# Compliers, defiers and the evaluation of information interventions

Marco Castillo* Mofioluwasademi Ige† Manu Navjeevan‡§

October 18, 2026

**Abstract**

When beliefs are heterogeneous, information provision may induce responses in opposite directions, and standard intent-to-treat estimates may fail to capture the full effect of an information intervention. In this paper, we show how elicited prior beliefs can be used to test for a uniform response to information and bound the share of agents whose decisions change. We demonstrate that testing for a uniform response to information is equivalent to the problem of testing for instrument validity in LATE models, with elicited prior beliefs in place of observed outcomes. Under random assignment, the joint distribution of reported prior beliefs and choices yields sharp lower bounds on flows into and out of each choice. For a binary choice, these lower bounds are exact under a testable behavioral model in which reported priors coincide with underlying beliefs and the information intervention causes agents to update their beliefs towards a common signal. We apply these results to an information RCT in Ibadan, Nigeria, where adolescents chose whether to register for senior high school and which track to pursue. Among students whose earnings priors lie on the same side of every signal provided by the intervention, we reject a uniform response to information at the 10 percent level. In a rebalanced sample, we estimate that at least 14.6 percent of students changed continuation decisions upon receipt of new information — three times the 4.9 percentage-point decline in continuation. Accounting for changes in educational track, we estimate that at least 17 percent of students changed their schooling choice in response to the intervention.

JEL CLASSIFICATIONS: C93, J24, D83

KEYWORDS: human capital accumulation, costly information acquisition, randomized controlled trial

*Texas A&M University, CESIfo and IZA, e-mail: marco.castillo@tamu.edu

†Freddie Mac, e-mail: mofioluwasademi_ige@freddiemac.com

‡Texas A&M University, e-mail: mnavjeevan@gmail.com

§We gratefully acknowledge the support of the state government of Oyo, Nigeria for this research. The project was pre-registered at the AEA RCT Registry (AEARCTR-0004839) and received human subjects approval from Texas A&M University's IRB (IRB2018-0424D). Texas A&M University and the University of Melbourne provided funding for this project. All authors declare no conflict of interest.

# 1 Introduction

At age 15 or earlier, about half of adolescents worldwide decide whether to continue their education and, if so, along which path (e.g., academic or vocational).[1] These decisions require careful consideration of the available options as outcomes can differ dramatically across decisions (Lai, Sadoulet and de Janvry, 2009; Lucas and Mbiti, 2012). However, existing evidence on the effectiveness of information provision in altering educational choices is mixed.[2] The mixed empirical record admits two possible interpretations. Either information barriers are a minor determinant of schooling choices, or the intent-to-treat estimates typically reported in information intervention studies obscure nonuniform responses to information. Since information can be most valuable when agents are near indifference between alternatives (de Lara and Gossner, 2020), information may generate responses in opposite directions exactly when it is most valuable. Distinguishing between these interpretations is thus crucial for evaluating information RCTs.

To illustrate, consider the intent-to-treat estimate from this paper's information RCT, described in more detail below, in which we randomly provide information about returns to schooling to adolescents in Ibadan, Nigeria. In a sample rebalanced to align the treatment and control distributions of reported priors, the information intervention reduced continuation by 4.9 percentage points. This intent-to-treat estimate of $-4.9$ percentage points is consistent with a variety of causal interpretations depending on the degree of heterogeneity in students' responses to information provision. If the response to information is uniform, i.e., the information provision weakly moves all students away from continuing education, the $-4.9$ percentage point intent-to-treat estimate is interpreted as 4.9 percent of students induced not to register by information provision and none induced to register. However, the observed choice probabilities in this sample are also consistent with as many as 21.2% of students' choices being affected by information provision, with some students induced to register and others induced not to register. Because an RCT identifies only the marginal distribution of choices under each assignment, these interpretations cannot be distinguished from treatment and control choice shares alone.

In this paper, we show how elicited prior beliefs can be used to nonparametrically test for

[1] See https://data.worldbank.org/indicator/SE.COM.DURS.

[2] Jensen (2010) finds that information about the returns to education increases schooling. Other studies providing information about several aspects of educational investment report null as well as positive and negative results (Ajayi, Friedman and Lucas, 2020; Bergman, Denning and Manoli, 2019; Bettinger, Long, Oreopoulos and Sanbonmatsu, 2012; Busso, Dinkelman, Martinez and Romero, 2017; Goux, Gurgand and Maurin, 2017; Gurantz, Howell, Hurwitz, Larson, Pender and White, 2021; Hastings, Neilson and Zimmerman, 2015; Hoxby and Turner, 2015; Hyman, 2020; Loyalka, Liu, Song, Yi, Huang, Wei, Zhang, Shi, Chu and Rozelle, 2013; Barr, Bird, Castleman and Skimmyhorn, 2022).

a uniform response to information and to identify or bound the effect of information on the gross flows into and out of each choice.[3] In the case of a binary schooling decision, we show that a uniform response to information imposes the same restrictions on the joint distribution of reported prior beliefs and choices that the monotonicity condition of Imbens and Angrist (1994) imposes on the joint distribution of observed outcomes and treatment decisions in the LATE model. This equivalence allows existing instrument-validity tests (Kitagawa, 2015; Mourifie and Wan, 2017; Sun, 2023) to be applied to test for a uniform response to information with elicited prior beliefs in place of outcomes. When there are more than two possible choices, testing for a uniform response to information is equivalent to testing for unordered monotonicity (Heckman and Pinto, 2018; Sun, 2023).

In the absence of a uniform response to information, researchers can use reported prior beliefs to bound the share of students whose choices are changed by information provision, analogous to recent bounds on the shares of compliers and defiers in binary IV models (Kédagni and Mourifié, 2020; Kédagni, Wu and Cui, 2025; Noack, 2026). By comparing continuation rates between treatment and control students with the same reported prior beliefs, we obtain lower bounds on the shares induced to continue and not to continue schooling. Together, these bounds yield lower bounds on the share of agents whose decisions would change in response to information provision. For each curriculum track, we similarly bound the share induced to choose that track and the share induced to choose another option instead. Combining these bounds with observed choice shares gives a lower bound on the total share whose schooling choice changes, including students who would continue under either assignment but choose different tracks.[4] The bounds for each choice are sharp relative to the joint distribution of reported priors and choices used to construct them. All of these bounds rely only on random assignment to each information treatment arm and do not require assumptions about how students update their beliefs in response to information or how reported beliefs are related to observed choices.

Our lower bounds on the shares induced to continue and not to continue schooling are equalities under a behavioral model where students update their beliefs toward a common signal and reported priors coincide with underlying beliefs. In this model, information affects choices only through beliefs. Relative expected returns rise when a student's underlying continuation prior lies below the signal and fall when it lies above, and higher relative returns make students more willing to continue schooling. Updating toward the same signal allows a nonuniform response to information; students whose underlying priors lie below the signal

[3] By elicited priors, we mean students' reports of their prior beliefs. These reports need not coincide with their underlying beliefs.

[4] For related bounds on changes between discrete choices in a randomized welfare experiment, see Kline and Tartari (2016).

may be induced to continue schooling, while those whose priors lie above it may be induced not to continue.[5] The model also implies a testable restriction: students with reported priors below the signal are weakly more likely to continue schooling when given information, while those with reported priors above it are weakly less likely to do so.[6] The estimated pattern across reported priors is consistent with this restriction, which we interpret as suggestive evidence for the model.

Our empirical setting is a field experiment in Ibadan, Nigeria, involving more than 3,600 students who were 14 years old on average and were deciding whether to continue to senior high school and which curriculum track to pursue. Students in randomly selected schools received information on earnings and college-admission probabilities across curriculum tracks.[7] We elicited prior and posterior beliefs about earnings and educational choices in the treatment schools and a subset of control schools, and asked about willingness to pay (WTP) for information before assigning the information treatment. More than one year later, we recorded administrative outcomes on continuation and track choice.

The empirical analysis gives three main results. First, reported earnings beliefs appear to move toward the information provided, i.e, agents who are overly optimistic revise their beliefs downwards upon receiving information while pessimistic agents revise upwards. Second, among students whose earnings priors for every curriculum track lie on the same side of the corresponding signals, a group for whom our model predicts the direction of response, we reject a uniform response at the 10 percent level. The same test does not reject in the full sample, which illustrates how the model can guide where to look for offsetting responses. Third, using reported priors for both alternatives raises the estimated lower bound on the share of students whose continuation decisions change from 6.7 percent to 14.6 percent. The latter is three times the 4.9 percentage point decline in continuation in the rebalanced sample. Accounting for changes in curriculum track, we estimate that at least 17 percent of students changed their schooling choice in response to the intervention.

Our analysis builds on the literature on testable implications of monotonicity in instrumental variables models (Balke and Pearl, 1997; Heckman and Vytlacil, 2005; Kitagawa, 2015; Mourifie and Wan, 2017; Sun, 2023) and on bounds that allow for violations of monotonicity (Kédagni et al., 2025; Noack, 2026). Elicited beliefs have been used to study educational

[5] This is similar to the local monotonicity condition of Dahl, Huber and Mellace (2023), which separates compliers and defiers by their potential outcomes.

[6] Under nondifferential misreporting satisfying a monotone-likelihood-ratio condition, the effect conditional on the reported prior still changes sign at most once, from positive to negative, although the crossing can move.

[7] According to Nigeria's Joint Admissions and Matriculation Board, one in three students who apply to college is admitted; see Section 3 for details.

and economic choices (Jensen, 2010; Hoxby and Turner, 2015; Wiswall and Zafar, 2015a; Bursztyn, Gonzalez and Yanagizawa-Drott, 2020), and conditioning on elicited priors is common in information-provision experiments.[8] Balla-Elliott (2026) studies the causal effects of beliefs on outcomes and proposes an estimator of their unweighted average under heterogeneous belief updating. We instead use elicited priors to test for a uniform response to information and bound the share of agents whose choices change.

The remainder of the paper is structured as follows. Section 2 develops the theoretical framework. Section 3 describes the Ibadan field experiment and the elicitation of beliefs and WTP. Section 4 presents the empirical results. Section 5 concludes. The appendices provide proofs, additional empirical results, the WTP analysis, and the elicitation instructions.

## 2 Theoretical framework

This section develops the econometric framework used in the empirical analysis. Section 2.1 shows how reported prior beliefs can be used to test for a uniform response to information, and Section 2.2 derives lower bounds on the flows into and out of each schooling choice. These results require only random assignment and consistency. Section 2.3 then gives a model of schooling choices and belief updating under which the flows operate through beliefs and, when priors are accurately reported, the bounds for binary choices are exact. In the appendix, we use the decision model to study willingness to pay for information.

### 2.1 Testing for a uniform response to information

When an information intervention induces some students to continue schooling and others not to continue, the intent-to-treat effect understates, in absolute value, the share whose choices change. We show that a uniform response to information is empirically testable using data on elicited prior beliefs. The resulting test is analogous to existing tests in the LATE literature that jointly test instrument exogeneity and monotonicity (Balke and Pearl, 1997; Kitagawa, 2015; Mourifie and Wan, 2017; Sun, 2023).

For each student $i = 1, \ldots, n$, we assume that the researcher observes $D_i \in \{0, 1\}$, a binary variable indicating whether the student continues to senior high school, and $Y_i = (Y_{i,0}, Y_{i,1})$, where $Y_{i,0}$ and $Y_{i,1}$ are the student's reported prior beliefs about earnings if they discontinue

[8] See, for example, Cantoni, Yang, Yuchtman and Zhang (2019), Bursztyn et al. (2020), and Coibion, Gorodnichenko and Weber (2022); Haaland, Roth and Wohlfart (2023) review this practice in information-provision experiments.

and continue education, respectively. These reports need not coincide with the student's underlying beliefs. The researcher also observes $Z_i \in \{0, 1\}$, indicating whether the student receives the information intervention or not. Let $D_i(1)$ and $D_i(0)$ denote the choices student $i$ would make with and without the information treatment, so that $D_i = D_i(Z_i)$. Throughout, we maintain the following description of the experiment.

**Assumption 1** (Random Assignment and Consistency)**.** *(i)* $0 < \Pr(Z_i = 1) < 1$ *and* $Z_i \perp\!\!\!\perp (Y_i, D_i(0), D_i(1))$*, and (ii)* $D_i = D_i(Z_i)$*.*

Part (i) formalizes random assignment, including independence of the reports elicited before information is provided. Part (ii) states that observed choices are the potential choices at the realized assignment. Neither condition requires accurate reporting or restricts how students update their beliefs.

We say that the response to the information intervention is *uniform* if the intervention moves every student's choice weakly in the same direction, that is, if either $D_i(1) \geq D_i(0)$ for all students or $D_i(1) \leq D_i(0)$ for all students. This is exactly the monotonicity condition of Imbens and Angrist (1994), with the information treatment playing the role of the instrument. Define

$$M^+ = \Pr(D_i(1) = 1,\, D_i(0) = 0) \quad \text{and} \quad M^- = \Pr(D_i(0) = 1,\, D_i(1) = 0),$$

the shares of students induced to continue and not to continue education, respectively. Under Assumption 1, the intent-to-treat effect identifies the difference between $M^+$ and $M^-$,

$$\Pr(D_i = 1 \mid Z_i = 1) - \Pr(D_i = 1 \mid Z_i = 0) = M^+ - M^-, \tag{1}$$

while the sum $M^+ + M^-$ describes the gross share of students who are induced to change their schooling decision by the information intervention. A uniform response sets one of the two flows to zero: no student is moved in the direction opposite the average effect, and the intent-to-treat effect coincides, in absolute value, with the total share of students whose schooling choices change. In contrast, when both flows are positive, the absolute intent-to-treat effect understates that share.

The following proposition gives the testable restriction corresponding to the instrument-validity test of Kitagawa (2015).

**Proposition 1** (Testability of a Uniform Response)**.** *Suppose Assumption 1 holds. Then,*

*for every Borel set $B$ of reported prior beliefs,*

$$\begin{aligned}&\Pr(Y_i \in B,\, D_i = 1 \mid Z_i = 1) - \Pr(Y_i \in B,\, D_i = 1 \mid Z_i = 0) \\ &\quad = \Pr(Y_i \in B,\, D_i(1) = 1,\, D_i(0) = 0) - \Pr(Y_i \in B,\, D_i(1) = 0,\, D_i(0) = 1).\end{aligned} \tag{2}$$

*If the response to the information treatment is uniform, then the left-hand side of (2) is weakly positive for every $B$ or weakly negative for every $B$. In particular, if it is strictly positive for some set $B$ and strictly negative for another set $B'$, then $M^+ > 0$ and $M^- > 0$.*

Proposition 1 refines the intent-to-treat decomposition in (1), which follows by taking $B$ to be the entire support of $Y_i$ in (2). Under a uniform response, one of the two terms on the right-hand side of (2) is zero for every $B$, so the left-hand side has the same sign for all sets $B$; an analogous statement holds for students with $D_i = 0$. These are the restrictions tested by the instrument-validity procedures mentioned above, with reported prior beliefs playing the role of the outcome. Rejecting uniformity requires ruling out both directions. A significant one-sided intent-to-treat effect can rule out one direction; otherwise, we test both distributional orderings. Under Assumption 1, such a rejection provides evidence that the information intervention induced some students to continue education and others not to continue — in the language of the LATE model, compliers and defiers coexist.

**Remark 1.** Proposition 1 holds with any variable $X_i$ which satisfies Assumption 1 in place of $Y_i$, so in principle any baseline covariate can be used to implement the test.[9] However, an arbitrary baseline covariate may not have predictive power as to the direction of a student's response to information — two students with identical covariates may hold different underlying priors and react to the same signal in opposite ways — and the resulting test may lack power to detect violations of a uniform response. Reported priors may be informative about this direction because the effect of information on expected returns depends on students' underlying beliefs. Section 2.3 gives a model in which this dependence can be made explicit. If the decision to continue education depends on additional beliefs, e.g., earnings from alternative options, reports of these beliefs can also be used in the test. Multiple reported priors play the role of a vector-valued outcome in the instrument-validity test of Kitagawa (2015); we return to the case of multiple beliefs and multiple choices in Section 2.2.

[9] Frandsen, Lefgren and Leslie (2023) use covariates independent of treatment assignment to test for average monotonicity in "judge fixed effect" designs.

