents who are classified as conservatives. Figure A1 shows a graphic representation of this classification. For downward revisions, the updating type is defined analogously. This classification, obviously, involves some subjectivity in how the band is defined. An alternative criteria that involves no subjectivity is to classify any insufficient (excessive) response relative to the Bayesian benchmark as conservative (alarmist). Reducing the bandwidth around the Bayesian benchmark to zero would ensure that almost all of the sample is classified as non-Bayesian updaters.

4.2.2 Heuristics for Own Major

Using the classification in equation (3), we first determine the distribution of respondents' types based on their earnings updating in their own (intended) major. Table 7 reports the distribution of types separately for the Specific and General treatments. Looking across column (1), we see that nearly a fifth of the sample respondents are non-updaters, i.e., they don't change

their self beliefs on receipt of information. Among the respondents who revise their beliefs, the most common heuristic is either Bayesian or excessive (Alarmist) updating, with about a fifth of the sample using each of these heuristics. A substantial proportion of respondents are conservative in their updating. Finally, 15-20% of the respondents update in a way that cannot be rationalized by our updating model.

Because our design included different kinds of public information, from general information about earnings for all workers to more specific information about earnings by particular gender and major, some individuals could find the general information not relevant but the majors specific information relevant. We might expect then the relative share of Conservatives to be larger in the General treatment. That is, however, not the case: The type distribution in the Specific and General treatments is very similar, and the relative share of Conservatives is only marginally greater in the General treatments.

The remaining columns of Table 7 report the distribution of heuristics for various subsamples. Columns (2) and (3) show the gender-specific distribution of types. In the Major Specific treatments, women, relative to men, are more likely to update; and conditional on updating, more likely to be Bayesian or Alarmist. The reverse patterns are observed for the General treatments. Overall, we cannot conclude that there are any systematic differences by gender. Columns (4) and (5) show the type distribution for freshmen and upperclassmen (sophomores and juniors). Two differences between the two groups are of note. First, 25-30% of upperclassmen do not update versus about 20% of freshmen. This suggests that through their more extensive college experience, upperclassmen have gathered more private information about their own future earnings. Second, conditional on updating, the most common heuristic for freshmen is Alarmist updating, while upperclassmen are most likely to be Bayesian. We do not find any notable differences in updating in the General treatments by ability, but do find that high ability respondents are more likely to react excessively to information in the Specific treatments (columns 6 and 7 of Table 7).

The last two columns of Table 7 show the distribution of types for respondents with positive errors (i.e., those who underpredict population earnings) and negative errors, respectively. The most common heuristic for respondents who make both negative and positive errors is Alarmist updating, i.e., they respond excessively to the information. However, we see that students are more than twice as likely to be conservative in their updating when their population error is negative compared to when it is positive. Therefore, this suggests that there is valence-based updating. Students tend to react (excessively) when the information is good news, i.e., when they receive the news that population earnings are higher than their priors.

4.2.3 Heuristics for Other Majors

We also collected data on earnings revisions in the four other major categories that the student reports is not their primary "intended major." Updating patterns in these other major categories may be different from those in own major if respondents have different levels of private information in other majors. As in the case of own major earnings, we first compute the Bayesian posterior for earnings in each of these other majors using equation (2), and then classify the respondent's type in each major using the classification outlined in equation (3). In order to analyze how updating heuristics in these other majors compare with the heuristic in own major, we collapse the respondent's type in these other majors into one type, which includes each of the types in equation (3) as well as an additional "Mixed" type, which denotes updating when the respondent uses a mixture of heuristics across the 4 majors. In order to complete this categorization, we created an algorithm for how different combinations of types map into a single type. Details of this procedure are provided in the Appendix.

Table 8 reports the joint distribution of types in own major and other majors. Each cell reports the percentage of the sample that falls in that group. With regard to updating heuristics in other majors, the most common heuristic is Mixed, i.e., respondents use a combination of heuristics when updating earnings in the four major categories. In the Specific treatments, about 35% of the respondents can be classified as either Bayesian or Alarmist in the other majors, versus about 20% of respondents in the General treatments. On the other hand, nearly 30% of respondents either do not update or update Conservatively in the General treatments, versus less than 10% of the respondents in the Specific treatments.

If we restrict our sample to respondents who use a heuristic other than Mixed in other majors, an interesting pattern stands out: Students are more likely to use the same heuristic in the other majors that they use in their own major, as indicated by the larger values in the diagonal cells. For example, students who are Alarmist in their own major are twice as likely to be Alarmist in their updating in other majors. This suggest that there is some consistency in updating heuristics across majors.

4.3 Robustness Checks: Alternative Specifications

4.3.1 Actual Errors and Updating Heuristic

As shown in Section 3, there is substantial variation in our sample in population errors, i.e., in the difference between perception of population earnings and true population earnings. In the analysis above, we do not use the errors that students make in population earnings when categorizing their updating heuristics. This could be problematic for the interpretation of our results. For example, it could be the case that students whom we classify as Conservative

in their updating had fairly accurate expectations of population earnings, which were then already incorporated in their self beliefs. Therefore, we find that they react less than the Bayesian amount to the provided information simply because they already knew the information treatment. Conversely, we may simply be classifying students who had very inaccurate perceptions of population earnings as Alarmists, since presumably the information that we provide would be most valuable to that group.

In order to test whether that is the case, Table 9 regresses the absolute value of the respondents' population errors in each major category onto their updating type in that major. More specifically, we regress the absolute value of the error onto a constant term and dummies for each of the other heuristics excluding Bayesian. The constant term shows the mean absolute value of the error for respondents who are classified as Bayesian (the omitted category), while the parameter estimates on the dummies are the additive mean errors for students who are classified as using that heuristic. In column (1) of the table, we pool all majors together, i.e., we have 5 observations per respondent.⁹ The mean absolute population error for a Bayesian updater is \$16,124. Except for the coefficient on Non-Updater, none of the other dummies are statistically significant. The column also reports the p-value of a test for the joint significance for all the covariates excluding the constant term. We reject the null that these covariates are jointly significant, indicating that errors are similar in magnitude, regardless of the heuristic used by the student. These results suggest that our classification procedure is not a mere consequence of the magnitude of the error that the student makes.

The remaining columns of Table 9 report the same regression as in column (1), but for each major separately. None of the parameter estimates on the terms excluding the constant are significant at levels of 95% or higher. We reject the null of the joint significance of these covariates for each of the major categories. We conclude that our classification of updating behavior is not systematically related to population errors.¹⁰

⁹We pool the Major Specific and General treatments together since results are qualitatively similar in both cases (results available from the authors upon request).

¹⁰A possible alternate is to use the population error – which is a measure of the relevance of the information – directly in the Bayesian updating model. That is, to use population error to proxy for *Info* in equation (1). However, since the Bayesian posterior is a convex combination of the prior and the signal, using the population error is not very meaningful. To illustrate this, consider a respondent with self beliefs of \$75,000 and population beliefs of \$100,000. If the true population earnings are \$125,000, this respondent has a population error of \$25,000. Using the population error instead of population earnings in the updating model, the Bayesian posterior would be a convex combination of self beliefs (\$75,000) and population error (\$25,000), which at most can be \$75,000. However, if the respondent finds information about population earnings relevant for self earnings, she should be revising her self earnings upwards. Therefore, we do not directly use the population errors when classifying updating heuristics.

4.3.2 Effective Information and Updating

As another robustness check of our classification algorithm, we analyze the relationship between each of the updating types and response to *effective* information. We define effective information as the information content in the information that we provide to the respondent, i.e., $\text{Info}_{im}^{Effect} = \text{True Population Earnings}_m - \text{Beliefs about Population Earnings}_{im}$. This is analogous to how we define population earnings error. We define the effective response, R_{im}^{Effect} , for respondent i in major m as:

$$R_{im}^{Effect} = \frac{\text{Post}_{im}^{Observ} - \text{Prior}_{im}}{\text{True Pop. Earnings}_m - \text{Beliefs about Pop. Earnings}_{im}}.$$

The effective response, R_{im}^{Effect} , is essentially the elasticity of self earnings revision in response to effective information. For logical updating, this metric should be positive. If our updating model accurately characterizes the respondent's heuristics, we should observe that the effective response is larger (smaller) for respondents who we classify as Alarmists (Conservatives), relative to someone classified as a Bayesian.

Another reason for this check is to understand the updating of respondents who we categorize as "Confused". We define Confused as those respondents who update in a direction opposite to that prescribed by Bayesian updating. For example, consider a male respondent who reports average self earnings in Economics to be \$50,000, and is then informed that average population earnings in Economics are \$74,542. Our updating model would imply upward revision in self earnings, with the magnitude of the revision depending on the uncertainty in the self earnings distribution. However, if the respondent's prior belief about population earnings in Economics were \$100,000, then this information—which reveals to the respondent that actual population earnings are lower than his priors—should cause the respondent to revise downward. While this updating is rational, our belief-updating model would categorize such a respondent as confused. Note that in this stylized example, the effective response of this respondent, $R^{effective}$, would be positive. Therefore, if we find that, the effective response of respondents whom we classify as confused is positive, then such updating is rational.

