Online Appendix for "Sustaining Honesty in Public Service: The Role of Selection"

Sebastian Barfort   Nikolaj A. Harmon   Frederik Hjorth
Asmus Leth Olsen

Contents

A Online Appendix

A.1 Econometric Details

This section presents additional details and discussion regarding the econometric methods we apply to the data from our experimental dice game. In subsection A.1.1 we present a longer derivation of the estimated individual cheat rates that we use throughout the main analysis and in subsection A.1.2 we discuss how we use these estimates to examine the relationship between dishonesty, job preferences and other characteristics. In subsection A.1.3, we further characterize the measurement error in the estimated cheat rates. The remaining subsections focus on how to identify and estimate the full distribution of dishonesty. Subsection A.1.4 discusses identification of the cheat rate distribution, while subsection A.1.5 derives the specific Maximum Likelihood estimator we use to examine the distribution. Finally, subsection A.1.6 derives an estimator that jointly estimates the distribution of dishonesty and job preferences.

We use the same setup and notation as in the main text but repeat it here for convenience: The data contain information on a random sample of N respondents indexed by i. For each respondent we observe whether they report a win in each of the K different rounds of the dice game, which we index by k. We let y_{ik} be an indicator variable for whether respondent i reported winning in round k and let $Y_i = \sum_{k=1}^{K} y_{ik}$ denote the total number of reported wins. The probability of winning truthfully is independent across rounds equal to p^*, however, individuals may cheat and report a win regardless of the actual outcome. Individual i dishonestly reports a win some fraction $\theta_i \in [0, 1]$ of the time and reports the truth otherwise. We refer to θ_i as individual i's cheat rate.

For the purpose of this appendix, we also introduce some additional notation: We let F denote the distribution of cheat rates in the population, $\theta_i \sim F$. We let ψ_i denote individual i's expected share of reported wins across multiple rounds of the dice game, $\psi_i = E\left(\frac{Y_i}{K}|\theta_i\right)$ and let G denote the distribution of ψ_i in the population of interest, $\psi_i \sim G$. Finally, we let X_i denote some characteristic of individual i that is observed in the data (typically an indicator for whether individual i prefers a public service career).

A.1.1 Estimating Individual Dishonesty

We start by providing a more detailed derivation of the estimator of individual cheat rates that serves as our measure of dishonesty throughout the main analysis. The starting point is to note that in a given round, a respondent will win truthfully and report a win with probability p^*, or, if he does not win truthfully, he will dishonestly report a win with probability θ_i. Accordingly, the probability of observing a win for individual i, conditional on his cheat rate is:

$$P\left(y_{ik} = 1 \middle| \theta_i\right) = p^* + (1 - p^*)\theta_i$$

Since of course $P\left(y_{ik} = 1 \middle| \theta_i\right) = E\left(y_{ik} \middle| \theta_i\right)$, we can rearrange the above to get:

$$\theta_i = \frac{1}{1 - p^*} E\left(y_{ik} \middle| \theta_i\right) - \frac{p^*}{1 - p^*}$$

Replacing the expectation $E\left(y_{ik} \middle| \theta_i\right)$ by the corresponding population moment $\frac{1}{K} Y_i$ then yields the method of moments estimator (denoted $\widehat{CheatRate}_i$ in the main text):

$$\hat{\theta}_i = \frac{1}{1 - p^*} \frac{1}{K} Y_i - \frac{p^*}{1 - p^*}$$

We make two remarks regarding this estimator here:

First, it is straightforward to check that the estimator is unbiased:

$$E\left(\hat{\theta}_i | \theta_i\right) = \frac{1}{1 - p^*} \frac{1}{K} \sum_{k=1}^{K} E\left(y_{ik} \middle| \theta_i\right) - \frac{p^*}{1 - p^*} = \theta_i$$

Second, it is worth noting that the estimated cheat rate, $\hat{\theta}_i$, will be negative for any respondent who reports winning fewer than $K \frac{p^*}{1 - p^*}$ times, in spite of the fact that in fact $\theta_i \geq 0$ by assumption. It is possible to define different estimators that are non-negative, however, these estimator will generally not be unbiased. As we shall see in the next section, unbiasedness is particularly important given the analysis we conduct.