## 2.2 Bounds on changes in schooling choices

The logic underlying the testable restriction in Proposition 1 can also be used to derive lower bounds on the gross flows into and out of each schooling choice. Throughout this subsection we assume that each reported prior $Y_{i,d}$, $d \in \{0,1\}$, is continuously distributed, and we write $f_{Y_d}$ for its marginal density with respect to Lebesgue measure.[10] The distribution of interest pairs this continuous report with the discrete choice $D_i$, so we work with *subdensities*: for each $d = 0, 1$, let $q_d$ and $p_d$ be the (identified) nonnegative functions characterized by

$$\int_B q_d(y_d)\,dy_d = \Pr(Y_{i,d} \in B,\, D_i = d \mid Z_i = 0)$$
$$\int_B p_d(y_d)\,dy_d = \Pr(Y_{i,d} \in B,\, D_i = d \mid Z_i = 1)$$

for every Borel set $B$. Equivalently, $q_d(y_d) = f_{Y_d}(y_d)\Pr(D_i = d \mid Y_{i,d} = y_d, Z_i = 0)$ and $p_d(y_d) = f_{Y_d}(y_d)\Pr(D_i = d \mid Y_{i,d} = y_d, Z_i = 1)$. Analogously, let $m^+$ and $m^-$ denote the subdensities of the compliers and defiers at the reported continuation prior $Y_{i,1}$, that is the nonnegative functions characterized by

$$\int_B m^+(y_1)\,dy_1 = \Pr(Y_{i,1} \in B,\, D_i(1) = 1,\, D_i(0) = 0)$$
$$\int_B m^-(y_1)\,dy_1 = \Pr(Y_{i,1} \in B,\, D_i(0) = 1,\, D_i(1) = 0)$$

for every Borel set $B$, so that $\int m^+(y_1)\,dy_1 = M^+$ and $\int m^-(y_1)\,dy_1 = M^-$.

Under Assumption 1, random assignment makes the marginal density $f_{Y_d}$ common to both treatment arms, so that the difference $p_1(y_1) - q_1(y_1)$ can be decomposed in terms of the complier and defier subdensities

$$p_1(y_1) - q_1(y_1) = m^+(y_1) - m^-(y_1).$$

Because $m^+$ and $m^-$ are nonnegative, the positive part of $p_1 - q_1$ is a lower bound on the gross inflow density and the negative part on the gross outflow density, $\max\{p_1(y_1) - q_1(y_1), 0\} \le m^+(y_1)$ and $\max\{q_1(y_1) - p_1(y_1), 0\} \le m^-(y_1)$. Integrating yields bounds on the complier and defier shares,

$$\int \max\{p_1(y_1) - q_1(y_1), 0\}\,dy_1 \le M^+ \quad \text{and} \quad \int \max\{q_1(y_1) - p_1(y_1), 0\}\,dy_1 \le M^-.$$

[10] If $Y_{i,d}$ has a density with respect to some other dominating measure $\mu$, the Lebesgue integrals below can be read as integrals against $\mu$.

Adding these bounds gives a lower bound on the gross share of students whose choices change, $M^+ + M^-$. This lower bound exceeds the absolute intent-to-treat effect whenever both integrals in the above display are positive. As with the test of a uniform response, the bounds use only random assignment and consistency with no restrictions needed on how beliefs or decisions respond to information provision.

The analysis above extends to settings with multiple schooling choices and a vector of reported prior beliefs. For example, in the empirical intervention described below, we collect reported earnings priors for several curriculum tracks. Let $D_i(z) \in \{0, 1, \ldots, K\}$ and let $Y_i = (Y_{i,0}, Y_{i,1}, \ldots, Y_{i,K})$ collect the student's reported earnings priors over the $K+1$ alternatives.[11] Assume $Y_i$ is continuously distributed with joint density $f_Y(\cdot)$ with respect to Lebesgue measure on $\mathbb{R}^{K+1}$. For each alternative $k$, define $q_k$ and $p_k$ as the subdensities of the vector of reported priors $Y_i$ on the event $\{D_i = k\}$ in the control and treatment arms, respectively. That is, $q_k$ and $p_k$ are the nonnegative functions on $\mathbb{R}^{K+1}$ characterized by

$$\int_B q_k(y)\,dy = \Pr(Y_i \in B,\ D_i = k \mid Z_i = 0)$$
$$\int_B p_k(y)\,dy = \Pr(Y_i \in B,\ D_i = k \mid Z_i = 1)$$

for every Borel set $B \subseteq \mathbb{R}^{K+1}$. Equivalently, $q_k(y) = f_Y(y)\Pr(D_i = k \mid Y_i = y, Z_i = 0)$ and $p_k(y) = f_Y(y)\Pr(D_i = k \mid Y_i = y, Z_i = 1)$. Let $m_k^+$ and $m_k^-$ denote the subdensities of the students drawn into and pushed out of track $k$, characterized by

$$\int_B m_k^+(y)\,dy = \Pr(Y_i \in B,\ D_i(1) = k,\ D_i(0) \neq k)$$
$$\int_B m_k^-(y)\,dy = \Pr(Y_i \in B,\ D_i(0) = k,\ D_i(1) \neq k)$$

The same cancellation as in the binary case yields, for any $k = 0, 1, \ldots, K$,

$$p_k(y) - q_k(y) = m_k^+(y) - m_k^-(y).$$

Since $m_k^+$ and $m_k^-$ are nonnegative, the positive part of $p_k - q_k$ is again a lower bound on the density of students switching into $k$ at $y$, and the negative part is a lower bound on the density of students switching out of $k$. Both bounds are exact at $y$ when $m_k^+(y)m_k^-(y) = 0$, that is, when students reporting $y$ do not move both into and out of $k$. Integrating over $\mathbb{R}^{K+1}$ delivers the population bounds, stated next with Assumption 1 applied to the belief

[11] It is not important that the dimension of elicited prior vector is the same as the number of potential choices. In our empirical analysis we use multiple measures of prior beliefs in a binary choice model. The analysis of Proposition 2 holds with any vector of beliefs $Y_i$ satisfying Assumption 1.

vector $Y_i$.

**Proposition 2** (Sharp Lower Bounds on Gross Flows)**.** *Suppose Assumption 1 holds with the vector $Y_i$ in place of $(Y_{i,0}, Y_{i,1})$ and $Y_i$ has density $f_Y(\cdot)$ with respect to Lebesgue measure. For each alternative $k$, define*

$$M_k^+ = \Pr(D_i(1) = k, D_i(0) \neq k), \qquad M_k^- = \Pr(D_i(0) = k, D_i(1) \neq k).$$

*Then,*

$$\int \max\{p_k(y) - q_k(y), 0\} dy \leq M_k^+$$

*and*

$$\int \max\{q_k(y) - p_k(y), 0\} dy \leq M_k^-.$$

*These lower bounds are jointly sharp in the sense that there exists some joint distribution of reported prior beliefs and potential choices consistent with the observable distribution and Assumption 1 under which both inequalities hold with equality for every $k$. If, moreover, for almost every report $y$, there are no simultaneous switches into and out of $k$,*

$$\Pr(D_i(1) = k, D_i(0) \neq k \mid Y_i = y) \Pr(D_i(0) = k, D_i(1) \neq k \mid Y_i = y) = 0,$$

*then both inequalities hold with equality for every joint distribution of reported prior beliefs and potential choices consistent with the observable distribution, Assumption 1, and this no-offsetting condition.*

Proposition 2 generalizes the binary case presented above, in which $M_1^+ = M^+$ and $M_1^- = M^-$. The lower bounds and their equality condition also apply when $Y_i$ is replaced by any subvector, including a single prior. These bounds do not require rational updating or restrictions on the channel through which assignment changes choices. With several choices, a uniform response requires that, for each $k$, either $M_k^+ = 0$ or $M_k^- = 0$: no alternative both gains and loses students because of the intervention. The direction of the response may differ across alternatives. This is the unordered monotonicity condition for a binary instrument (Heckman and Pinto, 2018; Sun, 2023), and Proposition 1 applies separately to each choice indicator.

The bounds also restrict the joint distribution of choices under the two assignments. To see this, collect the potential choices in a transition matrix whose $(k, \ell)$ entry $a_{k\ell} = \Pr(D_i(0) = k,\ D_i(1) = \ell)$ is the share of students who would choose alternative $k$ absent the intervention and $\ell$ under it. Table 1 illustrates this matrix for the choice among dropout ($k = 0$), Arts ($k = 1$), and combined Commerce and Science ($k = 2$): its diagonal collects the students

whose choice is unchanged and its off-diagonal entries the switchers. Random assignment identifies the two margins — the row sums are the control choice shares $\Pr(D_i(0) = k)$ and the column sums the treatment choice shares $\Pr(D_i(1) = \ell)$. However, the marginal distributions alone do not point identify the off-diagonal entries since only a single potential outcome is observed for each student.

For each alternative $k$, the share switching into $k$, $M_k^+$, is the sum of the off-diagonal entries in column $k$, and the share switching out of $k$, $M_k^-$, the sum of those in row $k$; i.e., in Table 1, $a_{0,1} + a_{0,2}$ is the share switching from dropout into continuation and $a_{1,0} + a_{2,0}$ the share switching from continuation into dropout. The density integrals of Proposition 2 bound these sums of off-diagonal entries from below. In our empirical application, we report a transition matrix that minimizes the sum of the off-diagonal entries subject the non-negativity constraints, the observed marginal distributions, and the bounds from Proposition 2. This is a linear program. Because the true joint distribution must satisfy these restrictions, this minimum is a lower bound on the gross share of students whose schooling choices change.

Table 1: Transition matrix and observed choice margins

| | Treatment | | | |
|---|---|---|---|---|
| Control | Dropout | Arts | Comm & Sci | Total |
| Dropout | $a_{0,0}$ | $a_{0,1}$ | $a_{0,2}$ | 0.08 |
| Arts | $a_{1,0}$ | $a_{1,1}$ | $a_{1,2}$ | 0.39 |
| Commerce & Science | $a_{2,0}$ | $a_{2,1}$ | $a_{2,2}$ | 0.53 |
| Total | 0.13 | 0.36 | 0.51 | 1.00 |

Notes: Marginal choice shares are from the experiment described in Section 3, estimated on the belief subsample of Table 9; entries may not sum to totals due to rounding.

We next give a model in which a student's prior beliefs determine the direction in which information can change their schooling decision. When these beliefs are accurately reported, the bounds in Proposition 2 hold with equality for the binary continuation decision.

### 2.3 Schooling choices and belief updating

Consider a student choosing whether to continue to senior high school ($D_i = 1$) or discontinue education ($D_i = 0$). Normalizing the utility of discontinuing to zero, write the utility of continuing under assignment $z$ as $U_i(z) + e_i$. Here $U_i(z)$ is the component of expected relative utility that may change with information, such as the expected earnings gain from continuing, while $e_i$ captures considerations held fixed across assignments, such as the consumption value of schooling. The student continues whenever $U_i(z) + e_i \geq 0$.

Holding $e_i$ fixed, increasing $U_i(z)$ can move choices only toward continuation. This individual monotonicity restriction underlies within-subject belief experiments such as Wiswall and Zafar (2015b). To use it to separate compliers and defiers, we need to determine the direction in which information changes each student's expected return to continuation.

Let $\tilde{Y}_i = (\tilde{Y}_{i,0}, \tilde{Y}_{i,1})$ denote student $i$'s beliefs about earnings from discontinuing and continuing education at the time of educational decision and $\tilde{Y}_i(z) = (\tilde{Y}_{i,0}(z), \tilde{Y}_{i,1}(z))$ denote the latent potential time-of-decision beliefs the student would hold under information assignment $z$. These two quantities are related via

$$\tilde{Y}_i = \tilde{Y}_i(Z_i) = Z_i\tilde{Y}_i(1) + (1 - Z_i)\tilde{Y}_i(0).$$

The reports $Y_i$ used in Sections 2.1 and 2.2 are measures of $\tilde{Y}_i(0)$, the beliefs held in the absence of the information intervention, i.e, "prior" beliefs. The analysis of this section does not assume any measure of the "posterior" beliefs $\tilde{Y}_i(1)$ is observed by the researched. Define the relative expected return by $R_i(z) = \tilde{Y}_{i,1}(z) - \tilde{Y}_{i,0}(z)$. The following assumption collects the restrictions of the decision model.

**Assumption 2** (Decision Model)**.** *Assume that*

*(i) All treated students receive the same signal $S$ about earnings from continuing education, and updating satisfies*

$$\text{sign}(R_i(1) - R_i(0)) = \text{sign}(S - \tilde{Y}_{i,1}(0)).$$

*(ii) Expected relative utility satisfies $U_i(z) = \nu(R_i(z))$ for an increasing, nonconstant function $\nu : \mathbb{R} \to \mathbb{R}$. Equivalently, potential schooling decisions satisfy*

$$D_i(z) = \mathbf{1}\{\nu(R_i(z)) + e_i \geq 0\}, \qquad z \in \{0, 1\}. \tag{3}$$

*(iii) Information assignment is independent of potential beliefs and idiosyncratic preferences:*

$$Z_i \perp\!\!\!\perp (\tilde{Y}_i(0), \tilde{Y}_i(1), e_i).$$

Assumption 2(i) requires that information raise relative expected returns to continuing education when students' prior beliefs about expected returns to continuing education are below the signal and lower the relative expected returns when these prior beliefs are above the signal. Receiving information about the expected returns from continuing education may also affect a students perceived returns from discontinuing education, say by raising or lowering

a student's beliefs about overall average income. The restriction in Assumption 2(i) does not rule this out, but places restrictions on how continuation and noncontinuation earnings beliefs jointly respond to information provision.

Assumption 2(ii) requires that the expected utility of continuing education is increasing in the relative expected returns. This in turn implies the Roy-model (Roy, 1951) type selection rule for the schooling choice given in (3) — a student continues to senior high school when the expected utility gain from continuing is nonnegative.[12] Finally, Assumption 2(iii) maintains that assignment to the information treatment is independent of both potential beliefs and the idiosyncratic shock.

For a student with $\tilde{Y}_{i,1}(0) < S$, Assumption 2(i)-(ii) imply that $D_i(1) \geq D_i(0)$ while, for a student with $\tilde{Y}_{i,1}(0) > S$, the inequality is reversed. Assumption 2(iii) then yields the following restriction on the observed conditional choice probabilities.

**Proposition 3** (Location of Responses)**.** *Under Assumptions 1 and 2, for almost every* $\tilde{y}_1 < S$*,*

$$\Pr(D_i = 1 \mid \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = 1) \geq \Pr(D_i = 1 \mid \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = 0).$$

*The inequality is reversed if* $\tilde{y}_1 > S$*.*

Proposition 3 does not imply a uniform response to information, students on opposite sides of the signal can therefore respond in opposite directions. Conditional on the underlying continuation prior $\tilde{y}_i$, however, responses are uniform. If this prior is accurately reported, i.e $Y_{i,1} = \tilde{Y}_{i,1}(0)$, conditioning on its report rules out offsetting flows and restriction needed for the bounds in Proposition 2 to hold with equality is satisfied.

**Corollary 1** (Exact Bounds under a Common Signal)**.** *Suppose Assumptions 1 and 2 hold,* $Y_{i,1} = \tilde{Y}_{i,1}(0)$ *almost surely, and* $Y_{i,1}$ *has a density with respect to Lebesgue measure. Then, for Lebesgue-almost every* $y_1$*,*

$$\Pr(D_i(1) = 1, D_i(0) = 0 \mid Y_{i,1} = y_1) \Pr(D_i(0) = 1, D_i(1) = 0 \mid Y_{i,1} = y_1) = 0,$$

*and the bounds of Proposition 2 applied with* $Y_i = Y_{i,1}$ *hold with equality:*

$$\int \max\{p_1(y_1) - q_1(y_1), 0\}\, dy_1 = M^+ \quad \text{and} \quad \int \max\{q_1(y_1) - p_1(y_1), 0\}\, dy_1 = M^-.$$

[12] See Heckman and Honoré (1990) for the empirical content of the Roy model and Heckman and Vytlacil (2005) for an analysis of treatment effects in selection models. Vytlacil (2002) relates latent-index selection models to the independence and monotonicity conditions of the LATE framework.