Table 10 reports the median effective response, R^{Effect} , by updating heuristic. The first row pools all the majors together and shows that the median effective response is 0.98 for Alarmists, compared to 0.41 for Bayesians. The response to effective information is unit elastic for Alarmists, and inelastic for Bayesians and Conservatives. On the other hand, the median (and mean) effective response for Confused is negative. That is, respondents whom we categorize as Confused are updating, on average, in a way that cannot be rationalized even after controlling for the information content of the signals that they receive. There is substantial variation in effective response as indicated by the large standard deviations. To test for whether the

distribution of R^{Effect} varies statistically between Bayesians and the other updating heuristics, the table also reports non-parametric tests for equality of the means and medians, as well as the Kolmogorov-Smirnov test for the equality of distributions. We find that estimates of R^{Effect} for Confused and Alarmists are statistically different from those of Bayesians. The remaining rows of the table show the corresponding statistics separately by major, and the same patterns emerge.

Overall, this shows that our updating model and classification of heuristics is quite reasonable. Respondents whom we characterize as Confused are, on average, updating in a manner that cannot be rationalized even if we control for their priors about population beliefs. Similarly respondents classified as Alarmists have a significantly higher effective response, compared to Bayesians and Conservatives.

5 Discussion: Behavior and Welfare Gains from Information Revelation?

Next, we explore the extent of the updating behavior by assessing whether the earnings updating spills over into beliefs about future actions, such as the student's future choice of major. We also assess whether there is evidence of welfare gains as result of our information treatments. In addition, we conduct an exercise to see if welfare is higher if individuals are all assumed to update in a Bayesian fashion, rather than as we observe them.

5.1 Major Choice Beliefs

A natural question to ask is whether our information treatments have an impact on students' beliefs about their future choice of college major.¹¹ Recall that our respondents are current college students, the majority of whom are freshman or sophomore students. Along with questions on earnings beliefs, our survey also asked respondents to provide the expected future percent chance (0 – 100) they would graduate in each of the 5 different major categories.¹² These questions about major choice were asked at all 3 stages of the survey, before and after the information treatments. For each respondent i , we calculated the absolute value of the change in the percent chance of graduating with each major m as $|prob_{im}^{post} - prob_{im}^{prior}|$, where

¹¹Wiswall and Zafar (2011) explore this issue in detail using the experimentally-generated panel of beliefs and probabilistic choices to estimate a rich model of college major choice without imposing any parametric assumptions on the taste distributions.

¹²Self beliefs about the probability of graduating with a major in each of the categories were elicited as follows: "What do you believe is the percent chance (or chances out of 100) that you would either graduate from NYU with a major in the following major categories or that you would never graduate/drop-out (i.e., you will never receive a Bachelor's degree from NYU or any other university)?"

$prob_{im}^{prior}$ is the initial stage belief about the probability of graduating in major m , prior to any information revelation; and $prob_{im}^{post}$ is either the intermediate or final belief, after information revelation.

Table 11 reports various statistics for the distribution of beliefs about graduating with different majors. About half of all respondents changed their beliefs about the percent chance they would graduate with a particular major. The mean of the absolute value of the change varies from 3.65 to 7.28 percent for the college major categories, with small mean changes of 1.61-1.67 for the not graduate category. For all majors, the mean change is largest from the initial to intermediate stages, but there is still additional updating in beliefs at the final stage after the second round of information treatments. With the large standard deviations (relative to means) we see evidence of substantial heterogeneity in the responsiveness of college major beliefs to the information treatments. We conclude that the information treatments we provided were meaningful enough not only to shift beliefs about self earnings but also for some individuals to update their expected probabilities of completing particular types of degrees.

5.2 Welfare

To provide some sense of the magnitude in updating our information treatments induced, we next provide a measure of welfare changes caused by the information treatment. In general, we would expect that at least some respondents to our survey are better off through exposure to previously unknown information. As discussed above, many of the individuals in our survey respond to the information treatments by updating their beliefs about future college major choices. Under the assumption that earnings are the main determinant of college major choice, we can compute the welfare change for respondent i as a result of our information experiment as follows:

$$\Delta \text{Welfare}_i \equiv \sum_m (prob_{im}^{post} * earn_{im}^{post} - prob_{im}^{prior} * earn_{im}^{post}), \quad (4)$$

where $prob_{im}^{post}$ ($prob_{im}^{prior}$) is the probability reported by i of majoring in major m after (before) the information on population earnings is provided to them, and $earn_{im}^{post}$ is individual i 's updated beliefs about earnings in major m . $\sum_m (prob_{im}^{post} * earn_{im}^{post})$ is expected earnings after the information treatment, and $\sum_m (prob_{im}^{prior} * earn_{im}^{post})$ is expected earnings if the individual were to maintain the same college major choices as before the information treatment. $\Delta \text{Welfare}_i = 0$ if the survey participant does not update her expected future major choices at all. $\Delta \text{Welfare}_i > 0$ if the respondent updates her expected future major choices in such a way that her expected earnings increase. While defining welfare on the basis that age 30 earnings are the only determinant of major choice is clearly restrictive (Arcidiacono, 2004; Zafar, 2010;

Beffy et al., 2011; Gemicci and Wiswall, 2011), the point of this exercise is simply to provide some sense of the magnitude of the change in students' choices using earnings as the metric.

The first column of Table 12 shows that the mean welfare change is \$432 in our sample: as a result of our information experiment, expected earnings at age 30 increase by \$432 due to the induced shift in expected college major choices. The majority, but not all of the change in expected earnings occurs between the initial and intermediate stage as the mean welfare change at the intermediate stage is \$331. Around 75 percent of respondents had non-negative changes in welfare ($\Delta \text{Welfare}_i \geq 0$) and the median change in welfare is zero since around half of all respondents do not change their choice probabilities. The increase in welfare, measured using expected earnings, is a consequence of some respondents adjusting their anticipated major choices as a result of the information treatments. While, on average, our information treatment increases welfare defined as perceived monetary returns to majors, an important question from a policy perspective is whether these gains will be actually realized. Since student outcomes are not observable, this is not directly testable.

5.3 Imposing Bayesian Updating

As we show in the previous section, there is considerable heterogeneity in belief updating, and the majority of the subjects in our information experiment are classified as non-Bayesian updaters. To provide some measure of the consequences of naively assuming all individuals update in a Bayesian fashion, we conduct an exercise in which we compute the gap between the expected earnings using observed revisions in our sample and the Bayesian-based expected earnings using the Bayesian benchmark. This gap, which we refer to as the "Bayesian welfare shortfall" is defined as:

$$\Delta \text{Welfare}_i^{Bayes} \equiv \sum_m \text{prob}_{im}^{post} * (\text{earnings}_{im}^{Bayes post} - \text{earnings}_{im}^{post}), \quad (5)$$

where $\text{earnings}_{im}^{Bayes post}$ is obtained from the updating model in equation (2). $\Delta \text{Welfare}_i^{Bayes} > 0$ implies that the Bayesian updating rule yields higher expected earnings than the actual update we observe the individual making. $\Delta \text{Welfare}_i^{Bayes} < 0$ implies that Bayesian rule is sub-optimal relative to the actual belief updating.

The second column of Table 12 calculates various statistics for the distribution of Bayesian welfare shortfall. We find that the mean Bayesian shortfall is substantial, with the average loss in expected earnings at age 30 of \$13,860. Reflecting the heterogeneity in updating heuristics we previously identified, more than a third of respondents would have received a positive gain in expected earnings from the assumption of Bayesian updating. However, the majority of respondents would have experienced a loss from the assumption of Bayesian updating. This

suggests that allowing for heterogeneous non-Bayesian updating, rather than naively imposing Bayesian updating, is an important modeling consideration with substantial differences in the implied welfare levels.

6 Conclusion

Expectations and aspirations have been shown to be important predictors of schooling choices, above and beyond other standard determinants of schooling (Jacob and Wilder, 2010). How students form these expectations is an important question for researchers and policy-makers alike, and remains an understudied area. This paper attempts to fill this gap by using an information experiment embedded in a survey. We find that students revise their beliefs of future earnings when provided with information on the population distribution of these characteristics. While there is substantial heterogeneity in students' response to information, it is correlated with the information content of the signals they receive, suggesting sensible updating on part of students. We also find substantial heterogeneity in updating heuristics used by our sample, with the majority of students classified as non-Bayesian updaters.