A.1.2 The Relationship between Dishonesty and Characteristics

The main aim of the paper is to estimate the relationship between dishonesty and job preferences or between dishonesty and other respondent characteristics. In this section, we discuss how this can be done in a regression framework using the individual estimated cheat rates from the previous section. For ease of exposition, we will frame the discussion specifically in terms of estimating the relationship between dishonesty and job preferences. Accordingly, in the rest of the section we will assume that X_i is some measure of individual i's preferences for public service, however, all the arguments go through if X_i is instead assumed to be some other characteristic of interest.

Throughout our analysis, we summarize the relationship between dishonesty and job preferences in the following linear regression:[1]

$$\theta_i = \beta_0 + \beta_1 X_i + \varepsilon_i \tag{1}$$

The object of interest here is the parameter β_1, which captures the relationship between individual i's cheat rate, θ_i and their job preferences, X_i. In the particular case where X_i is an indicator for whether i prefers a public service career, β_1 is just the difference in the mean cheat rate between individuals preferring a public service career and individuals preferring the private sector.

We can not directly estimate the regression above because we do not observe each individual's true cheat rate, θ_i. As discussed in the previous section, however, the data from the dice game allow us to construct an estimated cheat rate, $\hat{\theta}_i$, for each individual. As always, we can view this estimate as being equal to the true cheat rate plus a measurement error term ξ_i that is simply defined as $\xi_i \equiv \hat{\theta} - \theta_i$:

$$\hat{\theta}_i = \theta_i + \xi_i$$

Now, because $\hat{\theta}_i$ is unbiased for θ_i and because the measurement error in $\hat{\theta}_i$ stems

[1] More formally, we let β_0 and β_1 be defined in the usual (implicit) way by imposing $E(\varepsilon_i) = 0$ and $Cov(X_i, \varepsilon_i) = 0$ in (1).

solely from randomness in whether individuals' win truthfully or cheat in any specific round of the dice game, it follows that the measurement error ξ_i has mean zero and is mean independent of both θ_i and X_i:[2]

$$E\left(\xi_i|\theta_i, X_i\right) = 0$$

It follows that the measurement error ξ_i is classical.[3] As usual, we can therefore obtain a consistent estimate of our parameter of interest from a regression that uses the estimated cheat rate instead of the true cheat rate. Substituting $\theta_i = \hat{\theta}_i - \xi_i$ in (1) and rearranging we get:

$$\hat{\theta}_i = \beta_0 + \beta_1 X_i + (\varepsilon_i - \xi_i) \tag{2}$$

Because the measurement error is classical (in particular because $Cov(X_i, \xi_i) = 0$), OLS estimation of (2) will yield a consistent estimator for β_1 under the usual conditions. In other words, using the estimated cheat rate as the outcome variable instead of the true cheat rate does not affect the consistency of our estimates.

Note that the same is not true if we consider the reverse regression and instead regress of X_i on the estimated cheat rate $\hat{\theta}_i$. In this case our estimates will suffer from the usual attenuation bias. Accordingly, in our main analysis, we always focus on regressions that use the (estimated) cheat rate as the dependent variable.

A.1.3 The Degree of Measurement Error

The previous section showed that the measurement error in the individual estimated cheat rates do not affect the consistency of our estimates. Because the measurement error, ξ_i is absorbed into the composite error term in (2), however, the measurement error does

[2]Regardless of X_i, individual i's probability of winning truthfully in our dice game is p^* and his or her probability of being dishonest is θ_i. The conditional probability of reporting a win in the dice game is therefore $E\left(y_{ik}|\theta_i, X_i\right) = p^* + (1-p^*)\theta_i$. The same derivation that showed unbiasedness in Section A.1.1 therefore shows that $E\left(\hat{\theta}_i|\theta_i, X_i\right) = \theta_i$, which further implies $E\left(\xi_i|\theta_i, X_i\right) = 0$.

[3]The term "classical measurement error" is sometimes used to mean slightly differently things. Here we use it to refer to a situation in which the measurement error is uncorrelated with the true value and also uncorrelated with any other potential regressors.

increase the variance of the error term, which lowers precision and power. It is therefore of interest to examine the extent of the measurement error.