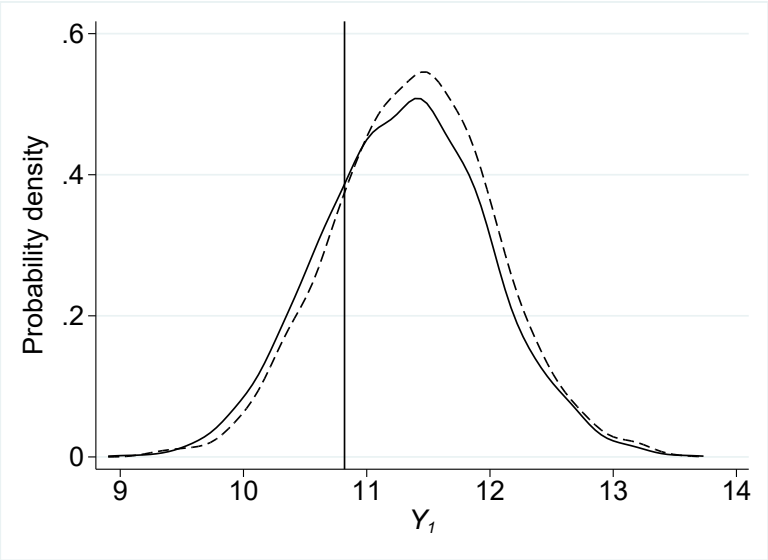



Figure 1: How information changes schooling choices in the simulation. The horizontal axis shows students' prior beliefs about earnings from continuing education ($Y_1$), which are accurately reported in the simulation. The solid curve is the treatment subdensity ($p_1$), and the dashed curve is the control subdensity ($q_1$); the area under each curve is the share of that group continuing to senior high school. The vertical line marks the earnings signal, $S = 10.8$. The area between the curves to its left represents students induced to continue (compliers); the area to its right represents students induced not to continue (defiers).

Under the conditions of Proposition 3 and Corollary 1, using choice probabilities conditional on the continuation prior alone is sufficient to identify the share of information compliers and defiers, so additional beliefs cannot tighten the population bounds. If beliefs are misreported, however, students who report the same continuation prior may update their true beliefs in opposite directions. Additional measures of students' underlying beliefs may then help tighten the bounds in Proposition 2. This motivates the comparison in Section 4 between bounds estimated using the continuation prior alone and those using priors for both alternatives.

The conditions of Proposition 3 and Corollary 1 also yield a testable restriction. Under these conditions, multiplying the conditional choice probabilities by the common marginal density of reported priors gives

$$\begin{cases} p_1(y_1) \geq q_1(y_1) & \text{if } y_1 < S, \\ p_1(y_1) \leq q_1(y_1) & \text{if } y_1 > S, \end{cases}$$

Thus $p_1 - q_1$ can be positive only below the signal and negative only above it. The one-sided tests described in Section 2.1 can be used to assess these restrictions.

Figure 1 illustrates the restriction using simulated data with $Y_i = \tilde{Y}_i(0)$. The subdensity

$p_1$, represented by a solid line, lies above $q_1$, represented with a dashed line, to the left of the signal $S$ and below $q_1$ to the right of the signal. The area between $p_1$ and $q_1$ below the signal equals the share of compliers, while the area above the signal equals the share of defiers. Because both groups are present, the absolute intent-to-treat effect is smaller than the share whose choices change.[13]

**Remark 2.** In our experiment, signals vary by gender and curriculum track. Suppose Assumption 2(i) holds with a student-specific signal $S_i$ in place of $S$ and priors are accurately reported. Students with priors below the smallest signal can then be induced only toward continuation, while those above the largest signal can be induced only away from continuation. Between these values, a given prior can lie below the signal of one student and above that of another, so students reporting the same prior may respond in opposite directions. The pooled difference $p_1 - q_1$ then need not change sign exactly once, and a sign change need not occur at the average signal. When information concerns several tracks, the resulting belief revisions need not satisfy the sign restriction on relative returns either. We accordingly interpret the pooled density comparison as suggestive evidence for the model. Students whose priors for every track lie below the corresponding signals for their gender, or above all of them, are an exception, since all of their track beliefs are revised in the same direction. Section 4.3 uses these students to test for a uniform response.

### 2.3.1 A testable restriction under misreporting

In our experiment, elicited posterior beliefs differ from priors even among students who receive no new information (see Section 4.1). This suggests that reported beliefs may contain error. Even under a common signal, students who give the same report may then have underlying priors on opposite sides of the signal, allowing their responses to offset. The bounds in Proposition 2 remain valid, but need no longer hold with equality. Under certain conditions on the misreporting error, however, the behavioral model of Section 2.3 remains testable. We present these conditions and the testable restriction below.

To simplify notation, write $Y \equiv Y_{i,1}$ for the reported continuation prior and $\tilde{Y}(0) \equiv \tilde{Y}_{i,1}(0)$ for the underlying belief without information, and suppose $Y$ has a density with respect to Lebesgue measure. Let

$$Q(y) \coloneqq p_1(y) - q_1(y)$$

denote the difference between the treatment and control subdensities defined in Section 2.2.

[13] In the simulated data $\bar{Y}_1 = 11.34$, $\bar{Y}_0 = 9.87$, $\rho(Y_0, Y_1) = 0.19$, $S = 10.8$, $\theta = 0.8$ (the weight on the signal in the posterior; see Section 4.1) and $\Pr(D = 1 \mid \tilde{Y}_1, \tilde{Y}_0) = (1 + \exp(-1 - 5(\tilde{Y}_1 - \tilde{Y}_0)))^{-1}$. These values are calibrated to the data described below.

**Assumption 3** (Nondifferential Monotone Misreporting)**.** *Suppose that:*

*(i) Misreporting is nondifferential:*

$$Y \perp\!\!\!\perp (Z_i, \tilde{Y}_i(0), \tilde{Y}_i(1), e_i) \mid \tilde{Y}(0).$$

*(ii) For any $y_H > y_L$, the conditional densities of $\tilde{Y}(0)$ given these reports have the same support and*

$$\frac{f_{\tilde{Y}(0)|Y}(\tilde{y} \mid y_H)}{f_{\tilde{Y}(0)|Y}(\tilde{y} \mid y_L)}$$

*is weakly increasing in $\tilde{y}$.*

Part (i) requires that the misreporting error is independent of treatment assignment, potential beliefs, and other factors that may influence the students' decision, i.e "nondifferential misreporting". In other words, after conditioning on the students true beliefs, a student's report contain no further information about assignment, updating, or preferences. Part (ii) requires the reported beliefs to satisfy a monotone likelihood ratio property, analagous to that of Athey (2002). This assumption weakly orders the beliefs associated with different reports: as the report rises, the relative likelihood of high versus low underlying priors cannot fall.

Together, the conditions in Assumption 3 allow us to test the model presented in Section 2.3. This testable restriction is given below.

**Proposition 4** (Single Crossing under Misreporting)**.** *Under Assumptions 1 to 3, the observable density difference $Q(y)$ has at most one sign change on the support of $Y$, and if a sign change occurs it is from positive to negative. Equivalently, there exists a cutoff $S^*$, possibly at an endpoint of the support of $Y$, such that*

$$Q(y) \geq 0 \quad \textit{for almost every } y < S^*, \qquad Q(y) \leq 0 \quad \textit{for almost every } y > S^*.$$

Proposition 4 shows that, under Assumption 3, the single-crossing restriction of the decision model continues to hold when prior beliefs are misreported. Information weakly increases the probability of continuation conditional on reports below $S^*$ and weakly decreases it conditional on reports above $S^*$. The cutoff need not coincide with the signal $S$, so it must be treated as unknown when testing this restriction. In particular, evidence that information reduces continuation at lower reported priors and increases it at higher reported priors would reject the model and reporting assumptions.

This restriction concerns the average response conditional on reported priors. Students with the same reported prior may still respond in opposite directions, so the restriction does not imply that the bounds in Proposition 2 hold with equality. The bounds remain valid under Assumption 1 alone.

**Remark 3.** The proof of Proposition 4 uses only the fact that the function

$$f(\tilde{y}) = \Pr(D_i(1) = 1 \mid \tilde{Y}(0) = \tilde{y}) - \Pr(D_i(0) = 1 \mid \tilde{Y}(0) = \tilde{y})$$

changes sign at most once, from positive to negative. The reporting assumptions preserve this sign pattern when averaging over the underlying beliefs associated with each report.

# 3 Experimental design and implementation

We provide background on the study setting and details about the data collection.

## 3.1 Background

The secondary education system in Nigeria is divided into junior high school and senior high school. After spending six years in primary school, students attend three years of junior high school and can spend three years in senior high school. Students can choose between an academic or a vocational curriculum at the end of junior high school. Students wishing to proceed with the academic option in senior high school have three curriculum options: arts, commercial/social sciences, and science. The vocational track also provides students with different areas for specialization. The diversity of the course curricula in senior high school allows students to choose their future career paths. Students are exposed to both mainstream academic courses and vocational courses. At the end of junior high school, or grade 9, students take a statewide examination—the Basic Education Certificate Examination (BECE)—which allows the transition to the next level of schooling.

## 3.2 Overview

The present study recruited students in their last year of junior high school who had to decide whether to continue to senior high school, go to vocational school/take an apprenticeship, or drop out of school entirely. The study occurred in Ibadan, the capital of Oyo, Nigeria, and Nigeria's third most populous city (3.2 million). The study was conducted with the approval

of the State of Oyo's Ministry of Education, Science, and Technology. The experiment had five stages: recruitment, baseline data collection, information provision in treated schools, collection of endline data, and collection of administrative data on educational choices.

In the first stage, schools were recruited and consent was obtained. Schools were assigned to three experimental conditions: treatment, impure control, and pure control. All students completed a baseline survey that collected demographic information, attitudes toward schooling, a proposed curriculum-track choice, career aspirations, and beliefs about educational outcomes. Students in the treatment and impure-control schools then completed three information-valuation tasks and a detailed belief module. Following this elicitation, students in the treatment schools received information on average earnings by track and college-admission probabilities. Students in the treatment and impure-control schools then answered the same belief questions a second time. This design follows the college-choice study of Wiswall and Zafar (2015a). The survey ended by collecting hypothetical time preferences and incentivized risk preferences.

### 3.3 Measures of prior and posterior beliefs

We elicited subjective probabilities over educational outcomes: stopping after junior high school, entering vocational school or an apprenticeship, leaving or completing senior high school, and leaving or completing college in a stated curriculum track. Students also reported employment and earnings beliefs by path, including the probabilities of earning at least N50,000, N100,000, and N200,000. Parallel questions asked about a typical student. Table 2 provides the wording of one educational-attainment elicitation.

Table 2: Example of self-belief elicitation

| **What are the chances that you will** | **Number** |
|---|---|
| go to art class? | |
| go to science class? | |
| go to commercial/social science class? | |
| go to vocational school after JSS3? | |
| drop out of school after JSS3? | |
| TOTAL: THE TOTAL SHOULD ADD UP TO 100 | |

The information treatment included statistics on earnings and labor supply in Nigeria, together with population-level college application and admission rates. This information came from the Joint Admissions and Matriculation Board (JAMB) and Stutern (2018).[14] These

[14] A total of 5,219 Nigerians who graduated during 2013–2017 completed the Stutern survey. Data collec-

were the most recent available data on admissions and recent graduates' labor-market outcomes when we designed the study.

Section 3.5 presents the information provided in the study. The intervention simplified the approach of Wiswall and Zafar (2015a); see Haaland et al. (2023) for a review of related information experiments. Repeated belief measures allow us to compare revision toward the information in the treatment group with revision caused by repetition, reflection, or reporting error in the impure-control group. This comparison provides evidence on the updating-toward-the-signal condition in Assumption 2(i); it does not by itself validate the other restrictions of the common-signal model.

### 3.4 Elicitation of willingness to pay for information

We elicited WTP for different pieces of information using a multiple price list (MPL), a discretization of the Becker–DeGroot–Marschak (BDM) mechanism. The MPL included prices from N0 to N200 in increments of N25. Participants responded to three MPLs. The first asked for their WTP for information on college admissions, the second for information on wages by curriculum track, and the third for both pieces. The bundle provides a weak scope check: it should not be valued less than either component.

To embed the WTP elicitation in the information RCT, we drew prices from the binary set $\{0, 250\}$. These draws were fixed at the classroom/school level to avoid the expected spillover effects if randomization were done at the individual level. As a consequence, either all of the students in a school were assigned to the information treatment group or they were all assigned to a no-information group. Since we could not ask students to pay for information with their out-of-pocket money, we provided all the students in our study with N200[15] that they could use in different experimental tasks, including the WTP tasks. They were told that payments would be calculated based on their choices in one of the tasks chosen at random.[16] Providing participants with money is consistent with common practice in experimental economics.

Previous research using the BDM shows that both the distribution of prices (Mazar, Koszegi

tion took place between February 8 and May 15, 2018. Stutern.com recruited respondents by email and social media, and offline surveys in Edo, Enugu, Ibadan, Imo, and Kaduna reached graduates in less-connected locations.

[15]This amount is about 1.1% of Nigeria's minimum wage and enough to cover a student's lunch. Our study balances the need for the salience of payoffs and the risk associated with transferring money to minors.

[16]In practice, one of the three MPLs was chosen randomly to determine the cost of information. If the price was N0, we provided all the information. We did this because we could not detect the effects of different sets of information since the treatments were assigned at the school level to avoid potential contamination.

and Ariely, 2014) and the upper bound of the distribution of prices (Bohm, Linden and Sonnegard, 1997) can affect elicited valuations. To avoid these issues in our experiment, we indicated to participants that the prices could take values as low as N0 and higher than N200. However, we cannot test whether the price set altered the average valuation of information.[17]

### 3.5 Randomization and implementation

**Information intervention:** The intervention provided information to students in randomly selected schools. Students in the treatment schools received information about average wages, the percent working full time, and the percent earning more than N50,000 and N100,000 for the different curricula/tracks, as well as the proportion of males (females) who applied and were admitted to college across the three tracks. Figure 2 shows how the information was presented; it was done in this way to make it easy to understand. We consulted with the State's Ministry of Education to ensure accurate information.

In the state of Oyo, there are over 600 public secondary schools. The sample of schools includes 115 coeducational junior high schools out of a universe of 133 in Ibadan city. These schools are evenly distributed in four areas of the city across the 5 local government areas. The subset of schools randomized into the study was visited by enumerators carrying official letters from the government, our IRB approval letter, a study overview/permission letter to obtain permission to visit the school, and a proposed date and time to visit.

The preregistered study planned for 32 schools: 16 treatment schools and 16 control schools, with 5,200 students.[18] We divided the control group so that half would complete the detailed belief module and half would not. This created the impure- and pure-control groups and allowed us to distinguish the information effect from any effect of repeated belief elicitation. We planned for 16 treatment, 8 pure-control, and 8 impure-control schools. Students were scheduled to take their exit examinations in June or July 2020 and to choose a senior-high-school track by September 2020. The study was designed to detect a five percentage point change in dropout with 80% power at the 5% significance level.[19]

The first stage of the study was conducted between November 8 and December 3, 2019. To allow for potential school attrition, we recruited additional schools and retain the 36 that

[17] Different approaches to eliciting valuations that would further minimize these issues (e.g., Allcott, Braghieri, Eichmeyer and Gentzkow, 2020; Mosquera, Odunowo, McNamara, Guo and Petrie, 2020) were not feasible because students needed to be provided with money to participate.

[18] AEA RCT Registry AEARCTR-0004839.

[19] Administrative dropout data implied an intraclass correlation of 0.02.

## Admissions information



**Arts**

- 15% of girls applied to Arts in university
- 10% of boys applied to Arts in university
- Of the girls that applied, 32% got admission
- Of the boys that applied, 30% got admission




**Science**

- 53% of girls applied to science in university
- 60% of boys applied to science in university
- Of the girls that applied, 26% got admission
- Of the boys that applied, 29% got admission




**Social-sciences**

- 33% of girls applied to social-science in university
- 30% of boys applied to social-science in university
- Of the girls that applied, 26% got admission
- Of the boys that applied, 27% got admission

Source: JAMB and CINFORES (2017)

## Salary information

### *Female*




This information is from a survey of university graduates in Nigeria. Among recent ***female*** graduates from a university with a Bachelor's degree in each of the above fields:

| | | | |
|---|---|---|---|
| • The percentage that are working full time is | **44.6%** | **45.6%** | **45.9%** |
| • The average monthly salary of those that are working full time is | **₦59,158** | **₦68,740** | **₦64,696** |
| • The percentage of those who are working full time that earn more than N50,000 monthly | **45.74%** | **56.2%** | **52.23%** |
| • The % of those who are working full time that earn more than N100,000 monthly is | **5.26%** | **7.85%** | **4.86%** |

### *Male*




This information is from a survey of university graduates in Nigeria. Among recent ***male*** graduates who just graduated from a university with a Bachelor's degree in each of the above fields:

| | | | |
|---|---|---|---|
| • The percentage that are working full time is | **48.54%** | **50.4%** | **49.35%** |
| • The average monthly salary of those that are working full time is | **₦53,100** | **₦77,849** | **₦71,924** |
| • The % of those who are working full time that earn more than N50,000 monthly is | **44%** | **60%** | **53.16%** |
| • The % of those who are working full time that earn more than N100,000 monthly is | **6%** | **25.4%** | **8.42%** |

| | | | |
|---|---|---|---|
| Among recent university graduates who received a Bachelor's degree in the above fields<br>The percentage of those who are women is | **65.52%** | **35.64%** | **56.62%** |

Source: JAMB and CINFORES (2017)

Figure 2: Information provided

completed data collection. There were 18 treatment schools ($N = 1{,}925$), 6 impure-control schools ($N = 658$), and 12 pure-control schools ($N = 1{,}054$). The main deviations from the original protocol were a reduction in the implementation time allowed by schools[20] and the delay of student examinations until August 2020 due to Covid-19. Figure 3 presents the design and implementation graphically.[21]

Table 3 presents basic statistics for the sample and its comparability across the three groups.[22] The average age of participants was 14 years. There were slightly fewer females in the sample than males. About 40 percent declared themselves to be Christian and had roughly four siblings. Almost 80 percent of the students lived with both of their parents. Around 10 percent of the students declared they had repeated at least one grade. The three groups are balanced in all variables we checked, except for the number of females. The pure control group had slightly more females than the information treatment group.