One policy implication of our results is almost immediate: Students respond to information about the population distribution of earnings by revising their beliefs as well as expected future choices. Since expectations play a critical role in decision-making under uncertainty and, in particular, for human capital decisions which have substantial economic consequences (Cunha et al., 2005), the large errors in population beliefs in our sample – even one comprised primarily of high ability students – suggests a role for information campaigns focused on providing accurate information on returns to schooling. While such campaigns have been conducted in developing countries (Jensen, 2010; Nguyen, 2010), our results make a case for such interventions in developed countries as well.

While there are large gender differences in composition of college majors (Zafar, 2011; Gemici and Wiswall, 2011), we do not find gender differences in information processing. Studies have shown that men tend to be more overconfident than women in a wide variety of settings (Barber and Odean, 2001; Niederle and Vesterlund, 2007). Possible mechanisms through which this may happen are gender differences in information acquisition and/or information processing. In our experimental setup, students don't have a choice to acquire information – they are simply given some information. In real instances, people choose when to acquire information based on the expected (perceived) costs and benefits of the information acquisition (e.g., whether to speak with a career counselor about earnings in different fields). The selective information acquisition process could result in different expectations updating, even if there are no differences in information-processing. In our study, we cannot address gender differences in information

acquisition. However, our findings rule out gender differences in information processing as a possible explanation. This is at odds with Mobius et al. (2011) who find substantial gender differences in both information processing and information acquisition. Possible explanations for these different findings could be that students in our study estimate absolute earnings, not relative performance as in their study, and that the two study designs have very different setups and information structures.

Another notable finding is that response to information is asymmetric and that, when information is bad news, students are likely to discount it. These findings support recent theoretical work on economic decisions involving uncertainty and belief formation over quantities of importance to the individual, such as future earnings. In these models, beliefs affect utility directly and not only through their impact on decision making. These models of ego or anticipatory utility predict that information processing would deviate from Bayesian updating towards optimism (Brunnermeier and Parker, 2005; Koszegi, 2006). Our findings are in line with this bias, and have implications for field studies and other interventions in which information or feedback is disseminated to respondents, particularly in the context of human capital investment decisions.

Finally, how students revise their beliefs and choices in a framework like ours where information is presented to them may be very different from the change in their actions if they were to acquire the information themselves (Hertwig et al., 2004). While it is challenging to identify changes in information sets in actual panels because of various confounding factors, an important question for future research is to explore how students' beliefs and choices evolve over longer time horizons, and in settings where they self select information.

References

- [1] Altonji, Joseph. 1993. "The Demand for and Return to Education when Education Outcomes are Uncertain." *Journal of Labor Economics*, 11: 48-83.
- [2] Arcidiacono, Peter. 2004. "Ability Sorting and the Returns to College Major." *Journal of Econometrics*, 121(1-2): 343-375.
- [3] Arcidiacono, Peter, Joseph Hotz, and Songman Kang. 2011. "Modeling College Major Choices using Elicited Measures of Expectations and Counterfactuals." *Journal of Econometrics*, forthcoming.
- [4] Armantier, Olivier, and Nicolas Treich. 2011. "Eliciting Beliefs: Proper Scoring Rules, Incentives, Stakes and Hedging." Working Paper.
- [5] Attanasio, Orazio, and Katja Kaufmann. 2011. "Education Choices and Returns on the Labour and Marriage Markets: Subjective Expectations, Gender and Decision Making." Working Paper.

- [6] Barber, Brad, and Terrance Odean, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." *The Quarterly Journal of Economics*, 116 (1): 261–292.
- [7] Beffy, Magali, Denis Fougere, and Arnaud Maurel. 2011. "Choosing the Field of Study in Post-Secondary Education: Do Expected Earnings Matter?" *The Review of Economics and Statistics*, Forthcoming.
- [8] Bettinger, Eric, Bridget Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2011. "The Role of Simplification and Information in College Decisions: Results from the H&R Block FAFSA Experiment" Working Paper.
- [9] Betts, Julian. 1996. "What do Students Know about Wages? Evidence from a Survey of Undergraduates." *Journal of Human Resources*, 31(1): 27-56.
- [10] Blau, Francine, and Marianne Ferber. 1991. "Career Plans and Expectations of Young Women and Men." *Journal of Human Resources*, 26: 581-607.
- [11] Brunnermeier, Markus, and Jonathan Parker. 2005. "Optimal Expectations." *American Economic Review*, 95(4): 1092–1118.
- [12] Cameron, Trudy. 2005. "Updating Subjective Risks in the Presence of Conflicting Information: An Application to Climate Change." *Journal of Risk and Uncertainty*, 30: 63-97.
- [13] Cunha, Flavio, James Heckman, Lance Lochner and Dimitriy Masterov. 2005. "Interpreting the evidence on life cycle skill formation." In *Handbook of the Economics of Education*, eds. E. Hanushek and F. Welch.
- [14] Dinkelman, Taryn, and Claudia Martínez. 2011. "Investing in Schooling in Chile: The Role of Information about Financial Aid for Higher Education." Working Paper.
- [15] Delavande, Adeline. 2008. "Measuring Revisions to Subjective Expectations." *Journal of Risk and Uncertainty*, 36(1): 43-82.
- [16] Dominitz, Jeff and Charles Manski. 1996. "Eliciting Student Expectations of the Returns to Schooling." *Journal of Human Resources*, 31: 1-26.
- [17] Dominitz Jeff. 1998. "Earning Expectations, Revisions, and Realizations." *Review of Economics and Statistics*, 80: 374-388.
- [18] Eil, David, and Justin Rao. 2011. "The Good News-Bad News Effect: Asymmetric Processing of Objective Information About Yourself." *American Economic Journal: Microeconomics*, 3(2): 114–138.
- [19] El-Gamal, Mahmoud, and David Grether. 1995. "Are People Bayesian? Uncovering Behavioral Strategies." *Journal of the American Statistical Association*, 90 (432): 1137-1145.
- [20] Gemici, Ahu, and Matthew Wiswall. 2011. "Evolution of Gender Differences in Post-Secondary Human Capital Investments: College Majors at the Intensive Margin." Working Paper.

- [21] Giustinelli, Pamela. 2010. "Uncertain Outcomes and Child-Parent Decision Making in Curriculum Choice: What Data Do We Need to Tell Them Apart?" Working Paper.
- [22] Grether, David. 1980. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." *The Quarterly Journal of Economics*, 95 (3): 537–57.
- [23] Grossman, Zachary, and David Owens. 2011. "An Unlucky Feeling: Overconfidence and Noisy Feedback." Working Paper.
- [24] Hastings, Justine, and Jeffrey Weinstein. 2008. "Information, School Choice and Academic Achievement: Evidence from Two Experiments." *Quarterly Journal of Economics*, 123(4): 1373-1414.
- [25] Hertwig, Ralph, Greg Barron, Elke Weber, and Ido Erev. 2004. "Decisions from Experience and the Weighting of Rare Events." *Psychological Science*, 15 (8): 534-539.
- [26] Jacob, Brian, and Tamara Wilder. 2010. "Educational Expectation and Attainment." NBER Working Paper 15683.
- [27] Jensen, Robert. 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics*, 125(2): 515-548.
- [28] Kahneman, Daniel, and Amos Tversky. 1982. "Subjective probability: A judgement of representativeness." in Daniel Kahneman, Paul Slovic and Amos Tversky (eds) *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- [29] Karni, Edi, and Zvi Safra. 1995. "The Impossibility of Experimental Elicitation of Subjective Probabilities." *Theory and Decision*, 38: 313-320.
- [30] Kaufmann, Katja. 2010. "Understanding the Income Gradient in College Attendance in Mexico: The Role of Heterogeneity in Expected Returns to College." Working Paper.
- [31] Koszegi, Botond. 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association*, 4(4): 673–707.
- [32] Manski, Charles. 1989. "Schooling as Experimentation: A Reappraisal of the Post-Secondary Drop-out Phenomenon." *Economics of Education Review*, 8(4): 305-312.
- [33] Manski, Charles. 2004. "Measuring Expectations." *Econometrica*, 72(5): 1329-1376.
- [34] Mobius, Markus, Muriel Niederle, Paul Niehaus, and Tanya Rosenblat (2011). "Managing Self-Confidence: Theory and Experimental Evidence." Working Paper.
- [35] Montmarquette, Claude, Kathy Cannings, and Sophie Mahseredjian. 2002 "How Do Young People Choose College Majors?" *Economics of Education Review*, 21(6): 543-556.
- [36] Nguyen, Trang. 2010. "Information, Role Models and Perceived Returns to Education: Experimental Evidence from Madagascar," MIT Working Paper.
- [37] Niederle, Muriel, and Lise Vesterlund. 2007. "Do Women Shy away from Competition? Do Men Compete too Much?" *Quarterly Journal of Economics*, 122(3): 1067-1101.

- [38] Smith, Herbert, and Brian Powell. 1990. Great Expectations: Variations in Income Expectations Among College Seniors. *Sociology of Education*, 63 (July): 194-207.
- [39] Stinebrickner, Todd, and Ralph Stinebrickner. 2010. "Learning about Academic Ability and the College Drop-out Decision." University of Western Ontario Working Paper 20086.
- [40] Stinebrickner, Todd, and Ralph Stinebrickner. 2011. "Math or Science? Using Longitudinal Expectations Data to Examine the Process of Choosing a College Major." Working Paper.
- [41] Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science*, 185(4157): 1124-1131.
- [42] Winkler, Robert, and Allan Murphy. 1970. "Nonlinear Utility and the Probability Score." *Journal of Applied Meteorology*, 9: 143-148.
- [43] Wiswall, Matthew. 2006. "Skill Composition and the Returns to Post-Secondary Education in the United States." Working Paper.
- [44] Wiswall, Matthew, and Basit Zafar. 2011. "Determinants of College Major Choice: Identification Using an Information Experiment." Staff Report No. 500. Federal Reserve Bank of New York.
- [45] Viscusi, W. Kip, and Charles O'Connor. 1984. Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers? *American Economic Review*, 74(5): 942-956.
- [46] Viscusi, W. Kip. 1997. Alarmist Decisions with Divergent Risk Information. *The Economic Journal*, 107(445): 1657-1670.
- [47] Zafar, Basit. 2010. "College Major Choice and the Gender Gap." Staff Report No. 364, Federal Reserve Bank of New York.
- [48] Zafar, Basit. 2011. "How do College Students Form Expectations." *Journal of Labor Economics*, 29(2): 301-348.

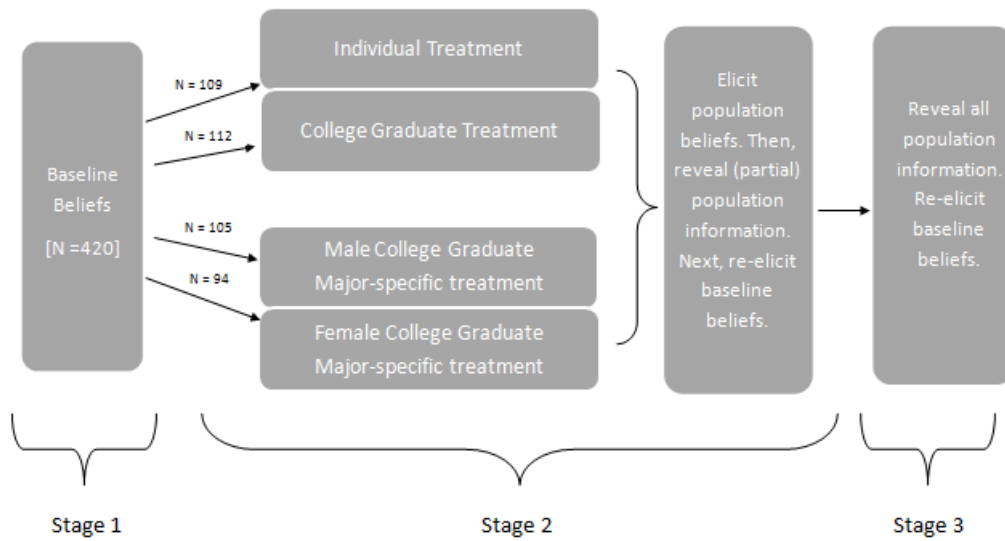


Figure 1: Survey Outline

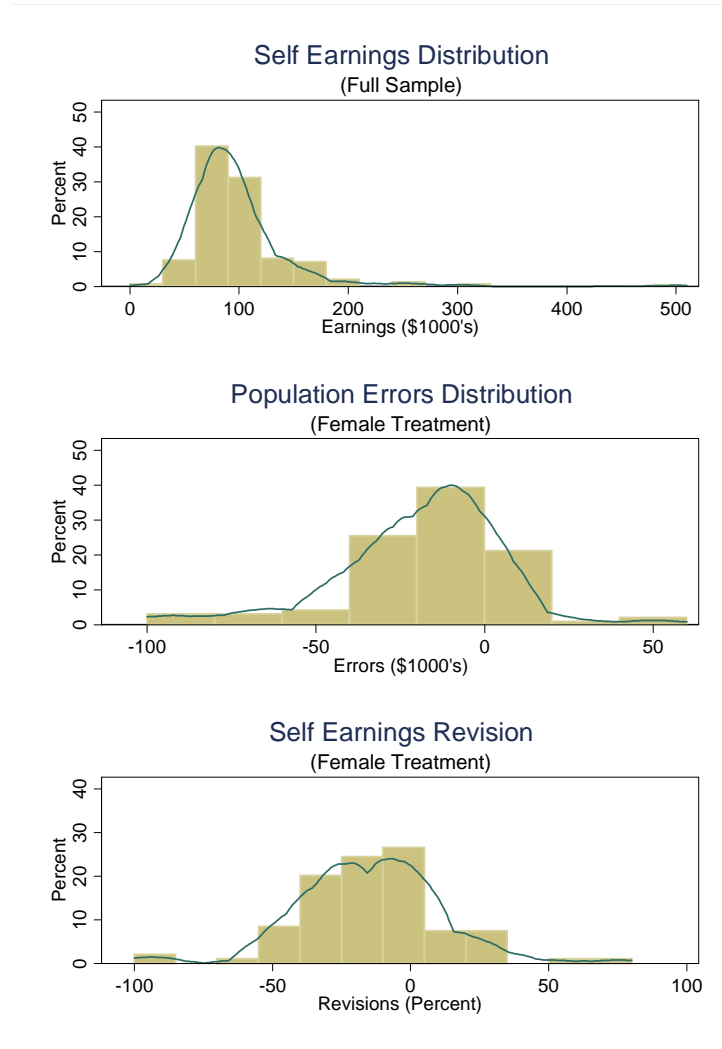


Figure 2: Beliefs about earnings in Economics/Business (in 000s of dollars). Top panel shows the self beliefs about earnings in econ/business at the initial stage for all respondents. Middle panel shows the errors in population beliefs (true female earnings in econ/business - population beliefs about female graduates in econ/business) for respondents in the Female Treatment. The bottom panel shows the percent revision of self beliefs of earnings in econ/business (intermediate-initial self beliefs) for respondents in the Female treatment.

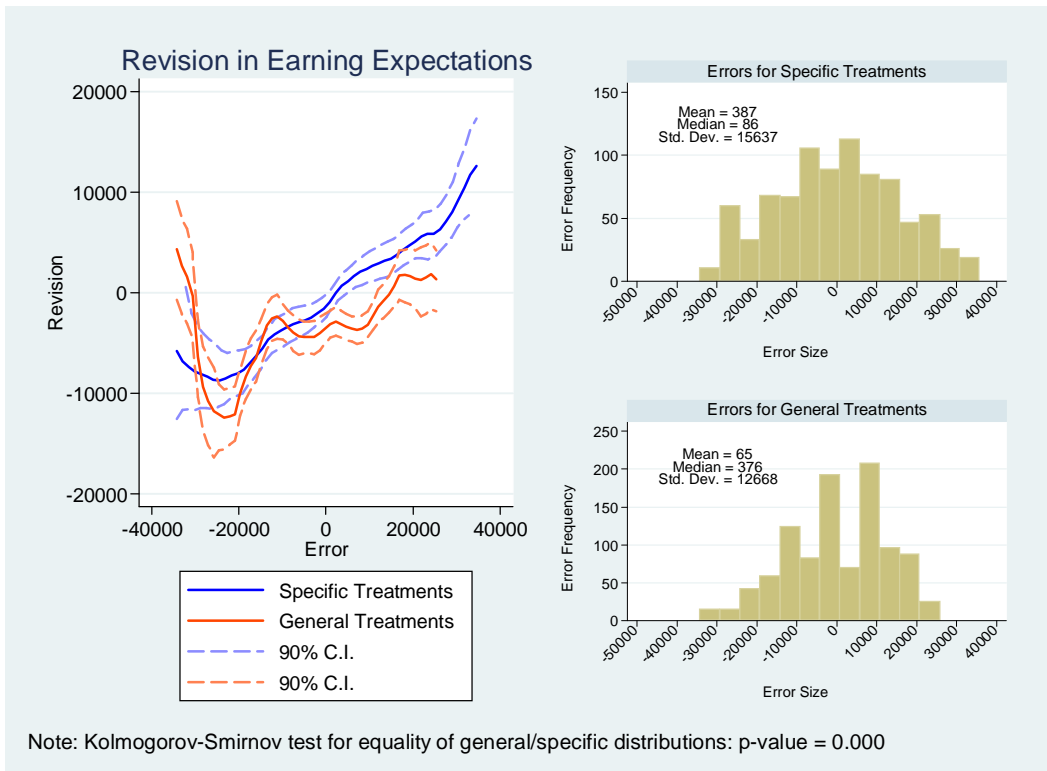


Figure 3: Local linear regression of self earnings revisions on population errors, for General and Specific treatments.