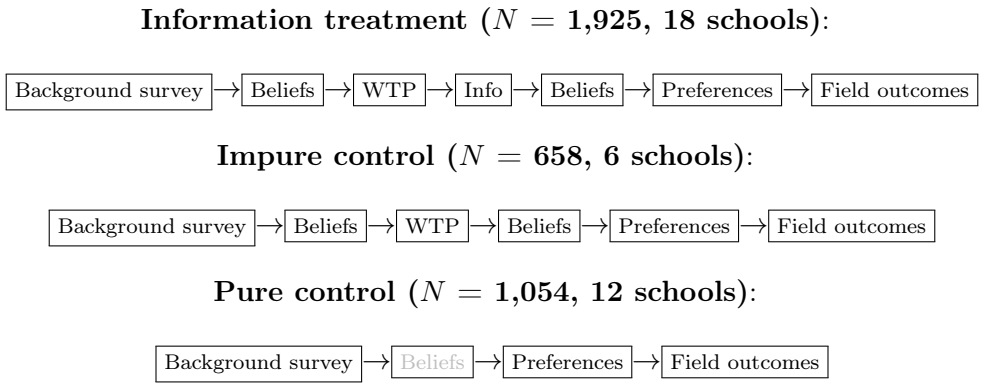



Figure 3: Study implementation

Notes: A subset of the belief questions and preference measures was collected from all participants. Preferences were collected at the end of the survey. Students were provided with N200 for the WTP (where applicable) and preference elicitations.

### 3.6 Outcomes collected

The study has two main classes of outcomes: beliefs about educational choices and realized educational choices. Beliefs were collected before and after the information intervention,

[20] We reclassified three schools from impure control to pure control because time constraints prevented collection of the detailed belief data. Two schools assigned to the information treatment (184 students) could not be visited for the information session and are counted as pure controls.

[21] A 2018 pilot assessed the feasibility of WTP elicitation. It covered 195 students in three schools, including two single-sex schools excluded from this study. Pilot and study WTP responses were comparable; results are available from the authors upon request.

[22] All classrooms in the last year of junior high school of the participating schools were visited by enumerators.

Table 3: Characteristics by treatment group

| Variable | (1) Cntrl (0) | (2) Info (1) | (3) Imp. Cntrl (2) | (4) (1)-(0) | (5) (2)-(0) | (6) (1)-(2) |
|---|---|---|---|---|---|---|
| Age | 13.978 | 14.018 | 14.189 | -0.041 | 0.083 | 0.007 |
| | (1.591) | (1.398) | (1.549) | (0.141) | (0.183) | (0.203) |
| Female | 0.498 | 0.448 | 0.465 | -0.037 | -0.029 | -0.029 |
| | (0.500) | (0.497) | (0.499) | (0.028) | (0.050) | (0.039) |
| Mother's years of education | 12.511 | 12.484 | 12.644 | -0.079 | 0.443 | -0.457 |
| | (3.677) | (3.540) | (3.216) | (0.264) | (0.359) | (0.310) |
| Christian | 0.394 | 0.364 | 0.394 | -0.030 | 0.021 | -0.052 |
| | (0.489) | (0.481) | (0.489) | (0.058) | (0.097) | (0.074) |
| No. of siblings | 3.985 | 3.892 | 3.930 | -0.072 | -0.075 | -0.030 |
| | (3.118) | (2.636) | (2.290) | (0.187) | (0.203) | (0.205) |
| Two-parent household | 0.786 | 0.785 | 0.757 | 0.010 | -0.014 | 0.007 |
| | (0.411) | (0.411) | (0.429) | (0.019) | (0.023) | (0.026) |
| Grades | 0.617 | 0.616 | 0.605 | 0.004 | -0.005 | 0.001 |
| | (0.144) | (0.141) | (0.135) | (0.011) | (0.016) | (0.013) |
| Repeated a grade | 0.097 | 0.093 | 0.156 | -0.027 | 0.057 | -0.048 |
| | (0.296) | (0.290) | (0.363) | (0.022) | (0.042) | (0.033) |
| Observations | 1,054 | 1,925 | 658 | 3,637 | 1,528 | 2,767 |

Notes: Columns 1-3 show the average by treatment group. The numbers in parentheses are used to indicate pairwise comparisons between the groups in columns 4-6. Missing data items are replaced with the variable's mean over the entire sample.

and administrative educational outcomes were collected more than a year after the baseline survey.

**Earnings beliefs:** Students reported expected monthly earnings after dropping out of senior high school, after finishing senior high school, and after completing a degree in arts, science, or commerce. They reported each path at graduation and at ages 30 and 50, as well as the probability that earnings would exceed predetermined thresholds. After winsorizing the reports, we define the continuation index $Y_1$ as the log of the arithmetic mean across the senior-high-school and degree reports at the three ages. We construct the noncontinuation index $Y_0$ from the three reports for the junior-high-school/dropout path. These indexes summarize reported monthly earnings priors; they do not measure realized lifetime earnings or fully describe beliefs about each alternative. Table A.1 compares the component priors across the treatment and impure-control groups. The impure-control means for senior high school at ages 21 and 30 are higher; Section 4 describes how the density analysis addresses prior imbalance.

**Enrollment beliefs:** At baseline, students reported their subjective probability of continuing beyond junior high school, choosing each senior-high-school and college track (arts, commerce, science, or vocational education), and leaving after junior high school, senior high school, or college. Repeated reports measure how these beliefs changed after information provision.

**Attendance and continuation:** We obtained school administrative data on attendance and enrollment. In particular, we recorded whether a student took a junior high school exit exam, the grade obtained, and whether they registered for senior high school. The binary outcome in the main analysis equals one for senior-high-school registration and zero for nonregistration; we refer to its complement as dropout or noncontinuation. Noncontinuation can include a vocational-school or apprenticeship path as well as leaving school. For students to qualify for admission to senior high school and higher education, nationwide examinations are held each year. Because exam scores determine a student's future educational choices, schools tend to stress memorization of facts rather than creative problem-solving. Students must pass at least six subjects to proceed to senior high school at the same or a different institution.

**Curriculum choice:** We obtained administrative data from schools on high school track choices. The curriculum tracks include arts, commerce, and science. All senior secondary students must study English, mathematics, one science, and one Nigerian language course.[23] The remaining electives are selected based on the student's interest in the sciences, the social sciences, or the arts.

The state of Oyo does not have a centralized system with all students' data. To minimize potential biases due to non-response, we visited each high school in the study to collect information on registration. This allowed us to cross-check whether students changed schools after junior high school and their decisions. We also conducted a phone survey for all students not found in the school records. This procedure allowed us to determine the outcomes for over 95% of the sample. We do not find significant differences in missing data across the treatments.

# 4 Results

## 4.1 Information updating

We use a normal–normal learning model to benchmark information updating. Suppose a student's prior about log earnings on a given path is normally distributed, $\mathcal{N}(\mu_0, \sigma_0^2)$, and that the student reports its mean, $\mu_0$. The student then receives a signal, $S$, which they interpret as their log earnings plus independent noise, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. The Bayesian posterior

[23] Science is not required for non-science tracks.

is normal with mean

$$\mu_1 = (1-\theta)\mu_0 + \theta S,$$

where $\theta = \frac{\sigma_0^2}{\sigma_0^2+\sigma^2}$. If the student again reports the mean, the revision between the two reports is

$$\mu_1 - \mu_0 = \theta(S - \mu_0).$$

The coefficient $\theta$ is the weight placed on the signal; a more precise prior implies a smaller value.

Figure A.2 shows average earnings priors by educational choice and reporting age. Expected earnings rise with age and schooling and are lowest for junior high school and highest for a science degree. Figure A.3 plots age-30 log-earnings priors relative to the gender-specific signal for the arts, science, and commerce tracks. Priors vary substantially around the signals, and most students expect earnings above the information provided. This heterogeneity leaves room for the intervention to move choices in both directions.

Repeated elicitation in the impure-control group helps distinguish updating caused by the information from revision caused by reflection, reporting error, or mean reversion. Table 4 regresses the change in log earnings beliefs on the gap between the log signal and the initial log belief, interacted with information assignment.[24] The regressions use reports at ages 21 and 30, the ages closest to the information shown to students; gender-specific signals provide additional variation. The interaction coefficient ranges from 0.12 for science at age 30 to 0.24 for commerce at age 30. Pooling questions gives similar estimates. These magnitudes are close to the 0.26–0.37 range reported by Hjort, Moreira, Rao and Santini (2021).[25]

Table A.2 applies the same specification to beliefs about earning above stated thresholds. These probabilities respond less than earnings levels, although the normal-learning benchmark may fit bounded probabilities poorly. The positive interaction coefficients are consistent with the updating-toward-the-signal condition in Assumption 2(i).

### 4.2 Treatment effects on senior-high-school continuation

We next turn to realized educational choices. Table 5 reports linear probability models of senior high school continuation on the experimental assignments. Students in the impure- and pure-control groups are 3.8 and 3.7 percentage points more likely to continue than treated

[24] We winsorize beliefs at 1%. Log beliefs are approximately normal. We did not elicit prior admission probabilities, so we cannot estimate the same specification for that information.

[25] Revision in the impure-control group indicates measurement error in reported beliefs. Fuster, Perez-Truglia, Wiederholt and Zafar (2022) find a smaller version of the same pattern for housing-price beliefs.

Table 4: Belief updating

| | Salary at 21 | | | Salary at 30 | | |
|---|---|---|---|---|---|---|
| | Arts | Comm | Science | Arts | Comm | Science |
| (signal − prior) | | | | | | |
| × Info treatment | 0.212 | 0.198 | 0.239 | 0.234 | 0.245 | 0.126 |
| | (0.046) | (0.045) | (0.045) | (0.043) | (0.046) | (0.043) |
| Observations | 2207 | 2206 | 2187 | 2228 | 2226 | 2202 |
| Adj R2 | 0.364 | 0.340 | 0.340 | 0.369 | 0.308 | 0.348 |

Notes: Salary at 21 for Arts refers to the answer to the question: "Imagine that you enrolled in the arts track and studied one of the arts courses at university, and tell me how much you think you would be paid monthly if you have just graduated and now work full time?" The labels for the other columns follow the same pattern. The dependent variable is the difference between the belief elicited the second time and the belief elicited the first time. Beliefs are winsorized at 1% and expressed in logs. "Info treatment" equals 1 if the participant was provided with information and 0 otherwise.

students, respectively. These estimates are indistinguishable ($p = 0.944$), and combining the controls implies that information lowers continuation by 3.7 percentage points from a control senior-high-school nonregistration rate of 9.4 percent. A randomization test designed for a small number of clusters gives $p = 0.077$, so the estimate is statistically distinguishable from zero at the 10 percent but not the 5 percent level.[26] An inverse-probability weighting adjustment for missing outcomes gives a similar estimate of 3.4 percentage points (s.e. 2.2 percentage points).[27,28]

Columns (3) and (4) interact information assignment with WTP for the information bundle. The negative interaction shows that information reduces continuation more strongly among students with higher WTP. The unconditional estimate is a 3.7 percentage point decline in continuation; column (3) implies a decline of about 9.5 percentage points at a WTP of N200. The pattern is robust to the additional moderators in column (4).

### 4.3 Evidence against a uniform response to information

The absolute intent-to-treat effect equals the gross share whose choice changes only under a uniform response. We implement the diagnostic in Section 2.1 using the summary priors

[26] We implement the randomization test proposed by Canay, Romano and Shaikh (2017). Because it requires treatment effects at the cluster level, we group adjacent schools and reduce the number of clusters from 36 to 18. Across 1,000 random pairings, the mean $p$-value is 0.077 (s.e. 0.033).

[27] We predict attrition using age, sex, mother's years of schooling, Christianity, average grades, and indicators for two-parent households and grade repetition.

[28] Two of the twenty schools assigned to the information treatment did not receive it, so Table 5 reports the effect of information actually received. The intent-to-treat effect of planned assignment, with the covariates of Table 3, is 3.4 percentage points (cluster-robust p = .09, randomization p = .15). The two noncompliant schools have no belief data, so within the belief sample planned and actual assignment coincide.

Table 5: Treatment effects and the value of information: continuing to senior high school

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Impure control | 0.038 | | | |
| | (0.028) | | | |
| Pure control | 0.037 | | | |
| | (0.020) | | | |
| Info treatment | | -0.037 | 0.031 | 0.003 |
| | | (0.020) | (0.033) | (0.061) |
| WTP/100 | | | 0.038 | 0.049 |
| | | | (0.007) | (0.011) |
| WTP/100 × Info treatment | | | -0.063 | -0.075 |
| | | | (0.014) | (0.017) |
| Constant | 0.869 | 0.906 | 0.863 | 0.891 |
| | (0.017) | (0.010) | (0.026) | (0.031) |
| Observations | 3473 | 3473 | 2279 | 2135 |

Standard errors in parentheses. Clustered at the school level.

Notes: The dependent variable equals 1 if a student registers for senior high school and 0 otherwise. The table reports linear probability models. Columns (3) and (4) use the treatment and impure-control groups (24 of the 36 schools), because WTP was elicited only in these groups. Column (4) includes indicators for female, two-parent households, suspension, and grade repetition; average grades and lottery investment; and their interactions with information assignment. Standard errors are clustered at the school level.

$Y_1$ and $Y_0$ defined in Section 3.6 and the age-30 priors for each track. Beliefs were elicited before information provision in the 18 treatment schools, the 6 impure-control schools, and 2 pure-control schools whose students also reported priors. The diagnostic does not require a correspondence between these summary measures and a particular signal. The $Y_1$ measure is balanced across the two groups.[29]

As a preliminary diagnostic, we examine WTP under the maintained hypothesis of a uniform response toward senior-high-school noncontinuation. For this diagnostic, we use the indicator for noncontinuation, $1 - D_i$, as the schooling outcome; students induced not to register are then compliers. Under monotonicity and random assignment, their characteristics can be recovered separately from treatment and control observations and should agree (e.g., Heckman and Pinto, 2018). Table 6 instead gives significantly higher implied complier WTP in the control group than in the treatment group, and several estimates exceed the N200 elicitation support. These failures cast doubt on the uniform-response restriction and show that the monotonicity-based mixture recovery is unstable.

[29] With school-level inference, neither the difference in means ($p = 0.41$) nor the difference in distributions ($p = 0.58$) is significant. Treated students do report a lower noncontinuation prior $Y_0$ (permutation $p = 0.01$ for the difference in distributions), which Section 4.4 addresses by rebalancing.

Table 6: Willingness to pay for information by compliers under monotonicity

| | In Control | In Treatment | p-value | Combined |
|---|---|---|---|---|
| Admission & Earnings | 345.80 | 132.71 | 0.01 | 239.26 |
| Admission rates | 292.33 | 132.21 | 0.02 | 212.27 |
| Earnings | 251.04 | 95.43 | 0.02 | 173.24 |

Notes: The table reports monotonicity-based estimates recovered separately from the control and treatment groups. WTP was elicited on a N0–N200 price list. Estimates outside that support indicate failure of the maintained mixture interpretation.

We formally test the sign restrictions in Proposition 1. Because information was assigned by school, we compute all $p$-values for these tests at the level of the 26 schools in the belief sample, either by permuting information assignment across schools or with procedures that cluster at the school level. A uniform response requires either $M^+ = 0$ or $M^- = 0$, so we test the implication of each hypothesis and reject a uniform response at a given level only when both are rejected at that level. The procedure of Kitagawa (2015), which compares the treatment and control subdensities over every interval of priors, never rejects $M^+ = 0$ (Appendix Table A.3).[30] With 26 schools, a test that searches over so many comparisons can have little power against violations confined to part of the prior distribution.