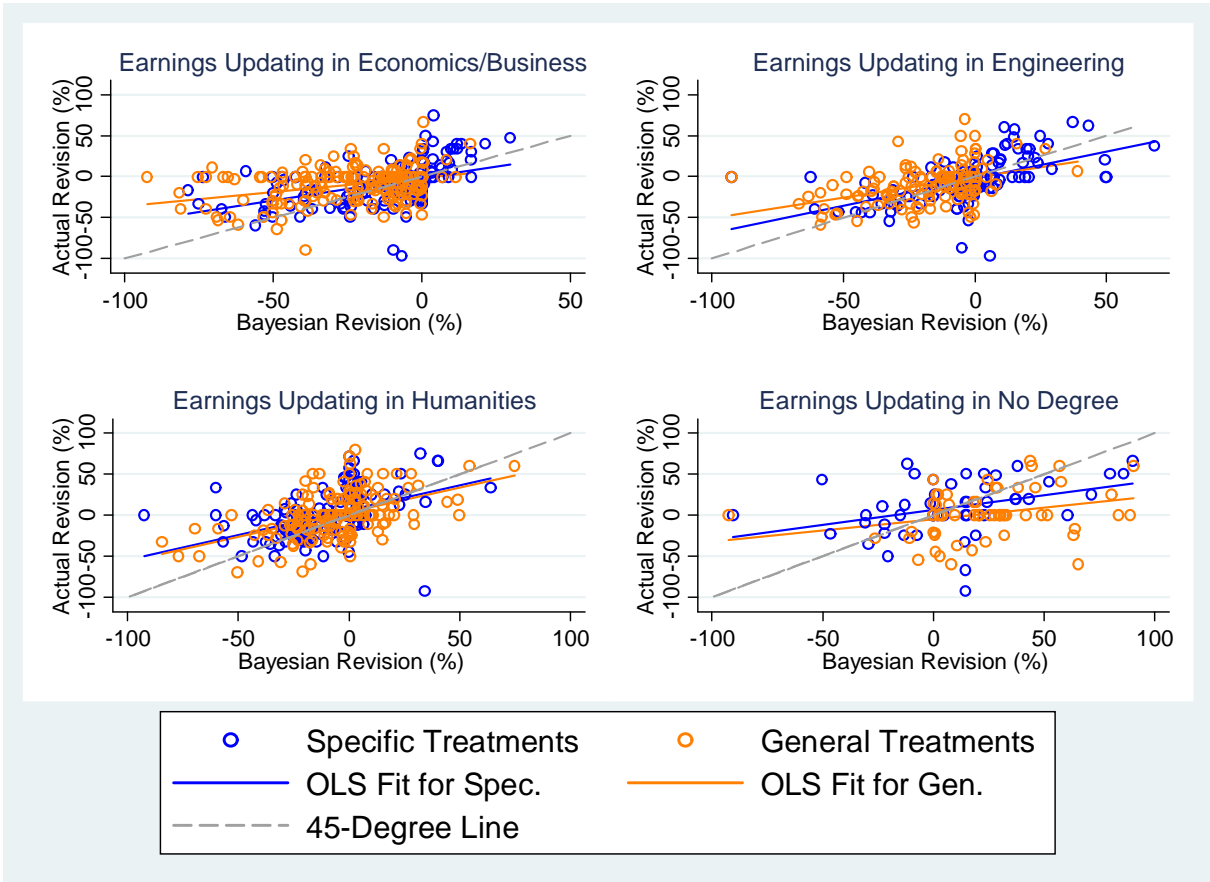


Figure 4: Actual revision $\left(\frac{\text{observed posterior}-\text{observed prior}}{\text{observed prior}}\right)$ versus Bayesian revision $\left(\frac{\text{Bayesian posterior}-\text{observed prior}}{\text{observed prior}}\right)$ in 4 different major categories. Also shown are the lines of best fit for the Specific and General treatments, and a 45-degree line).

Table 1: Information revealed in the Treatments

All Individuals Treatment

The following information is from the US Census Bureau.

Among all individuals (including college and non-college graduates) aged 30:

The percentage that are working full time is	59.80%
The percentage of those that are working full time who are women is	42.70%
The average annual earnings of those that are working full time is	\$45,726
The percentage of those that are working full time that earn more than \$35,000 per year is	59.00%
The percentage of those that are working full time that earn more than \$85,000 per year is	7.30%

College Treatment

The following information is from the US Census Bureau.

Among all college graduates currently aged 30:

The percentage that are working full time is	69.80%
The percentage of those that are working full time who are women is	52.80%
The average annual earnings of those that are working full time is	\$60,376
The percentage of those that are working full time that earn more than \$35,000 per year is	80.70%
The percentage of those that are working full time that earn more than \$85,000 per year is	14.80%

Female Major Specific Treatment

The following information is from the US Census Bureau.

Among all female college graduates aged 30 who received a Bachelor's degree in major (M):

	Econ	Eng	Hum	Nat	No Grad
The percentage that are working full time is	60.6%	72.8%	52.3%	55.3%	51.6%
The average annual earnings of those that are working full time is	\$60,730	\$75,086	\$49,154	\$60,021	\$34,603
The percentage of those that are working full time that earn more than \$35,000 per year is	85.5%	99.0%	72.2%	84.0%	44.9%
The percentage of those that are working full time that earn more than \$85,000 per year is	27.5%	26.9%	8.0%	8.5%	1.6%

Male Major Specific Treatment

The following information is from the US Census Bureau.

Among all male college graduates aged 30 who received a Bachelor's degree in major (M):

	Econ	Eng	Hum	Nat	No Grad
The percentage that are working full time is	93.5%	91.6%	77.6%	81.9%	72.1%
The average annual earnings of those that are working full time is	\$74,542	\$82,377	\$52,937	\$72,583	\$47,803
The percentage of those that are working full time that earn more than \$35,000 per year is	92.4%	95.2%	78.8%	90.6%	65.2%
The percentage of those that are working full time that earn more than \$85,000 per year is	31.5%	33.6%	8.7%	24.2%	5.7%

Also Revealed to All Respondents in Final Stage

The percentage of those who are women is	34.70%	18.20%	55.20%	48.00%	42.30%
--	--------	--------	--------	--------	--------

Table 2: Sample Characteristics

Number of respondents:		420
Num of respondents by Treatment:		
Male Treatment		105
Female Treatment		94
College Treatment		112
Individuals Treatment		109
School year:		
Freshman		40.48%
Sophomore		36.19%
Junior		23.33%
Mean Age		20.15
	(std.)	(1.16)
Female		63.33%
Race:		
White		38.57%
Non-Asian Minority		15.95%
Asian		45.48%
Parents' Characteristics:		
Mean Parents' Income		151.04
	(std.)	(152.56)
Mother B.A. or More		71.22%
Father B.A. or More		75.30%
Ability Measures:		
Mean SAT Math Score		701.03
	(std.)	(77.37)
Mean SAT Verbal Score		684.51
	(std.)	(70.75)
Mean GPA		3.48
	(std.)	(0.32)
Intended/Current Major:		
Economics		30.24%
Engineering		5.00%
Humanities		47.85%
Natural Sciences		16.90%
(Intend to) Double Major		36.84%

Table 3: Baseline Beliefs about Self, and Population Earnings (in 000s of Dollars)

	Self Beliefs			Prob of earning ^c		Pop. Earning Beliefs: ^d			Absolute Pop. Earning Errors: ^g		
	All Beliefs ^a (1)	(in major) ^b (2)	All (3)	All (4)	Male T ^e (5)	Female T (6)	General T ^f (7)	Male T (8)	Fem T (9)	General T (10)	
Economics	100.03 [90]	111.31*** [95]	83.55 [95]	67.04 [70]	95.10 [80]	80.90 [75]		31.28 [15.46]	23.89 [19.27]		
Engineering	(70.81) 88.78 [80]	(77.45) 124.38 [95]	(26.29) 81.6 [93]	(22.59) 61.34 [65]	(58.82) 91.13 [75]	(28.10) 77.56 [70]		(53.84) 30.99 [17.38]	(24.98) 21.04 [14.91]		
Humanities	(64.24) 64.16 [60]	(188.93) 61.81 [60]	(26.37) 74.37 [80]	(23.33) 41.35 [40]	(93.51) 64.76 [60]	(46.51) 62.29 [55]		(88.61) 19.89 [12.94]	(41.49) 20.23 [10.85]		
Natural Science	(53.97) 82.26 [70]	(27.12) 119.99*** [85]	(24.04) 77.9 [85]	(22.36) 54 [50]	(36.50) 76.79 [70]	(64.10) 71.04 [65]		(32.77) 23.04 [17.42]	(62.21) 20.30 [10.02]		
No Degree	(80.02) 37.48 [30]	(152.09) [50]	(25.56) 48.71 [50]	(24.49) 15.25 [10]	(42.56) 38.30 [40]	(50.05) 36.98 [30]		(35.96) 14.94 [14.80]	(46.72) 15.01 [9.60]		
All Indiv. ^h			(28.21)	(17.87)	(14.85)	(45.21)		(9.30)	(42.68)		
College ⁱ											
Observations	420	<i>j</i>	420	420	105	94		105	94		

Mean reported in the first cell. Median reported in square brackets []. Standard Deviation reported in parentheses ().

^a Beliefs reported in the Initial Stage about self earnings in '000s of dollars.

^b Beliefs of earnings of respondents who report to be majoring in that major (in 000s). This column also reports the pairwise test of whether mean of these respondents is equal to that reported by those not majoring in that major: *** Sig. at the 1% level.

^c Probability (on a 0-100 scale) of annual income at age 30 being at least \$35,000, and at least \$85,000 in each of the major categories.

^d Beliefs reported about the earnings of current 30 year olds working in the labor force (in 000s of dollars).

^e Male T (Female t) column refers to the Male (Female) treatments. In these, respondents reported the population beliefs for male (female) workers.

^f General T refers to the two general treatments— Individual and College Treatments.

^g Absolute Population Earning Error in major $m = |\text{True Population Earnings in } m - \text{Beliefs about pop earnings in } m|$.

^h In the Individual Treatment, students reported population beliefs about all Individuals currently in the full-time labor force.

ⁱ In the College Treatment, students reported population beliefs about College graduates currently in the full-time labor force.

^j There are 140 students (intending to) majoring in Econ, 24 in Engineering, 227 in Humanities, and 76 in Natural Sciences.

^k 109 students reported beliefs in the Individual treatment; 112 reported beliefs in the College treatment.

Table 4: Heterogeneity in (Absolute) Errors

Dependent Variable: Absolute Population Error	General Treatments			Specific Treatments		
	All	Positive Error [⊗]	Negative Error	All	Positive Error	Negative Error
	(1)	(2)	(3)	(4)	(5)	(6)
Female	3201.8 (3490.8)	587.1 (774.1)	12958.5 (8277.3)	8166.8** (3202.7)	1073.3 (1162.4)	13328.8** (6281.3)
High Ability ^a	-12067.2*** (3752.7)	-2541.3*** (817.4)	-17581.3* (9346.6)	8644.5** (3405.5)	-2229.5* (1212.6)	19773.5*** (6444.8)
Sophomore	-11389.3*** (3743.6)	1592.5* (896.4)	-22895.5*** (8026.4)	5540.5 (3619.3)	2244.1* (1311.4)	8581.5 (6747.4)
Junior	-12307.0*** (4388.4)	-1081.0 (945.5)	-12988.1 (11536.9)	5259.8 (3981.4)	-351.5 (1464.3)	8215.5 (7434.1)
Non-Asian minority	-1156.2 (4854.0)	255.1 (1060.4)	5343.3 (11545.2)	-7305.1 (4732.5)	2296.3 (1870.3)	-16222.5* (8441.3)
Asian	9767.9*** (3640.0)	-64.29 (832.6)	25199.2*** (8170.0)	-4959.4 (3388.3)	235.6 (1233.5)	-12350.1* (6482.5)
Gender Matches ^b				-2962.6 (3110.9)	-606.4 (1150.2)	-5151.1 (5809.6)
In-Major ^c				6556.3 (4233.9)	42.42 (1868.4)	8225.1 (6935.2)
College Treatment	9338.2*** (3294.6)	-593.5 (743.8)	27090.1*** (7497.8)			
Economics				8772.3* (4904.8)	1126.0 (2287.4)	10494.8 (7782.8)
Engineering				8964.9* (5146.4)	4255.9** (2054.9)	21663.8** (9707.4)
Natural Science				3978.9 (5026.0)	2153.3 (2095.2)	8277.0 (8719.8)
No Degree				-2100.4 (5245.4)	1943.8 (2091.2)	-4790.5 (10086.1)
Constant	15296.0*** (4498.4)	10676.4*** (1025.7)	4410.1 (10568.8)	10733.2** (5362.5)	11506.3*** (2280.8)	11286.2 (9044.5)
Obs.	1080	610	470	990	472	518
R-Squared	0.037	0.035	0.079	0.031	0.035	0.057

Table reports pooled OLS estimates of the absolute error on demographics.

Absolute Error in major $m = |\text{True Population Earnings in } m - \text{Beliefs about pop earnings}|$.

Standard Deviations in Parentheses. *, **, *** represent significance at 10%, 5%, and 1%, respectively.

[⊗] Sample restricted to observations with positive population error (underestimation of population earnings)

^a High Ability is defined as SAT score > 1450; 123 of the 420 respondents are high ability.

^b Dummy that equals 1 if the respondent's gender is the same as that of the population workers about whom information is provided.

^c Dummy that equals 1 if the respondent's (intended) major is the same as the one for which beliefs are being reported.

Table 5: Percent Revisions in Self Earnings

	All (1)	Intermediate – Specific T (2)	Initial ^a General T (3)	All (4)	Final – Specific T (5)	Initial ^b General T (6)
Economics	-5.11 (6.95)	-1.88 (13.48)	-8.02 (5.10)	-8.04 (8.84)	-2.71 (17.06)	-12.83* (6.69)
Engineering	0.78 (6.95)	6.52 (13.48)	-4.39 (5.10)	1.99 (8.84)	9.88 (17.06)	-5.11** (6.69)
Humanities	2.19 (6.95)	0.59 (13.48)	3.64 (5.10)	5.65 (8.84)	5.50 (17.06)	5.78 (6.69)
Natural Science	-0.28 (6.95)	0.97 (13.48)	-1.40 (5.10)	3.15 (8.84)	5.17 (17.06)	1.33 (6.69)
No Degree	54.57 (6.95)	84.76 (13.48)	27.38* (5.10)	85.81 (8.84)	122.08 (17.06)	53.15* (6.69)
Num Obs.	2100	995	1105	2100	995	1105

The table reports the mean percent revisions for self beliefs for the various major categories. Standard errors in parentheses.

^a % Revision in self earnings from initial to intermediate stage: $\frac{\text{intermediate self belief} - \text{initial self belief}}{\text{initial self belief}} * 100$

^b Revision in earnings from initial to final stage, after all information in the four treatments has been revealed to students: $\frac{\text{final self belief} - \text{initial self belief}}{\text{initial self belief}} * 100$

*** Equality of % revision in self beliefs is rejected in the Specific and General treatments (using a 2-tailed t-test).

Table 6: Self Earnings Updating and Population Errors

Dependent Variable: Revisions in Self Earnings Beliefs (Intermediate – Initial)						
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Error ^a	0.034***					
	(0.009)					
Error × General T		-0.000011				
		(0.012)				
Error × Specific T		0.078***				
		(0.014)				
Error × 1(Error>0)			0.181***			
			(0.063)			
Error × 1(Error≤0)			0.0035			
			(0.0096)			
Panel B						
Error × Female				0.029***		
				(0.0094)		
Error × Male				0.092***		
				(0.030)		
Error × Freshman					0.024**	
					(0.012)	
Error × Sophomore					0.053***	
					(0.018)	
Error × Junior					0.047**	
					(0.022)	
Error × High Ability ^b						0.051***
						(0.018)
Error × Low Ability						0.027**
						(0.011)
Num. Obs	2100	2100	2100	2100	2100	2070

Table reports OLS estimates of regression of (intermediate-initial) revision of self beliefs on population errors by information type and individual characteristics. All regressions include a constant term and dummies for each of the covariates that are interacted with Error (not reported here).

Standard Deviations in Parentheses. *, **, *** represent significance at 10%, 5%, and 1%, respectively.

^a Error = Population Earnings Belief - True Population Earnings

^b High Ability is defined as SAT score > 1450; 123 of the 420 respondents are high ability.