Remark 1 suggests conditioning on students whose priors predict the direction of their response. With gender- and track-specific signals, the decision model predicts this direction only for students whose priors lie on the same side of every signal (Remark 2). We call a student a pessimist if their age-30 earnings priors for all three tracks lie below the corresponding signals for their gender, and an optimist if all three lie above them (see Figure A.3). Under the model, information can only move pessimists toward continuation and optimists away from it. Pessimists make up 20 percent of the belief sample and optimists 37 percent, with similar shares in the treatment and control groups. The model only guides the choice of subsample. The test itself relies only on random assignment and consistency (Assumption 1), under which a uniform response restricts every subsample defined by baseline variables.

Table 7 reports the test of Mourifie and Wan (2017), which conditions on $Y_1$ through a cubic series and imposes the inequalities of Proposition 1 for both $D_i = 1$ and $D_i = 0$. Among pessimists and optimists, both hypotheses are rejected at the 10 percent level under either form of inference, with $p$-values of 0.099 for $M^+ = 0$ and 0.031 for $M^- = 0$ using the critical values of Chernozhukov, Lee and Rosen (2013), and 0.070 and 0.032 using randomization

[30] We thank Zhenting Sun for providing the test code.

inference. Under Assumption 1, this is evidence that the intervention induced some students to continue and others not to continue. In the full sample, only the permutation $p$-values fall below 0.10, consistent with the loss of power Remark 1 describes. Appendix Table A.4 reports other specifications of the test.

Table 7: Test of a uniform response using prior beliefs

| | $H_0\colon M^+ = 0$ | | $H_0\colon M^- = 0$ | |
|---|---|---|---|---|
| | Perm. | CLR | Perm. | CLR |
| All students ($N$ = 2,485) | 0.094 | 0.164 | 0.088 | 0.167 |
| Pessimists and optimists ($N$ = 1,401) | 0.070 | 0.099 | 0.032 | 0.031 |

Notes: $p$-values for the test of Mourifie and Wan (2017) of the implications of Proposition 1 under $M^+ = 0$ and under $M^- = 0$, imposing the inequalities for both $D_i = 1$ and $D_i = 0$ and conditioning on $Y_1$ through a cubic series. Perm. reassigns information across the 26 schools in the belief sample, keeping 18 treated (499 permutations); it is exact for the narrower hypothesis that information changes no student's choice. CLR uses the simulated critical values of Chernozhukov et al. (2013) with a school-clustered covariance matrix (1,000 draws). Pessimists (optimists) report age-30 earnings priors below (above) the corresponding signal for their gender in all three tracks. The test is unweighted because it conditions only on $Y_1$, which is balanced across groups; Appendix Table A.4 also reports rebalanced versions.

The common-signal model adds a prediction about where the two responses occur. Figure 4 shows the estimated subdensities of the continuation prior and continuation under each assignment. The treatment subdensity lies above the control subdensity at lower priors and below it at higher priors, a pattern consistent with Proposition 3. The two cross near 11.2 (about N73,000), which is within the N53,100–N77,849 range of signals provided. Signals, however, vary by gender and track, and $Y_1$ aggregates several earnings reports, so the location of the pooled crossing does not formally test the common-signal proposition (Remark 2). Under the conditions of Proposition 4, reporting error can move the crossing but cannot create additional sign changes.

Appendix B.1 reports one-sided tests split at the mean experimental signal. The main prior measure detects an outflow above that value but not below it. Because neither heterogeneous signals nor measurement error makes the mean signal the exact pooled cutoff, these are suggestive checks on the shape rather than formal validation of the common-signal model.

### 4.4 Bounds on flows and a decomposition of the intent-to-treat estimate

Section 2.2 shows that the positive and negative parts of the treatment-control subdensity difference give lower bounds on the flows into and out of continuation. Table 8 reports plug-in estimates. We first rebalance the treated sample to match the control distribution

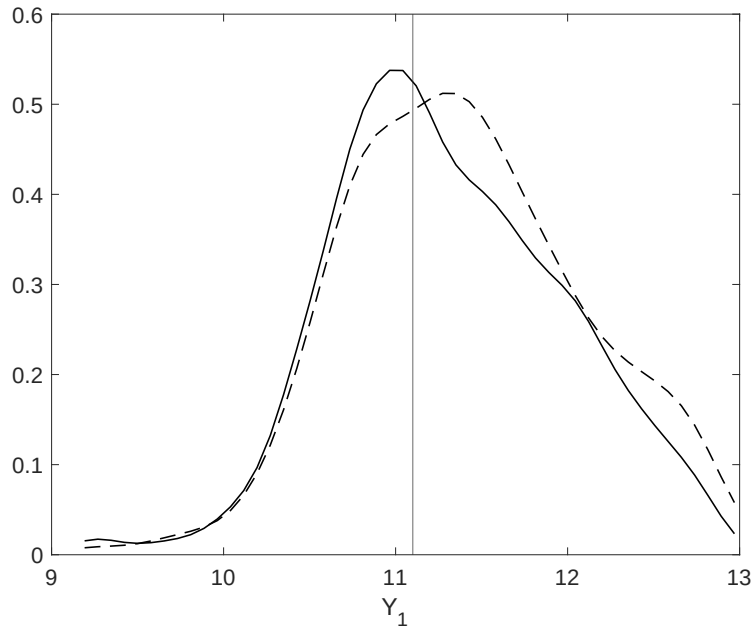



Figure 4: Estimated subdensities of prior beliefs and continuation by information assignment. Solid line for treatment; dotted line for control. Estimates are reweighted for assignment probabilities.

of priors and then estimate the subdensities by kernel methods.[31] The top panel uses $Y_1$ alone, while the bottom panel uses $(Y_0, Y_1)$.

Table 8: Estimated lower bounds on behavioral flows

| | One dimensional | | | | |
|---|---|---|---|---|---|
| | Switched in | Switched out | Out minus In | Out plus In | Fréchet UB |
| Mean | 0.009 | 0.058 | 0.049 | 0.067 | 0.212 |
| SE | 0.004 | 0.006 | 0.008 | 0.006 | 0.008 |
| | **Two dimensional** | | | | |
| | Switched in | Switched out | Out minus In | Out plus In | Fréchet UB |
| Mean | 0.049 | 0.097 | 0.049 | 0.146 | 0.212 |
| SE | 0.006 | 0.007 | 0.008 | 0.010 | 0.008 |

The "Out minus In" column is the net decline in continuation and equals the magnitude of the intent-to-treat effect in the rebalanced sample: 4.9 rather than 3.7 percentage points. It is the same in both panels because the net response does not depend on the priors used. The "Out plus In" column is an estimated lower bound on the gross response. It is 6.7 percent using $Y_1$ and 14.6 percent using $(Y_0, Y_1)$: 1.4 and 3.0 times the net response in the same sample, and 1.8 and 3.9 times the full-sample intent-to-treat estimate. The last column reports the Fréchet upper bound on the gross response; for a binary choice it is the largest value consistent with the treated and control choice shares alone.

The gap between the two panels is informative about the behavioral model. If every student received the same signal and reported their priors accurately, the continuation prior alone would determine the direction of each response (Corollary 1). Conditioning on the discontinuation prior as well could not uncover any additional offsetting flows, and the one- and two-dimensional population bounds would both equal the population flows. The larger two-dimensional estimates may instead reflect the gender- and track-specific signals of Remark 2, reporting error, or additional heterogeneity related to $Y_0$, in addition to sampling, smoothing, and rebalancing error. Because our experiment provides several signals, neither panel need equal the population flows, but both remain valid lower bounds.

[31] We bin beliefs and sample from the treatment group to match the control proportions. We retain samples in which the prior distributions do not differ significantly and estimate subdensities using Matlab's `mvksdensity`, Silverman's bandwidth rule, and an Epanechnikov kernel. Rebalancing adjusts for prior imbalance; it does not impose the common-signal or no-offsetting conditions.

### 4.5 Bounds on flows across educational choices

Proposition 2 gives lower bounds for several mutually exclusive educational paths. For each path, we estimate the joint subdensity of its earnings prior, the junior-high-school earnings prior, and the realized choice under each assignment.[32] We integrate the positive and negative parts of the subdensity difference to estimate lower bounds on the shares switching into and out of that path. We carry out this calculation separately for Arts, Commerce/Science, and dropout.[33] Since the choices are mutually exclusive, each student can enter at most one path under treatment.

These entry and exit bounds, together with the observed choice margins, restrict the transition matrix introduced in Table 1. We use a linear program to find the matrix with the smallest off-diagonal sum that satisfies these restrictions, and report this minimum-switching matrix in Table 9. Its last column is the control distribution of choices, its last row is the treatment distribution, and the off-diagonal sum is 17 percent. This sum estimates a lower bound on total switching. The individual cells, including the 0.07 Commerce/Science-to-dropout entry, are from the minimizing matrix; the population transitions are not point identified.

Table 9: Minimum-switching distribution of educational choices

| Control | Treated | | | |
|---|---|---|---|---|
| | Dropout | Arts | Commerce & Science | Total |
| Dropout | 0.02 | 0.02 | 0.04 | 0.08 |
| | [0.00,0.05] | [0.00,0.06] | [0.00,0.08] | [0.05,0.12] |
| Arts | 0.04 | 0.34 | 0.00 | 0.39 |
| | [0.01,0.08] | [0.31,0.38] | [0.00,0.01] | [0.35,0.42] |
| Commerce & Science | 0.07 | 0.00 | 0.47 | 0.53 |
| | [0.03,0.10] | [0.00,0.00] | [0.41,0.51] | [0.48,0.57] |
| Total | 0.13 | 0.36 | 0.51 | 1.00 |
| | [0.10,0.17] | [0.31,0.41] | [0.47,0.55] | |

Notes: The table reports the feasible transition matrix with the smallest off-diagonal sum subject to the observed choice margins and the estimated per-choice flow bounds. Brackets report bootstrap intervals from the original calculation. The individual transition probabilities are not point identified.

[32] For Arts, these objects are $f(Y_0, Y_{\text{Arts}}, \text{Arts} = 1 \mid Z = 0)$ and $f(Y_0, Y_{\text{Arts}}, \text{Arts} = 1 \mid Z = 1)$.

[33] We estimate Gaussian kernel subdensities over a $100 \times 100$ grid using Silverman's bandwidth rule and integrate numerically using the trapezoidal rule. Before estimation, we rebalance the data in $4 \times 4$ prior-belief bins. Confidence intervals use a clustered bootstrap.

# 5 Conclusions

Despite disparities in access to information about the returns to education and the long-term consequences of poor schooling choices, evidence on the effectiveness of educational information interventions is mixed. One interpretation is that information barriers are of secondary importance. Another is that a small average effect conceals offsetting responses, with some students induced to register for senior high school and others induced not to register.

We develop a diagnostic for this distinction. Under random assignment and consistency, a uniform response to information imposes the same joint-distribution restrictions used in instrument-validity tests, with elicited priors playing the role of outcomes. The positive and negative parts of the corresponding subdensity differences also give design-based lower bounds on the gross flows into and out of each choice. In the Ibadan experiment, we reject a uniform response at the 10 percent level among students whose priors lie on the same side of every signal provided by the intervention, the students for whom the model predicts the direction of response. The estimated lower bound based on priors for both continuation and noncontinuation is three times the net response in the rebalanced belief sample and nearly four times the full-sample intent-to-treat estimate. A minimum-switching construction that also incorporates curriculum-track choices gives an estimated lower bound of 17 percent.

These estimates measure changes in schooling choices, not welfare. Eliciting additional baseline beliefs can reveal offsetting responses that a larger sample alone would not detect.

## References


Ajayi, K.F., Friedman, W.H., Lucas, A.M., 2020. When information is not enough: Evidence from a centralized school choice system.

Allcott, H., Braghieri, L., Eichmeyer, S., Gentzkow, M., 2020. The Welfare Effects of Social Media. American Economic Review 110, 629–676.

Athey, S., 2002. Monotone comparative statics under uncertainty. Quarterly Journal of Economics 117, 187–223.

Balke, A., Pearl, J., 1997. Bounds on treatment effects from studies with imperfect compliance. Journal of the American Statistical Association 92, 1171–1176.

Balla-Elliott, D., 2026. Identifying causal effects in information provision experiments. The Review of Economics and Statistics Forthcoming.

Barr, A., Bird, K.A., Castleman, B.L., Skimmyhorn, W.L., 2022. Can information and advising affect postsecondary participation and attainment for non-traditional students? evidence from a large-scale experiment with the u.s. army.

Bergman, P., Denning, J.T., Manoli, D., 2019. Is Information Enough? The Effect of Information about Education Tax Benefits on Student Outcomes. Journal of Policy Analysis and Management 38, 706+.

Bettinger, E.P., Long, B.T., Oreopoulos, P., Sanbonmatsu, L., 2012. The Role of Application Assistance and Information in College Decisions: Results from the H & R Block Fafsa Experiment. Quarterly Journal of Economics 127, 1205–1242.

Blackwell, D., 1951. Comparison of experiments, in: In Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability. volume Berkeley, CA: University of California Press of *Annual Review of Economics*, pp. 93–102.

Bohm, P., Linden, J., Sonnegard, J., 1997. Eliciting reservation prices: Becker-DeGroot-Marschak mechanisms vs. markets. Economic Journal 107, 1079–1089.

Bursztyn, L., Gonzalez, A.L., Yanagizawa-Drott, D., 2020. Misperceived social norms: Women working outside the home in saudi arabia. American Economic Review 110, 2997–3029.

Busso, M., Dinkelman, T., Martinez, C.A., Romero, D., 2017. The effects of financial aid and returns information in selective and less selective schools: Experimental evidence from Chile. Labour Economics 45, 79–91.

Canay, I.A., Romano, J.P., Shaikh, A.M., 2017. Randomization tests under an approximate symmetry assumption. Econometrica 85, 1013–1030.

Cantoni, D., Yang, D.Y., Yuchtman, N., Zhang, Y.J., 2019. Protest as strategic games: Experimental evidence from hong kong's antiauthoritarian movement. Quarterly Journal of Economics 134, 1021–1077.

Chambers, C.P., Liu, C., Rehbeck, J., 2020. Costly information acquisition. Journal of Economic Theory 186.

Chernozhukov, V., Lee, S., Rosen, A.M., 2013. Intersection bounds: Estimation and inference. Econometrica 81, 667–737.

Coibion, O., Gorodnichenko, Y., Weber, M., 2022. Monetary policy communications and their effects on household inflation expectations. Journal of Political Economy 130, 1537–1584.

Dahl, C.M., Huber, M., Mellace, G., 2023. It is never too LATE: A new look at local average treatment effects with or without defiers. The Econometrics Journal 26, 378–404.

Frandsen, B., Lefgren, L., Leslie, E., 2023. Judging judge fixed effectst. American Economic Review 113, 253–277.

Fuster, A., Perez-Truglia, R., Wiederholt, M., Zafar, B., 2022. Expectations with endogenous information acquisition: An experimental investigation. Review of Economic and Statistics 104, 1059–1078.

Goux, D., Gurgand, M., Maurin, E., 2017. Adjusting Your Dreams? High School Plans and Dropout Behaviour. Economic Journal 127, 1025–1046.

Gurantz, O., Howell, J., Hurwitz, M., Larson, C., Pender, M., White, B., 2021. A National-Level Informational Experiment to Promote Enrollment in Selective Colleges. Journal of Policy Analysis and Management 40, 453+.

Haaland, I., Roth, C., Wohlfart, J., 2023. Designing Information Provision Experiments. Journal of Economic Literature 61, 3–40.

Hastings, J., Neilson, C.A., Zimmerman, S.D., 2015. The effects of earnings disclosure on college enrollment decisions.

Heckman, J.J., Honoré, B.E., 1990. The empirical content of the Roy model. Econometrica 58, 1121–1149.

Heckman, J.J., Pinto, R., 2018. Unordered Monotonicity. Econometrica 86, 1–35.

Heckman, J.J., Vytlacil, E., 2005. Structural equations, treatment effects, and econometric policy evaluation. Econometrica 73, 669–738.

Hirshleifer, J., 1971. Private and social value of information and reward to inventive activity. American Economic Review 61, 561–574.

Hjort, J., Moreira, D., Rao, G., Santini, J.F., 2021. How Research Affects Policy: Experimental Evidence from 2,150 Brazilian Municipalities. American Economic Review 111, 1442–1480.

Hoxby, C.M., Turner, S., 2015. What High-Achieving Low-Income Students Know About College. American Economic Review 105, 514–517.

Hyman, J., 2020. Can Light-Touch College-Going Interventions Make a Difference? Evidence from a Statewide Experiment in Michigan. Journal of Policy Analysis and Management 39, 159+.

Imbens, G., Angrist, J., 1994. Identification and Estimation of Local Average Treatment Effects. Econometrica 62, 467–475.

Jensen, R., 2010. The (perceived) returns to education and the demand for schooling. The Quarterly Journal of Economics 125, 515–548.