Table 7: Distribution of Updating Heuristics For Self Earnings in Own (Intended) Major

	All	Male	Female	Freshmen	Upper classmen	Low Ability ^a	High Ability	Positive Error ^b	Negative Error
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Specific Treatments								
Bayesian	23.1%	20.5%	24.8%	19.2%	31.4%	23%	23.8%	23.8%	22.8%
Alarmist	25.1%	23.1%	26.5%	32.1%	13.7%	21.5%	33.3%	28.6%	23.5%
Conservative	13.6%	12.8%	14.1%	19.2%	9.8%	16.3%	7.9%	1.6%	19.1%
Non-Updater	24.1%	26.9%	22.3%	18.0%	25.5%	24.4%	23.8%	25.4%	23.5%
Confused	14.1%	16.7%	12.4%	11.5%	19.6%	14.8%	11.1%	20.6%	11%
Obs.	199	78	121	78	51	135	63	63	136
	General Treatments								
Bayesian	18.1%	18.4%	17.9%	14.1%	23.4%	22.4%	8.3%	20.2%	15.5%
Alarmist	24.9%	29%	22.8%	28.3%	19.2%	23.7%	28.3%	24.2%	25.8%
Conservative	15.8%	21.1%	13.1%	17.4%	14.9%	14.7%	20%	12.1%	20.6%
Non-Updater	22.2%	19.7%	23.5%	20.7%	29.8%	19.9%	26.7%	25.0%	18.6%
Confused	19.0%	11.8%	22.8%	19.6%	12.8%	19.2%	16.7%	18.6%	19.6%
Obs.	221	76	145	92	47	156	60	124	97

The table reports the distribution (percent) of types, defined by updating heuristic.

See text for definition of each type.

^a Low ability is the subsample of respondents with SAT score ≤ 1450

^b Positive error is the subset of respondents who underestimated population earnings.

Table 8: Updating Type for Own (intended) Major and Other Majors

Type in own major:	Type in other majors:					Row Total ^b	
	Bayesian	Alarmist	Conservative	Non-Updater	Confused		Mixed
Bayesian	6.53%	2.51%	0%	0.5%	1.51%	12.06%	23.12% (46)
Alarmist	3.02%	11.56%	0.5%	0%	2.01%	8.04%	25.13% (50)
Conservative	1.01%	1.01%	3.02%	0.5%	3.02%	5.03%	13.57% (27)
Non-Updater	2.01%	3.52%	0.5%	3.02%	3.02%	12.06%	24.12% (48)
Confused	0.5%	2.51%	1.01%	0.5%	3.02%	6.53%	14.07% (28)
<i>Column Total^a</i>	<i>13.07% (26)</i>	<i>21.11% (42)</i>	<i>5.03% (10)</i>	<i>4.52% (9)</i>	<i>12.56% (25)</i>	<i>43.72% (87)</i>	<i>100% (199)</i>
			<i>Specific Treatments</i>				
Bayesian	2.71%	0.9%	2.71%	1.81%	1.81%	8.14%	18.1% (40)
Alarmist	1.36%	5.88%	1.36%	0.45%	4.52%	11.31%	24.89% (55)
Conservative	2.26%	0.9%	4.98%	1.81%	1.81%	4.07%	15.84% (35)
Non-Updater	1.81%	0.45%	3.62%	6.79%	2.71%	6.79%	22.17% (49)
Confused	0.45%	1.36%	2.26%	2.71%	6.79%	5.43%	19% (42)
<i>Column Total</i>	<i>8.6% (19)</i>	<i>9.5% (21)</i>	<i>14.93% (33)</i>	<i>13.57% (30)</i>	<i>17.65% (39)</i>	<i>35.75% (79)</i>	<i>100% (221)</i>

Each cell shows the percent of the sample that falls in that category. Number of observations in parentheses.

See text for how types are defined for both own major, and non-pursued majors.

^a Column Total shows the % of respondents who are classified as the column Type for updating earnings in their other majors.

^b Row Total shows the % of respondents who are classified as the row Type for updating earnings in their own (intended) major.

Table 9: Absolute Error (in each Major Category) vs. Type (in each Major Category)

Dependent Variable: Absolute Error in Population Earnings ^a						
	All Majors	Economics	Engineering	Humanities	Nat. Science	No Degree
	(1)	(2)	(3)	(4)	(5)	(6)
	<i>All Treatments</i>					
Alarmist	3108.6 (3284.9)	10592.9 (6976.1)	-10393.4 (8377.1)	-3060.6 (7636.8)	7718.3 (6726.3)	11258.9 (7389.5)
Conservative	2167.9 (4068.8)	8167.3 (8190.9)	-9901.6 (11128.8)	3355.1 (9533.4)	3636.4 (8875.9)	3753.7 (7829.8)
Non-Updater	7087.0** (3368.4)	13333.1* (7393.6)	-1545.0 (9483.9)	1470.3 (7622.1)	13333.4* (7214.0)	10108.6* (6073.9)
Confused	1210.5 (3681.0)	1057.0 (8660.0)	-1065.5 (10080.5)	4083.1 (8366.4)	178.7 (7474.1)	2249.2 (6682.1)
Constant	16124.0*** (2409.3)	14137.0*** (5386.3)	25703.7*** (6336.0)	17308.1*** (5799.5)	13576.3*** (5030.1)	10293.3** (4343.2)
F-test (p-value) ^b	0.219	0.174	0.437	0.823	0.283	0.184
Observations	2084	420	419	420	420	405

Table reports OLS estimates of regression of Abs. error in pop earnings onto the respondent's type in that major (excluded category is Bayesian).

Standard errors in parentheses. All regressors are dummy variables. Significance stars (*, **, ***) represent significance at the 10%, 5%, and 1% levels, respectively.

^a Absolute Population Earning Error in major $m = |\text{True Population Earnings in } m - \text{Beliefs about pop earnings in } m|$.

^b P-value for a test of the joint significance of all the covariates excluding the constant term.

Table 10: Response to Effective Information by Updating Type

	Bayesian	Alarmist	Conservative	Non-Updater	Confused
All Majors	0.56 [0.41] (3.56)	0.98*** [0.54]*** (4.3)	0.48 [0.12] (4.41)	0*** [0]*** (0)	-0.42*** [-0.47]*** (3.67)
<i>Num. Obs.</i>	463	539***	250	485***	347***
Economics	0.42 [-0.18] (3.94)	0.97** [0.58]** (4.32)	0.8 [-0.38] (5.17)	0*** [0]** (0)	-0.35** [-0.09] (3.64)
<i>Num. Obs.</i>	84	124**	64*	95***	53*
Engineering	0.27 [0.24] (3.92)	0.96*** [0.48]* (4.42)	0.43 [-0.07] (5.16)	0*** [0]* (0)	-0.23** [0.19] (3.94)
<i>Num. Obs.</i>	98	131***	47	79***	64**
Humanities	0.55 [0.68] (2.83)	1.09* [0.97] (4.31)	0.48 [0.6] (3.69)	0*** [0]*** (0)	-0.61*** [-0.91]*** (3.53)
<i>Num. Obs.</i>	80	109***	47	110***	74***
Natural Science	0.6 [0.41] (3.42)	1** [0.49] (4.23)	0.39 [-0.34] (4.67)	0*** [0]*** (0)	-0.52*** [-0.44]*** (4.03)
<i>Num. Obs.</i>	93	118**	44	88***	77***
No Degree	0.78 [0.8] (3.51)	0.82 [-0.08] (4.13)	0.43* [0.95] (2.48)	0*** [0]*** (0)	-0.39*** [-0.9]*** (3.2)
<i>Num. Obs.</i>	108	57**	48***	113***	79***

The table reports the median, mean, and standard dev of the response to effective info ($\frac{\text{Posterior-Prior}}{\text{Effective Info}}$) for each major by the updating type in that major. Mean in square brackets and std dev in parentheses. The table also reports pairwise tests of the equality of the median (Median test), the mean (Wilcoxon rank-sum test), and the distribution (Kolmogorov-Smirnov test) against the corresponding value for the Bayesian type. Stars reported on the median, mean, and sample size, respectively. ***, **, * Difference significant at the 1%, 5%, and 10% level, respectively.

Table 11: Impact of Information on Choices and Welfare

	Absolute Probability Change ^a				
	Economics	Engineering	Humanities	Natural Science	No Degree
Int. - Initial	5.16 [0] 47.38% (8.96)	3.65 [0] 45.24% (6.9)	7.06 [2.5] 54.76% (10.53)	4.58 [0] 47.86% (8.62)	1.61 [0] 28.10% (4.42)
Final - Initial	5.68 [0.5] 50% (8.96)	4.1 [0] 46.19% (6.9)	7.28 [3] 54.76% (10.53)	4.85 [0] 46.67% (8.62)	1.67 [0] 26.90% (4.42)

^a The first row shows the mean absolute change in choice probability. In the second row, [.] is the median absolute change in probability and the % is the proportion of respondents who change their probability in that stage relative to the initial stage. Standard Deviations of absolute change in probabilities reported in parentheses in third row.