Kédagni, D., Mourifié, I., 2020. Generalized instrumental inequalities: testing the instrumental variable independence assumption. Biometrika 107, 661–675.

Kédagni, D., Wu, H., Cui, Y., 2025. Robust identification in randomized experiments with noncompliance. Working paper, March 21. arXiv:2408.03530.

Kitagawa, T., 2015. A test for instrument validity. Econometrica 83, 2043–2063.

Kline, P., Tartari, M., 2016. Bounding the labor supply responses to a randomized welfare experiment: A revealed preference approach. American Economic Review 106, 972–1014.

Lai, F., Sadoulet, E., de Janvry, A., 2009. The adverse effects of parents' school selection errors on academic achievement: Evidence from the Beijing open enrollment program. Economics of Education Review 28, 485–496.

de Lara, M., Gossner, O., 2020. Payoffs-beliefs duality and the value of information. SIAM Journal on Optmization 30, 464–489.

Loyalka, P., Liu, C., Song, Y., Yi, H., Huang, X., Wei, J., Zhang, L., Shi, Y., Chu, J., Rozelle, S., 2013. Can information and counseling help students from poor rural areas go to high school? Evidence from China. Journal of Comparative Economics 41, 1012–1025.

Lucas, A.M., Mbiti, I.M., 2012. The Determinants and Consequences of School Choice Errors in Kenya. American Economic Review 102, 283–288.

Marbach, M., Hangartner, D., 2020. Profiling Compliers and Noncompliers for Instrumental-Variable Analysis. Political Analysis 28, 435–444.

Mazar, N., Koszegi, B., Ariely, D., 2014. True Context-dependent Preferences? The Causes of Market-dependent Valuations. Journal of Behavioral Decision Making 27, 200–208.

Mosquera, R., Odunowo, M., McNamara, T., Guo, X., Petrie, R., 2020. The economic effects of Facebook. Experimental Economics 23, 575–602.

Mourifie, I., Wan, Y., 2017. Testing local average treatment effect assumptions. Review of Economic and Statistics 99, 305–313.

Noack, C., 2026. Sensitivity of LATE estimates to violations of the monotonicity assumption. Working paper, March 15.

Roy, A.D., 1951. Some thoughts on the distribution of earnings. Oxford Economic Papers 3, 135–146.

Stutern, 2018. Nigerian graduate report: 2018. https://www.kaggle.com/stutern/nigerian-gradaute-report-2018.

Sun, Z., 2023. Instrument validity for heterogeneous causal effects. Journal of Econometrics 237.

Vytlacil, E., 2002. Independence, monotonicity, and latent index models: An equivalence result. Econometrica 70, 331–341.

Wiswall, M., Zafar, B., 2015a. Determinants of college major choice: Identification using an information experiment. The Review of Economic Studies 82, 791–824.

Wiswall, M., Zafar, B., 2015b. How do college students respond to public information about earnings? Journal of Human Capital 9, 117–169.

**APPENDICES INTENDED FOR ONLINE PUBLICATION**

# A Proofs of theoretical results

*Proof of Proposition 1.* For $z \in \{0,1\}$, parts (ii) and then (i) of Assumption 1 give

$$\begin{aligned}\Pr(Y_i \in B,\, D_i = 1 \mid Z_i = z) &= \Pr(Y_i \in B,\, D_i(z) = 1 \mid Z_i = z) \\ &= \Pr(Y_i \in B,\, D_i(z) = 1).\end{aligned}$$

Decomposing each probability by the value of the other potential choice,

$$\begin{aligned}\Pr(Y_i \in B,\, D_i(1) = 1) &= \Pr(Y_i \in B,\, D_i(1) = 1,\, D_i(0) = 1) \\ &\quad + \Pr(Y_i \in B,\, D_i(1) = 1,\, D_i(0) = 0), \\ \Pr(Y_i \in B,\, D_i(0) = 1) &= \Pr(Y_i \in B,\, D_i(0) = 1,\, D_i(1) = 1) \\ &\quad + \Pr(Y_i \in B,\, D_i(0) = 1,\, D_i(1) = 0).\end{aligned}$$

Differencing the two displays cancels the common term $\Pr(Y_i \in B,\, D_i(1) = 1,\, D_i(0) = 1)$ and gives (2).

If $D_i(1) \geq D_i(0)$ for all students, then $\Pr(Y_i \in B,\, D_i(0) = 1,\, D_i(1) = 0) = 0$ for every $B$, so the left-hand side of (2) equals $\Pr(Y_i \in B,\, D_i(1) = 1,\, D_i(0) = 0) \geq 0$. If $D_i(1) \leq D_i(0)$ for all students, the first term on the right-hand side of (2) is zero for every $B$ and the left-hand side is weakly negative. Finally, if the left-hand side of (2) is strictly positive at some set $B$ and strictly negative at another set $B'$, then

$$\begin{aligned}&M^+ \geq \Pr(Y_i \in B,\, D_i(1) = 1,\, D_i(0) = 0) > 0 \\ \text{and} \quad &M^- \geq \Pr(Y_i \in B',\, D_i(0) = 1,\, D_i(1) = 0) > 0.\end{aligned}$$

□

*Proof of Proposition 2.* By Assumption 1 and the definitions of the subdensities, for every $k$ and almost every $y$,

$$\begin{aligned}p_k(y) &= m_k^+(y) + f_Y(y) \Pr(D_i(0) = k,\, D_i(1) = k \mid Y_i = y), \\ q_k(y) &= m_k^-(y) + f_Y(y) \Pr(D_i(0) = k,\, D_i(1) = k \mid Y_i = y).\end{aligned}$$

Subtracting the two equations gives

$$p_k(y) - q_k(y) = m_k^+(y) - m_k^-(y).$$

Because $m_k^+$ and $m_k^-$ are nonnegative,

$$\max\{p_k(y) - q_k(y), 0\} \le m_k^+(y) \quad \text{and} \quad \max\{q_k(y) - p_k(y), 0\} \le m_k^-(y),$$

and integrating over $y$ gives the two lower bounds, since $\int m_k^+(y)\,dy = M_k^+$ and $\int m_k^-(y)\,dy = M_k^-$. If the no-offsetting condition holds, then $m_k^+(y)m_k^-(y) = 0$ almost everywhere, so both pointwise inequalities, and hence both bounds, hold with equality.

It remains to establish joint sharpness. Fix any observable distribution consistent with Assumption 1. We construct a joint distribution by assigning $\min\{q_k(y), p_k(y)\}$ to students who choose $k$ under both assignments and matching the remaining mass across different choices. For $a \in \mathbb{R}$, write $[a]_+ = \max\{a, 0\}$ and define

$$r(y) = \sum_{k=0}^{K} [p_k(y) - q_k(y)]_+ = \sum_{k=0}^{K} [q_k(y) - p_k(y)]_+,$$

where the equality follows from $\sum_{k=0}^K p_k(y) = \sum_{k=0}^K q_k(y) = f_Y(y)$ for almost every $y$. For $k, \ell \in \{0, \dots, K\}$, define the joint subdensities

$$h_{k\ell}(y) = \begin{cases} \min\{q_k(y), p_k(y)\}, & k = \ell, \\ \dfrac{[q_k(y) - p_k(y)]_+ [p_\ell(y) - q_\ell(y)]_+}{r(y)}, & k \neq \ell,\ r(y) > 0, \\ 0, & k \neq \ell,\ r(y) = 0. \end{cases}$$

Since $[q_k(y) - p_k(y)]_+ [p_k(y) - q_k(y)]_+ = 0$, summing over rows and columns gives

$$\sum_{\ell=0}^{K} h_{k\ell}(y) = \min\{q_k(y), p_k(y)\} + [q_k(y) - p_k(y)]_+ = q_k(y),$$
$$\sum_{\ell=0}^{K} h_{\ell k}(y) = \min\{q_k(y), p_k(y)\} + [p_k(y) - q_k(y)]_+ = p_k(y).$$

These identities also hold when $r(y) = 0$, since then $p_k(y) = q_k(y)$ for every $k$. The functions $h_{k\ell}$ are measurable and nonnegative, and their sum is $f_Y$. They therefore define

a joint distribution of $(Y_i, D_i(0), D_i(1))$ by

$$\Pr(Y_i \in B,\ D_i(0) = k,\ D_i(1) = \ell) = \int_B h_{k\ell}(y)\, dy$$

for every Borel set $B$. Draw $Z_i$ independently of these variables with its observed assignment probability and set $D_i = D_i(Z_i)$. The resulting distribution satisfies Assumption 1 and reproduces the observable distribution of $(Y_i, D_i, Z_i)$.

Under this construction,

$$m_k^+(y) = \sum_{\ell \neq k} h_{\ell k}(y) = p_k(y) - h_{kk}(y) = [p_k(y) - q_k(y)]_+,$$
$$m_k^-(y) = \sum_{\ell \neq k} h_{k\ell}(y) = q_k(y) - h_{kk}(y) = [q_k(y) - p_k(y)]_+.$$

Integrating over $y$ shows that the same joint distribution attains both lower bounds for every $k$, establishing joint sharpness. □

*Proof of Proposition 3.* Fix $(\tilde{Y}_{i,0}(0), \tilde{Y}_{i,1}(0)) = (\tilde{y}_0, \tilde{y}_1)$. If $\tilde{y}_1 < S$, Assumption 2(i) implies

$$R_i(1) \geq R_i(0) = \tilde{y}_1 - \tilde{y}_0,$$

and, because $\nu$ is increasing, Assumption 2(ii) then gives $D_i(1) \geq D_i(0)$ for every value of $e_i$. Combined with Assumption 2(iii) and $D_i = D_i(Z_i)$, this yields

$$\Pr(D_i = 1 \mid \tilde{Y}_i(0) = (\tilde{y}_0, \tilde{y}_1), Z_i = 1) \geq \Pr(D_i = 1 \mid \tilde{Y}_i(0) = (\tilde{y}_0, \tilde{y}_1), Z_i = 0).$$

If $\tilde{y}_1 > S$, the same argument gives $R_i(1) \leq R_i(0)$ and reverses the inequality.

It remains to integrate out $\tilde{Y}_{i,0}(0)$. For $z \in \{0, 1\}$,

$$\Pr(D_i = 1 \mid \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = z) = \int \Pr(D_i = 1 \mid \tilde{Y}_{i,0}(0) = \tilde{y}_0, \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = z) \, dF_{\tilde{Y}_{i,0}(0) \mid \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = z}(\tilde{y}_0),$$

and Assumption 2(iii) implies $F_{\tilde{Y}_{i,0}(0) \mid \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = 1} = F_{\tilde{Y}_{i,0}(0) \mid \tilde{Y}_{i,1}(0) = \tilde{y}_1, Z_i = 0}$. Integrating the pointwise inequality against this common distribution gives the result. □

*Proof of Corollary 1.* Because $Y_{i,1} = \tilde{Y}_{i,1}(0)$ almost surely, conditioning on the reported prior also fixes the underlying continuation prior. Fix $y_1 < S$. As in the proof of Propo-

sition 3, Assumption 2(i) implies $R_i(1) \geq R_i(0)$ for every student with $Y_{i,1} = \tilde{Y}_{i,1}(0) = y_1$, whatever the value of $\tilde{Y}_{i,0}(0)$, and Assumption 2(ii) then gives $D_i(1) \geq D_i(0)$ for every value of $e_i$. Hence

$$\Pr(D_i(0) = 1,\, D_i(1) = 0 \mid Y_{i,1} = y_1) = 0.$$

For $y_1 > S$, the same argument gives $\Pr(D_i(1) = 1,\, D_i(0) = 0 \mid Y_{i,1} = y_1) = 0$. The product of the two conditional probabilities is therefore zero for every $y_1 \neq S$, and $\Pr(Y_{i,1} = S) = 0$ because $Y_{i,1}$ has a density. This is the no-offsetting condition of Proposition 2 applied with $Y_i = Y_{i,1}$ and $k = 1$, so both bounds hold with equality. □

*Proof of Proposition 4.* Define

$$m(\tilde{y}) \coloneqq \Pr(D_i(1) = 1 \mid \tilde{Y}(0) = \tilde{y}) - \Pr(D_i(0) = 1 \mid \tilde{Y}(0) = \tilde{y}),$$

so that $m(\tilde{y}) \geq 0$ for $\tilde{y} < S$ and $m(\tilde{y}) \leq 0$ for $\tilde{y} > S$ by Proposition 3. Because $D_i(z)$ is a function of $(\tilde{Y}_i(z), e_i)$, Assumption 3(i) implies $\Pr(D_i(z) = 1 \mid Y, \tilde{Y}(0)) = \Pr(D_i(z) = 1 \mid \tilde{Y}(0))$ and, together with Assumption 2(iii), that $Z_i \perp\!\!\!\perp (Y, D_i(0), D_i(1))$. Using $D_i = D_i(Z_i)$, for $z \in \{0, 1\}$ and any $t$,

$$\begin{aligned}\Pr(D_i = 1,\, Y \leq t \mid Z_i = z) &= \mathbb{E}[\mathbf{1}\{Y \leq t\} D_i(z) \mid Z_i = z] = \mathbb{E}[\mathbf{1}\{Y \leq t\} D_i(z)] \\ &= \mathbb{E}[\mathbf{1}\{Y \leq t\} \Pr(D_i(z) = 1 \mid \tilde{Y}(0))].\end{aligned}$$

Differencing over $z$ and differentiating with respect to $t$,

$$Q(y) = f_Y(y)\, \mathbb{E}[m(\tilde{Y}(0)) \mid Y = y].$$

If the support of $\tilde{Y}(0)$ lies on one side of $S$, $\mathbb{E}[m(\tilde{Y}(0)) \mid Y = y]$ has a constant sign and the result is immediate. Otherwise, fix two reports $y_H > y_L$ and define

$$r(\tilde{y}) = \frac{f_{\tilde{Y}(0)|Y}(\tilde{y} \mid y_H)}{f_{\tilde{Y}(0)|Y}(\tilde{y} \mid y_L)},$$

which is weakly increasing by Assumption 3(ii). Let $c = r(S) \geq 0$, with a one-sided limit used if necessary. For $\tilde{y} < S$, $m(\tilde{y}) \geq 0$ and $r(\tilde{y}) \leq c$; for $\tilde{y} > S$, $m(\tilde{y}) \leq 0$ and $r(\tilde{y}) \geq c$. In either case $m(\tilde{y}) r(\tilde{y}) \leq c\, m(\tilde{y})$ for almost every $\tilde{y}$, so

$$\begin{aligned}\mathbb{E}[m(\tilde{Y}(0)) \mid Y = y_H] &= \int m(\tilde{y}) r(\tilde{y}) f_{\tilde{Y}(0)|Y}(\tilde{y} \mid y_L)\, d\tilde{y} \\ &\leq c \int m(\tilde{y}) f_{\tilde{Y}(0)|Y}(\tilde{y} \mid y_L)\, d\tilde{y} = c\, \mathbb{E}[m(\tilde{Y}(0)) \mid Y = y_L].\end{aligned}$$

Because $c \geq 0$, $\mathbb{E}[m(\tilde{Y}(0)) \mid Y = y_L] \leq 0$ implies $\mathbb{E}[m(\tilde{Y}(0)) \mid Y = y_H] \leq 0$: as $y$ increases, $\mathbb{E}[m(\tilde{Y}(0)) \mid Y = y]$ changes sign at most once, from positive to negative. Multiplying by the nonnegative density $f_Y(y)$ preserves this property, which is the stated single crossing of $Q$. □

# B ADDITIONAL MATERIAL

Table A.1: (Log) earnings beliefs by career track and age

| Variable | (1) Info (1) | (2) Imp. Cntrl (2) | (3) (2)-(1) |
|---|---|---|---|
| Junior HS at 21 | 9.170 | 9.405 | 0.235 |
| | (1.289) | (1.289) | (0.153) |
| Junior HS at 30 | 9.652 | 9.847 | 0.195 |
| | (1.187) | (1.160) | (0.117) |
| Junior HS at 50 | 10.014 | 10.145 | 0.131 |
| | (1.171) | (1.158) | (0.098) |
| Senior HS at 21 | 9.822 | 10.054 | 0.232 |
| | (1.062) | (1.064) | (0.119) |
| Senior HS at 30 | 10.202 | 10.382 | 0.180 |
| | (0.957) | (0.969) | (0.079) |
| Senior HS at 50 | 10.515 | 10.661 | 0.146 |
| | (0.988) | (1.035) | (0.103) |
| Arts at 21 | 10.702 | 10.794 | 0.092 |
| | (0.775) | (0.843) | (0.121) |
| Arts at 30 | 11.097 | 11.184 | 0.087 |
| | (0.826) | (0.911) | (0.140) |
| Arts at 50 | 11.511 | 11.522 | 0.011 |
| | (0.877) | (0.976) | (0.145) |
| Commerce at 21 | 10.891 | 10.804 | -0.087 |
| | (0.842) | (0.884) | (0.099) |
| Commerce at 30 | 11.229 | 11.229 | -0.000 |
| | (0.877) | (0.911) | (0.146) |
| Commerce at 50 | 11.606 | 11.560 | -0.046 |
| | (0.892) | (0.961) | (0.150) |
| Science at 21 | 11.036 | 11.105 | 0.068 |
| | (0.975) | (1.035) | (0.148) |
| Science at 30 | 11.378 | 11.397 | 0.019 |
| | (0.960) | (1.064) | (0.156) |
| Science at 50 | 11.671 | 11.675 | 0.004 |
| | (1.046) | (1.136) | (0.195) |
| Observations | 1,925 | 658 | 2,583 |

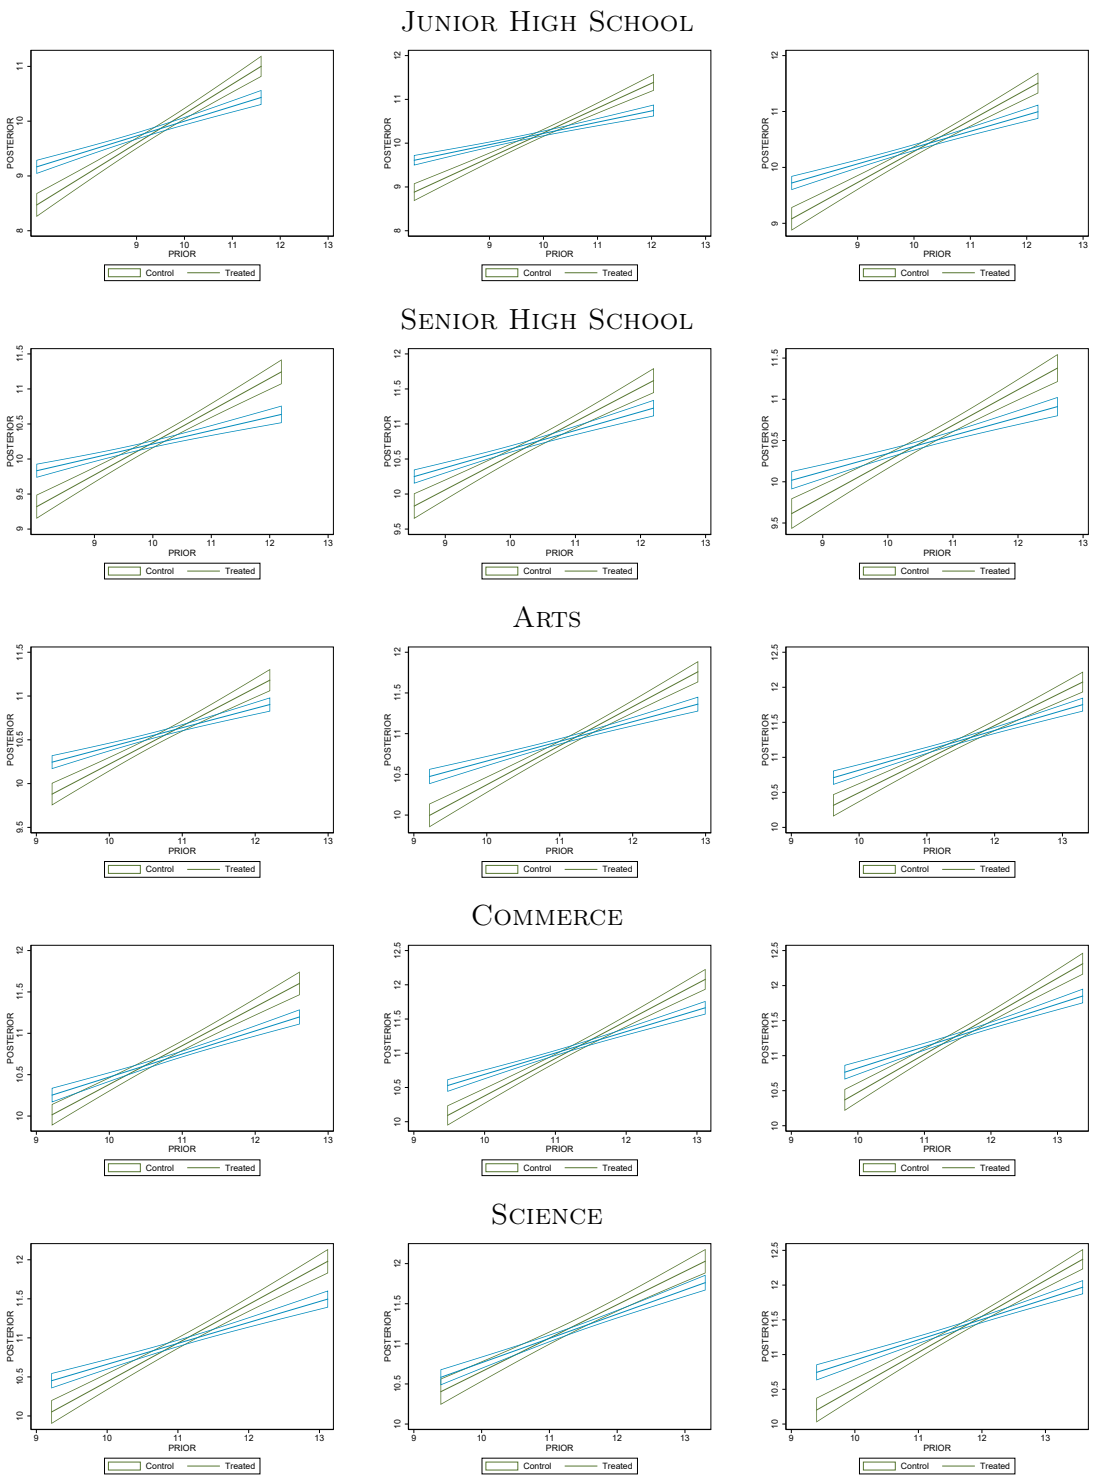



Figure A.1: Belief updating by treatment

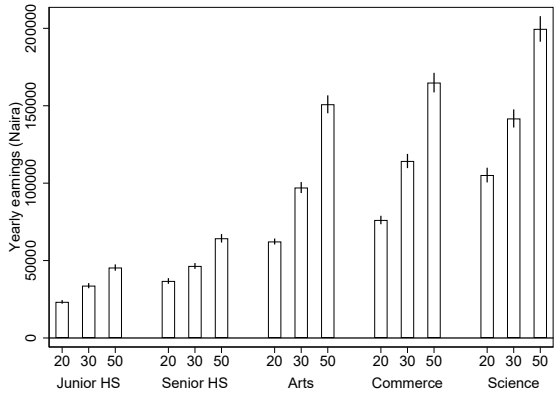



Figure A.2: Prior beliefs by educational choice and age

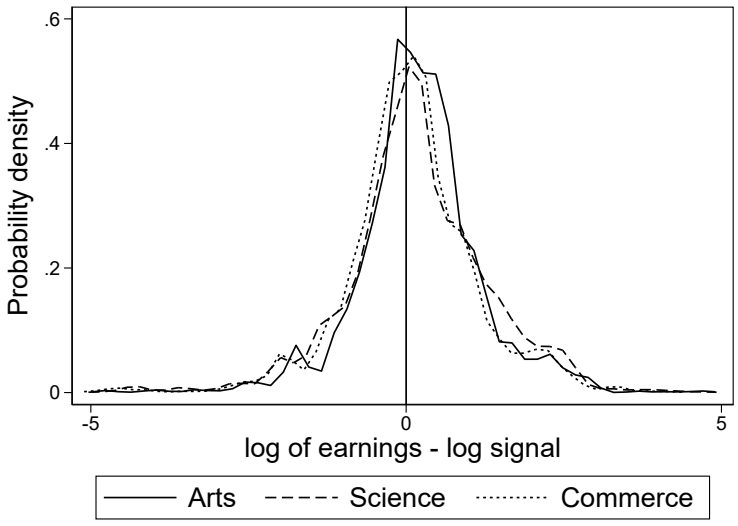



Figure A.3: Distribution of earnings beliefs at 30 relative to the corresponding signals

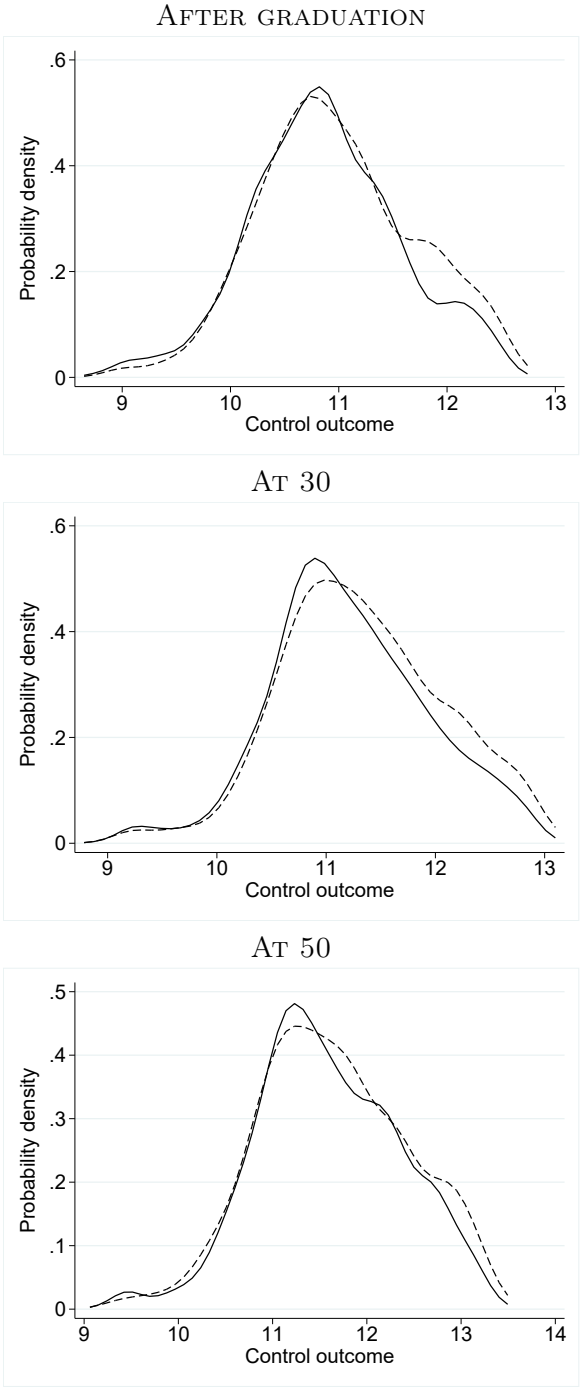



Figure A.4: Estimated subdensities of age-specific prior beliefs and continuation by information assignment. Solid lines for treatment; dotted lines for control. Estimates are reweighted for assignment probabilities.

Table A.2: Threshold probability updating

| | Pr(Salary at 30 ≥ 50,000) | | | Pr(Salary at 30 ≥ 100,000) | | |
|---|---|---|---|---|---|---|
| | Arts | Comm | Science | Arts | Comm | Science |
| (signal − prior) | | | | | | |
| × Info treatment | 0.146 | 0.083 | 0.065 | 0.180 | 0.152 | 0.205 |
| | (0.049) | (0.048) | (0.047) | (0.049) | (0.050) | (0.046) |
| Observations | 2133 | 2115 | 2121 | 2115 | 2108 | 2105 |
| Adj R2 | 0.324 | 0.345 | 0.369 | 0.391 | 0.364 | 0.334 |

Notes: The first column for Arts refers to the answer to the question: "What is the percentage chance that if you were working full time you would earn at least N50,000 monthly if you graduated from the arts track at university and were 30 years old?" The labels of the other columns follow the same pattern. The dependent variable is the difference between the belief elicited the second time and the belief elicited the first time. "Info treatment" equals 1 if the participant was provided with information and 0 otherwise.

Table A.3: Test of a uniform response using baseline beliefs (Kitagawa, 2015)

| | | $H_0\colon M^+ = 0$ | | $H_0\colon M^- = 0$ | |
|---|---|---|---|---|---|
| | $N$ | Perm. | School boot. | Perm. | School boot. |
| All students | | | | | |
| $Y_1$ | 2,485 | 0.942 | 0.804 | 0.345 | 0.363 |
| $(Y_0, Y_1)$ | 2,473 | 0.156 | 0.208 | 0.060 | 0.060 |
| Pessimists and optimists | | | | | |
| $Y_1$ | 1,401 | 0.938 | 0.792 | 0.964 | 0.896 |
| $(Y_0, Y_1)$ | 1,397 | 0.297 | 0.441 | 0.178 | 0.180 |

Notes: $p$-values for the test of Kitagawa (2015) of the implications of Proposition 1 for continuation ($D_i = 1$) under $M^+ = 0$ and under $M^- = 0$, with trimming constant $\xi = 0.07$. Perm. reassigns information across schools (499 permutations); it is exact for the narrower hypothesis that information changes no student's choice. School boot. is a recentered bootstrap that resamples schools within each group (499 draws). The two-dimensional rows evaluate rectangles on a grid of fifty equally spaced points for each prior. Pessimists and optimists are defined in Section 4.3.

Table A.4: Test of a uniform response using baseline beliefs (Mourifie and Wan, 2017)

| | | | Unweighted | | Rebalanced | |
|---|---|---|---|---|---|---|
| Basis | Restrictions | $H_0$ | CLR | Perm. | CLR | Perm. |
| All students | | | | | | |
| Linear in $Y_1$ | 1 | $M^+ = 0$ | 0.529 | 0.535 | 0.722 | 0.994 |
| Linear in $Y_1$ | 1 | $M^- = 0$ | 0.432 | 0.347 | 0.639 | 0.810 |
| Linear in $Y_1$ | 2 | $M^+ = 0$ | 0.807 | 0.735 | 0.958 | 0.996 |
| Linear in $Y_1$ | 2 | $M^- = 0$ | 0.424 | 0.144 | 0.521 | 0.218 |
| Cubic in $Y_1$ | 1 | $M^+ = 0$ | 0.645 | 0.866 | 0.774 | 1.000 |
| Cubic in $Y_1$ | 1 | $M^- = 0$ | 0.411 | 0.315 | 0.618 | 0.735 |
| Cubic in $Y_1$ | 2 | $M^+ = 0$ | 0.164 | 0.094 | 0.211 | 0.124 |
| Cubic in $Y_1$ | 2 | $M^- = 0$ | 0.167 | 0.088 | 0.038 | 0.034 |
| Linear in $(Y_0, Y_1)$ | 1 | $M^+ = 0$ | 0.222 | 0.144 | 0.809 | 0.994 |
| Linear in $(Y_0, Y_1)$ | 1 | $M^- = 0$ | 0.485 | 0.515 | 0.755 | 0.974 |
| Linear in $(Y_0, Y_1)$ | 2 | $M^+ = 0$ | 0.435 | 0.353 | 0.985 | 0.998 |
| Linear in $(Y_0, Y_1)$ | 2 | $M^- = 0$ | 0.394 | 0.295 | 0.635 | 0.575 |
| Quadratic in $(Y_0, Y_1)$ | 1 | $M^+ = 0$ | 0.407 | 0.447 | 0.877 | 1.000 |
| Quadratic in $(Y_0, Y_1)$ | 1 | $M^- = 0$ | 0.507 | 0.585 | 0.818 | 0.990 |
| Quadratic in $(Y_0, Y_1)$ | 2 | $M^+ = 0$ | 0.676 | 0.749 | 0.797 | 0.856 |
| Quadratic in $(Y_0, Y_1)$ | 2 | $M^- = 0$ | 0.047 | 0.054 | 0.254 | 0.236 |
| Pessimists and optimists | | | | | | |
| Linear in $Y_1$ | 1 | $M^+ = 0$ | 0.657 | 0.846 | 0.645 | 0.780 |
| Linear in $Y_1$ | 1 | $M^- = 0$ | 0.612 | 0.699 | 0.558 | 0.573 |
| Linear in $Y_1$ | 2 | $M^+ = 0$ | 0.909 | 0.936 | 0.908 | 0.906 |
| Linear in $Y_1$ | 2 | $M^- = 0$ | 0.645 | 0.427 | 0.655 | 0.435 |
| Cubic in $Y_1$ | 1 | $M^+ = 0$ | 0.744 | 0.988 | 0.701 | 0.942 |
| Cubic in $Y_1$ | 1 | $M^- = 0$ | 0.491 | 0.431 | 0.614 | 0.731 |
| Cubic in $Y_1$ | 2 | $M^+ = 0$ | 0.099 | 0.070 | 0.125 | 0.092 |
| Cubic in $Y_1$ | 2 | $M^- = 0$ | 0.031 | 0.032 | 0.021 | 0.026 |
| Linear in $(Y_0, Y_1)$ | 1 | $M^+ = 0$ | 0.353 | 0.277 | 0.783 | 0.966 |
| Linear in $(Y_0, Y_1)$ | 1 | $M^- = 0$ | 0.584 | 0.641 | 0.729 | 0.906 |
| Linear in $(Y_0, Y_1)$ | 2 | $M^+ = 0$ | 0.614 | 0.553 | 0.975 | 0.994 |
| Linear in $(Y_0, Y_1)$ | 2 | $M^- = 0$ | 0.793 | 0.798 | 0.697 | 0.647 |
| Quadratic in $(Y_0, Y_1)$ | 1 | $M^+ = 0$ | 0.469 | 0.513 | 0.844 | 0.988 |
| Quadratic in $(Y_0, Y_1)$ | 1 | $M^- = 0$ | 0.621 | 0.725 | 0.822 | 0.986 |
| Quadratic in $(Y_0, Y_1)$ | 2 | $M^+ = 0$ | 0.743 | 0.810 | 0.936 | 0.980 |
| Quadratic in $(Y_0, Y_1)$ | 2 | $M^- = 0$ | 0.335 | 0.345 | 0.586 | 0.619 |

Notes: $p$-values for the test of Mourifie and Wan (2017) of the implications of Proposition 1 under $M^+ = 0$ and under $M^- = 0$. With one restriction the test imposes only the inequality for continuation ($D_i = 1$); with two it adds the inequality for noncontinuation ($D_i = 0$). CLR uses the simulated critical values of Chernozhukov et al. (2013) with a school-clustered covariance matrix (1,000 draws). Perm. reassigns information across schools (499 permutations); it is exact for the narrower hypothesis that information changes no student's choice. Rebalanced estimates reweight treated students to match the control distribution of $(Y_0, Y_1)$ over a $4 \times 4$ grid of bins. Pessimists and optimists are defined in Section 4.3.