Table 12: Impact of Information on Choices and Welfare

	Welfare Change	
	Observed ^a	Bayesian Shortfall ^b
Int. - Initial	0.331 [0] 74.76% (5.06)	-15.46 [-3.47] 41.43% (70.74)
Final - Initial	0.432 [0] 77.38% (7.75)	-13.86 [-2.99] 44.05% (80.75)

^a Observed welfare change for individual i is defined as:

$$\sum_m (prob_{im}^{posterior} * earnings_{im}^{posterior} - prob_{im}^{prior} * earnings_{im}^{posterior})$$

^b Bayesian welfare shortfall for individual i is defined as:

$$\sum_m prob_{im}^{posterior} * (earnings_{im}^{Bayesian\ posterior} - earnings_{im}^{posterior})$$

Posterior is the updated belief reported in intermediate stage for the top panel and the final stage for the lower panel. Prior is the belief reported in the initial stage. See text for discussion of how the Bayesian posterior is calculated.

Welfare amounts are in 000s of dollars. The first row reports the mean observed welfare change; the second row reports the median change in [.] and the proportion of respondents with non-negative welfare change. Standard dev of welfare change reported in parentheses in third row.

A Appendix

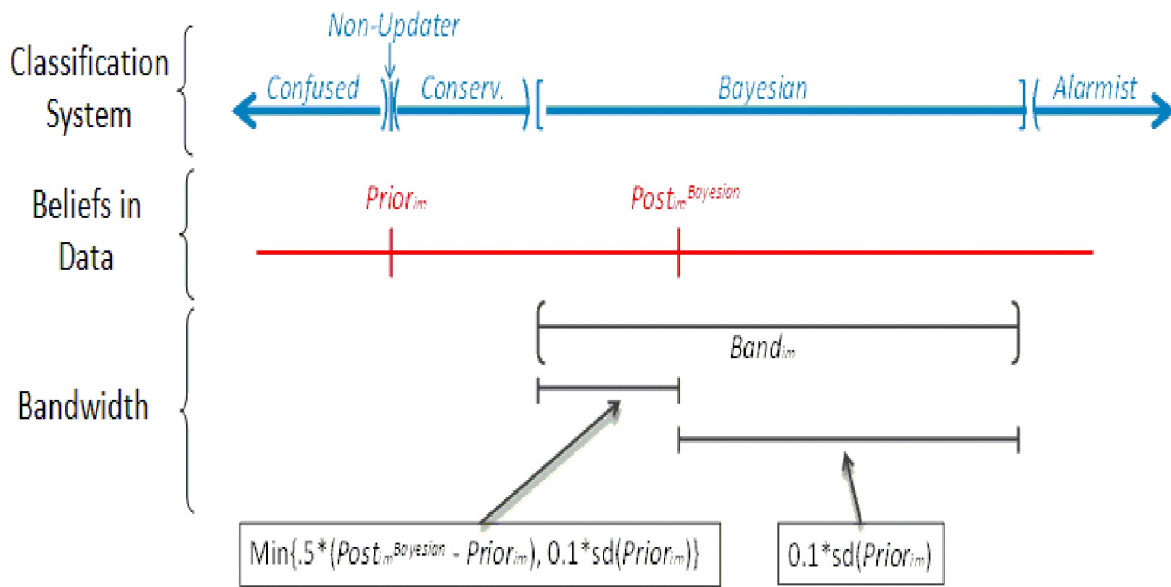


Figure A1: Classification of Heuristics. In this example, a Bayesian updater would revise upward on receipt of information.

Table A1: Asymmetric Response to Information

	(1)	(2)	(3)	(4)	(5)
Error ^a × General T × 1(Error>0)	0.258*** (0.0801)				
Error × General T × 1(Error<0)	-0.0111 (0.0124)				
Error × Specific T × 1(Error>0)	0.406*** (0.0649)				
Error × Specific T × 1(Error<0)	0.0482*** (0.0145)				
Error × Female × 1(Error>0)		0.319*** (0.0610)			
Error × Female × 1(Error<0)		0.0156 (0.00977)			
Error × Male × 1(Error>0)		0.474*** (0.0911)			
Error × Male × 1(Error<0)		0.000000373 (0.0363)			
Error × Freshman × 1(Error>0)			0.641*** (0.0905)		
Error × Freshman × 1(Error<0)			-0.0000848 (0.0122)		
Error × Sophomore × 1(Error>0)			0.197** (0.0775)		
Error × Sophomore × 1(Error<0)			0.0392** (0.0188)		
Error × Junior × 1(Error>0)			0.329*** (0.0964)		
Error × Junior × 1(Error<0)			0.0189 (0.0239)		
Error × High Ability × 1(Error>0)				0.720*** (0.109)	
Error × High Ability × 1(Error<0)				0.0127 (0.0185)	
Error × Low Ability × 1(Error>0)				0.263*** (0.0570)	
Error × Low Ability × 1(Error<0)				0.0118 (0.0109)	
Error × Economics × 1(Error>0)					0.480*** (0.119)
Error × Economics × 1(Error<0)					0.0371* (0.0215)
Error × Engineering × 1(Error>0)					0.541*** (0.0935)
Error × Engineering × 1(Error<0)					0.000764 (0.0170)
Error × Humanities × 1(Error>0)					0.122 (0.130)
Error × Humanities × 1(Error<0)					0.00481 (0.0205)
Error × Nat. sciences × 1(Error>0)					0.264** (0.105)
Error × Nat. sciences × 1(Error<0)					0.00949 (0.0219)
Error × No Degree × 1(Error>0)					0.206* (0.115)
Error × No Degree × 1(Error<0)					-0.00679 (0.0239)
Num. Obs	2100	2100	2100	2070	2100

Table reports OLS estimates of regression of (intermediate-initial) revision of self beliefs on covariates. All regressions include a constant term and dummies for each of the covariates that are interacted with Error. Standard Deviations in Parentheses. *, **, *** represent significance at 10%, 5%, and 1%, respectively.
^a Error = Population Earnings Belief - True Population Earnings

A.1 Information on Survey Design and Information Treatments

Description of data sources provide to survey respondents:

Sources:

1) CPS: The Current Population Survey (CPS) is a monthly survey of about 50,000 households conducted by the Bureau of the Census for the Bureau of Labor Statistics. The survey has been conducted for more than 50 years. The CPS is the primary source of information on the labor force characteristics of the U.S. population. The sample is scientifically selected to represent the civilian non-institutional population.

2) NSCG: The 2003 National Survey of College Graduates (NSCG) is a longitudinal survey, designed to provide data on the number and characteristics of individuals. The Bureau of the Census conducted the NSCG for the NSF (National Science Foundation). The target population of the 2003 survey consisted of all individuals who received a bachelor's degree or higher prior to April 1, 2000.

Methodology:

1) CPS: Our CPS sample is taken from the March 2009 survey. Full time status is defined as "usually" working at least 35 hours in the previous year, working at least 45 weeks in the previous year, and earning at least \$10,000 in the previous year. Average employment rates, average earnings, and percent with greater than \$35,000 or \$85,000 earnings is calculated using a sample of 2,739 30 year old respondents.

2) NSCG: We calculate inflation adjusted earnings using the Consumer Price Index. The salary figures we report are therefore equivalent to CPS figures in 2009 March real dollars. Full time status is defined as in the CPS sample. Given the need to make precise calculations for each field of study group, we use the combined sample of 30-35 year old respondents and age adjust the reported statistics for 30 year olds. This sample consists of 14,116 individuals. To calculate average earnings, we use an earnings regression allowing for separate age intercepts, one each for 6 ages 30-35. The predicted value of earnings from the regression is used as the estimate of average earnings for 30 year olds. For the percent full time employed, and percent with earnings greater than \$35,000 and \$85,000, we use a logit model to predict these percentages for 30 year olds and include a separate coefficient for each of the 6 ages 30-35.

A.2 Updating Heuristics in Other Majors

The classification system for the 4-major "Other majors" case is as follows. Using the method outlined in equations (2) and (3), we first identify the respondent's updating heuristic in each of the four majors. We then classify the updating heuristic in this "Other Majors" category as:

- CONFUSED if: the respondent uses the Confused heuristic in at least 2 of the 4 major

categories.

- NON-UPDATING if: the respondent does not update self earnings beliefs in at least 3 of the 4 major categories.
- ALARMIST if: the respondent is Alarmist in at least 3 of the 4 major categories, OR uses the Alarmist heuristic in 2 of the categories and Bayesian in the other two.
- BAYESIAN if: the respondent is Bayesian in at least 3 of the 4 major categories, OR uses the Bayesian heuristic in 2 of the four categories and the Conservative or Non-updating heuristic in the other two.
- CONSERVATIVE if: the respondent uses the Conservative or Non-Updating heuristic in at least 3 majors, but does not use the Non-Updating heuristic in more than 2 majors.
- MIXED if: the respondent uses at least one each from three of the following heuristics: Bayesian, Conservative, Confused, and Alarmist with at most 1 Confused; OR the respondent uses the Alarmist heuristic in 2 major categories and the Conservative heuristic in the other 2.