## B.1 Split-sample tests of the common-signal model

This appendix reports one-sided tests related to the density comparison in Section 4.3. Signals in the experiment differ across students and $Y_1$ aggregates several earnings reports, so a split at the mean signal is not a formal test of the common-signal proposition. We use the split only to describe where the treatment and control densities differ. Let $Z = 1$ denote information assignment and $D = 1$ continuation. For any interval $B$ in the support of $Y_1$, define

$$\phi(B) \;=\; \Pr(Y_1 \in B,\, D = 1 \mid Z = 0) \;-\; \Pr(Y_1 \in B,\, D = 1 \mid Z = 1).$$

The test statistic is $S = \sup_B \left[\sqrt{T_N}\,\phi(B)/\sigma(B)\right]$, where $\sigma(B)$ is a studentizing term and $T_N$ a sample-size adjustment. The null hypothesis being tested is

$$H_0 :\; \Pr(Y_1 \in B,\, D = 1 \mid Z = 0) \;\leq\; \Pr(Y_1 \in B,\, D = 1 \mid Z = 1) \quad \text{for all } B,$$

that is, that there is no set of prior beliefs on which information lowers the joint probability of holding those beliefs and continuing education. A rejection means that, for some range of priors, students who received information are less likely to continue education than students who did not. We compute $p$-values at the school level, both by permuting information assignment across schools and by a bootstrap that resamples schools within each group.

With a common signal and accurate reporting, the null holds below the signal and its reverse holds above it (Proposition 3). We split at the mean experimental signal to describe whether the observed outflow is concentrated among lower or higher priors.[34]

Table A.5 reports $p$-values for the aggregate continuation prior with $\xi = 0.07$. With the school bootstrap, the full-sample test does not reject ($p = 0.36$). Below the mean signal the test does not reject either ($p = 0.84$), while above it the test rejects at the 5 percent level ($p = 0.03$). Permutation $p$-values are similar. The detected outflow is therefore concentrated among higher priors.

[34] A rejection of $H_0$ above the split documents an outflow; it does not directly test the reverse ordering over every interval.

Table A.5: Split-sample tests of the common-signal model, main specification

| | $N$ | Perm. | School boot. |
|---|---|---|---|
| Full sample | 2,485 | 0.345 | 0.363 |
| $Y_1 < 11.1$ | 967 | 0.962 | 0.840 |
| $Y_1 > 11.1$ | 1,518 | 0.032 | 0.034 |

Notes: $p$-values for the one-sided test of Kitagawa (2015) described in the text, with trimming constant $\xi = 0.07$. Perm. reassigns information across schools (499 permutations). School boot. is a recentered bootstrap that resamples schools within each group (499 draws). $Y_1$ is the log of average earnings beliefs if continuing education. The sample is split at the arithmetic mean of the signals provided ($\approx 11.1$ in logs).

## B.2 Monotonicity-based response-type characteristics

Table A.6 reports characteristics recovered under a uniform response toward senior-high-school noncontinuation, using the corresponding monotonicity-based mixture formulas. The implied complier group accounts for about 4% of the sample, so we use 5,000 bootstrap draws for inference. The last column reports the share of draws in which the complier mean exceeds the noncomplier mean. Several differences are large, but the recovered mean of the binary "Monotonic" variable is 1.07, outside its $[0, 1]$ support. Together with the WTP support violations in Table 6, this shows that the recovered characteristics are unstable. We use the table to check the maintained uniform-response model; we do not interpret these characteristics literally.

Table A.6: Characteristics of response types

| | All | Always-Taker | Never-Taker | Compliers | Pr(Complier ≥ Non-Complier) |
|---|---|---|---|---|---|
| Percent of population | 1.00 | 0.09 | 0.87 | 0.04 | |
| | Mean of each characteristic | | | | |
| Age | 14.03 | 14.27 | 13.99 | 14.51 | 0.82 |
| Female | 0.47 | 0.52 | 0.45 | 0.83 | 0.89 |
| Mother ed. (years) | 12.54 | 13.04 | 12.51 | 11.81 | 0.38 |
| Christian | 0.38 | 0.35 | 0.36 | 0.87 | 0.98 |
| Siblings (No.) | 3.92 | 4.03 | 3.89 | 4.41 | 0.58 |
| Two-parent HH | 0.78 | 0.78 | 0.79 | 0.49 | 0.09 |
| Grades | 0.62 | 0.62 | 0.62 | 0.56 | 0.16 |
| Repeated a grade | 0.10 | 0.11 | 0.09 | 0.41 | 0.99 |
| Discount factor | 0.50 | 0.48 | 0.51 | 0.34 | 0.10 |
| Lottery investment | 40.41 | 38.07 | 41.16 | 27.16 | 0.11 |
| Monotonic | 0.69 | 0.60 | 0.69 | 1.07 | 0.95 |
| WTP admissions | 108.31 | 105.36 | 105.73 | 167.91 | 0.99 |
| WTP salaries | 103.83 | 106.47 | 102.16 | 137.77 | 0.93 |
| WTP both | 105.60 | 101.34 | 102.44 | 192.75 | 0.99 |

Notes: This analysis is done at the student level. Row 1 shows the share of each compliance group in the sample. The shares of each compliance group are slightly different for the data on WTP since this was collected only for the non pure control conditions. The remaining rows show the means of each student or school characteristic across the different subgroups. The last column shows the probability that the mean characteristic of the complier group is larger than the mean characteristic of non-compliers. P-values are calculated using 5000 bootstrap samples. The mean characteristic uses the approach suggested by Marbach and Hangartner (2020).

## B.3 Willingness to pay for information

Our analysis of WTP builds on the classical literature on the value of information (Blackwell, 1951; Hirshleifer, 1971; de Lara and Gossner, 2020) and relates valuations to heterogeneity in behavioral responses.

### B.3.1 A model of willingness to pay

We present a simple descriptive model of willingness to pay for information. In this model, students who are close to indifferent between continuing and discontinuing education may be willing to pay more for information, because a small change in their expected returns can reverse their preferred choice. In Appendix B.3.2, we informally test this implication in our experiment by examining how responses to information vary with WTP elicited before the intervention.

We build on the decision model in Section 2.3. Suppressing the student index, let $U$ denote the expected gain from continuing education before information, which may vary across students. A student with idiosyncratic preferences $e$ continues when $U + e \geq 0$ and is indifferent between the two choices when $U + e = 0$.

Students value information before observing the signal, taking $U$ and $e$ as given. Write $U(s) = U + \eta(s)$ for the expected gain from continuing education after receiving signal $s$ and assume $\mathbb{E}[\eta(s) \mid U, e] = 0$. The expectation is over signals under the student's prior, so the posterior expected gain averages to $U$.[35] The distribution of $\eta(s)$ may otherwise depend on both $U$ and $e$.

Assume that paying for information reduces utility by the same amount under either schooling choice. A student's willingness to pay for information is then

$$w(U, e) := \mathbb{E}[\max\{0, U + e + \eta(s)\} \mid U, e] - \max\{0, U + e\}.$$

This is the difference between expected utility with and without information, conditional on $U$ and $e$, which are known to the student. Convexity of $\max\{0, x\}$ and the conditional mean restriction imply $w(U, e) \geq 0$ by Jensen's inequality.

[35] Bayesian posterior beliefs average to the prior (Chambers, Liu and Rehbeck, 2020). Here we impose the corresponding restriction on expected gains, in addition to Assumption 2(i).

The conditional mean restriction also allows us to write

$$w(U,e) = \begin{cases} \mathbb{E}[\max\{0, U+e+\eta(s)\} \mid U, e], & U+e \leq 0, \\ \mathbb{E}[\max\{0, -U-e-\eta(s)\} \mid U, e], & U+e > 0. \end{cases}$$

If $U+e<0$, information is valuable when it makes continuation preferable; if $U+e>0$, it is valuable when it makes discontinuation preferable. Holding $e$ and the distribution of utility revisions fixed, individual WTP increases weakly as $U$ rises toward $-e$ from below and decreases weakly as $U$ rises above $-e$. It is therefore greatest at indifference, $U+e=0$. If no signal changes the preferred choice, information has zero value (de Lara and Gossner, 2020; Balla-Elliott, 2026).

To see this more formally, hold $e$ and the conditional distribution of $\eta(s)$ fixed as $U$ varies.[36] The function $w(U,e)$ is then Lipschitz in $U$, so for Lebesgue-almost every $U$ its derivative is

$$\frac{\partial w(U,e)}{\partial U} = \begin{cases} \Pr(U+e+\eta(s)>0 \mid U,e) \geq 0, & U+e<0, \\ -\Pr(U+e+\eta(s)<0 \mid U,e) \leq 0, & U+e>0. \end{cases}$$

The inequalities are strict when information has a positive probability of changing the preferred choice. These derivative signs need not hold if the conditional distribution of revisions changes with $U$.

### B.3.2 Willingness to pay by response group

WTP predicts which students respond. The treatment effect becomes more negative for continuation as WTP rises (Table 5). This pattern is consistent with the near-indifference logic in Section B.3.1.

A stronger exercise assigns a valuation to each response group. This calculation interprets the positive and negative subdensity differences as the full flows within the relevant prior regions and is therefore more model-dependent than the lower-bound results. Table A.7 compares continuers across assignments within the prior regions associated with each response. Under no offsetting, the treatment-control difference in the inflow region isolates students switched into continuation, while the control-treatment difference in the outflow region isolates students switched out. The law of total probability then gives the implied WTP in columns (3) and (6).

[36] A sufficient condition for the distribution of revisions to remain fixed as $U$ varies is $\eta(s) \perp\!\!\!\perp U \mid e$.

Table A.7: WTP BY PRIOR BELIEFS, ASSIGNMENT, AND RESPONSE GROUP

| | AC | AC+In | Switched-In | AC+Out | AC | Switched-Out | *p*-values | | |
|---|---|---|---|---|---|---|---|---|---|
| | (1) $Z=0$ | (2) $Z=1$ | (3) | (4) $Z=0$ | (5) $Z=1$ | (6) | (7) $H_0$:(1)=(2) | (8) $H_0$:(4)=(5) | (9) $H_0$:(3)=(6) |
| ALL INFORMATION: | | | | | | | | | |
| Mean | 112.4 | 106.0 | -97.9 | 113.5 | 101.8 | 212.7 | 0.083 | $< 0.001$ | $< 0.001$ |
| SD | 1.7 | 3.1 | 48.4 | 0.9 | 2.1 | 22.0 | | | |

Notes: AC denotes always-continuers, students who continue under both assignments. Columns (1)–(3) use students whose priors lie in the region where information moves students into continuation. Column (1) is the mean WTP of control continuers (always-continuers), column (2) that of treated continuers (always-continuers and students switched in), and column (3) the implied mean WTP of students switched in. Columns (4)–(6) make the corresponding comparison in the region where information moves students out of continuation. In this region, control continuers are always-continuers or students switched out, while treated continuers are always-continuers. WTP was elicited on a N0–N200 price list, and this range is not imposed as a boundary condition on the implied values.

Students switched out of continuation have an implied WTP of N213, and students switched in have an implied WTP of −N98. Neither value can be the mean of responses to a price list that runs from N0 to N200. Although the difference between the two groups is statistically significant, values outside the elicitation range suggest that differencing mixtures does not reliably recover group valuations. For this reason, we do not interpret these estimates as literal valuations, nor as evidence that students avoid information or are myopic.

# C Willingness to pay for information - Instructions

Let's play the following game! From time to time, we will stop during our explanation and allow you to ask questions. The purpose of this game is to help us understand how much value you place on certain information. Remember, if any part of the game makes you feel uncomfortable, you can talk to your school counselor or principal about it. They will be able to help you.

Consider that I want to sell information to you. The information can help you to make better decisions about choosing between science, arts or commercial class in SS1. This type of information is important because it can affect what you become in future. The information tells you the chances of getting admitted into higher institutions based on the type of class you choose or information on average salary for different professions and chances of working full time. For example, you will know the percentage of boys and girls that apply to study Arts, commercial and science courses and what percentage of them get admitted. Imagine you have N200, I would like to know if you will be willing to exchange the money for information. I will offer an amount of money as shown in the table below.

You will play three versions of the game, but only one will be used to pay you. Once you have made the decision for each of the three rounds, we will choose a number from one to three by randomly choosing from numbered balls in a bag. One of your classmates will be the one to pick the ball. The number on the ball chosen will determine which of one of the three versions of the game will be chosen to pay you. Next, I will present the class with 10 cards in a bag which represents prices drawn from N0 to more than N200, and someone in your class will be asked to pick one (the person will not know which card represents what price and I also do not know). The price on the card chosen will be used to determine if you get the information or not. If the price that is drawn from the bag is less than what you select as your value for the information, you will pay the drawn price and receive the information. If, however, the drawn price is strictly greater than what you choose as your valuation, then you do not get the information but keep your money. Think carefully about each decision.

Please listen carefully to the following example of this game: Dele is willing to buy the information at N125 and no more. So, he chooses "yes" for prices N0 - N125 (rows A-F) and chooses "no" for prices 150 and above (options G through I). We present him with a bag that has cards drawn from N0 to more than N200, he puts his hand in the bag and chooses price N0. Since the price N0, is always less than any amount he could have chosen, he will receive the information and keep his N200. And if N250 is chosen, he will not get

the information since it is larger than any amount he could have paid.

| **Do you accept the price?** | Yes | No |
|---|---|---|
| A: Price: N0 (means you pay: N0)<br>You receive the information AND keep a payment of N200 | X | |
| B: Price: N25 (means you pay: N25)<br>You receive the information AND keep a payment of N175 | X | |
| C: Price: N50 (means you pay: N50)<br>You receive the information AND keep a payment of N150 | X | |
| D: Price: N75 (means you pay: N75)<br>You receive the information AND keep a payment of N125 | X | |
| E: Price: N100 (means you pay: N100)<br>You receive the information AND keep a payment of N100 | X | |
| F: Price: N125 (means you pay: N125)<br>You receive the information AND keep a payment of N75 | X | |
| G: Price: N150 (means you pay: N150)<br>You receive the information AND keep a payment of N50 | | X |
| H: Price: N175 (means you pay: N175)<br>You receive the information AND keep a payment of N25 | | X |
| I: Price: N200 (means you pay: N200)<br>You receive the information AND keep a payment of N0 | | X |

Do you have any questions?

[QUESTION 1. Other questions are similar.]

Consider that I want to sell information to you.

**Information to sell: The chances of people getting admitted into university based on the type of class they choose in SS1. For example, you will know the percentage of boys and girls that apply to study Arts, commercial and science courses in university and what percentage of them get admitted. This information is from JAMB.**