We can examine the measurement error in our estimated cheat rate by examining the variance of the estimator. To do this we have to take a stronger stance on the dependence of cheating behavior across rounds of the dice game. We focus here on the case where cheating behavior is independent across time for a given individual.[4] In this case, for an individual with cheat rate θ_i, the total number of reported wins, Y_i, is simply the number of successes in K independent trials with success probability $p^* + (1 - p^*)\theta_i$. Conditional on θ_i, Y_i therefore follows a binomial distribution:

$$Y_i|\theta_i \sim B(K, p^* + (1 - p^*)\theta_i) \tag{3}$$

Recall that our estimated cheat rate for each individual is $\hat{\theta}_i = \frac{1}{1-p^*}\frac{1}{K}Y_i - \frac{p^*}{1-p^*}$. Applying the standard formula for the variance of a binomially distributed random variable along with some simple algebra then yields the following expression for the variance of the estimated individual cheat rate (and thus for the extent of measurement error):

$$Var\left(\hat{\theta}_i \middle| \theta_i\right) = Var\left(\xi_i \middle| \theta_i\right) = \frac{\theta_i(1-\theta_i)}{K} + \frac{p^*}{(1-p^*)}\frac{(1-\theta_i)}{K}$$

From the above expression we see that the measurement error in our measure of dishonesty is increasing in p^* and decreasing in K. This motivates the design of our dice game which has a relatively low win probability, $p^* = \frac{1}{6}$ and asks respondents to repeat the dice game many times over, $K = 40$ (as we shall see in the next section, asking respondents to repeat the dice game offers additional advantages if one wants to estimate the full distribution of dishonesty).

A.1.4 Identification of the Full Distribution of Dishonesty

In the preceding sections we showed how to examine whether dishonest differs across individuals with different observables X_i. Next, we turn to the more basic question of

[4] It is conceptually straightforward to do similar derivations when dishonesty exhibits time dependence, however, it requires that one is willing to specify the exact form of time dependence.

examining how much dishonesty differs across the population overall. Accordingly, in this section, we consider how to estimate the full distribution of cheat rates in the population, F. As in the preceding section, this requires that we take a stronger stance on the time dependence of cheating behavior so from now on we maintain the assumption that cheating behavior is independent across time for a given individual.[5]

We start our discussion by focusing solely on identification, that is we ask what can be identified if we had experimental data on all individuals in the population of interest rather than just our specific sample of respondents. It turns out that the answer to this question depends crucially on how many times respondent repeat the dice game in the experiment, K. At one extreme, if each respondents only plays one round of the dice game, the data is completely uninformative about the extent of heterogeneity in dishonesty: When $K = 1$ it can be shown that any observed outcome of the experiment is consistent with a "no heterogeneity" scenario in which all indviduals have the same cheat rate.[6]

At the other extreme, we might consider what happens when the number of repetitions in the dice game becomes very large. When the number of rounds in the dice game grows to infinity, the share of observed wins reported for individual i converges to to the expected share of wins for this individual, $\frac{Y_i}{K} \xrightarrow{p} \psi_i$ when $K \to \infty$. In this case, the experimenter is therefore able to observe the distribution of expected share of wins across individuals, G. It is easy to show that this non-parametrically identifies the full distribution of dishonesty, F.[7]

Perhaps unsurprisingly, the empirically relevant case where K is large but finite turns out to fall between these two extremes. In this case, the experimental data is informative about the extent of heterogeneity in dishonesty but the full distribution of dishonesty is not in general non-parametrically identified. To build an understanding of why this

[5]Again, if one is willing to specify the exact form of time dependence in cheating, it is conceptually straightforward to adapt the results and estimators we present here to a situation with time dependence.

[6]When each respondent only participates in one round of the dice game, the data observed by the experimenter is just the number of individuals reporting a correct and an incorrect guess, which is equivalent to observing just the share of participants who report a correct guess in their roll, $P(y_{i1} = 1)$. Now let x be some observed value of $P(y_{i1} = 1)$. If all individuals are assumed to have a cheat rate of $\bar{\theta} = \frac{1}{1-p^*}x - \frac{p^*}{1-p^*}$, this exactly implies $P(y_{i1} = 1) = p^* + (1-p^*)\bar{\theta} = x$.

[7]$\eta_i = E(y_{ik} = 1|\theta_i) = p^* + (1-p^*)\theta_i$ so η_i is a linear transformation of θ_i. Knowledge of the distribution of η_i, therefore also pins down the distribution of θ_i.

is the case, note that when respondents play K rounds of the dice game, the outcome variable observed in the experiment, Y_i, takes on $K+1$ possible values (0 reported wins, 1 reported win, ... , K reported wins) for each individual. Accordingly, it can be shown that the informativeness of the experimental data regarding the distribution of cheat rates, F, can be summarized by $K+1$ moment conditions.[8]

While the moment conditions are generally very informative about the shape of F, $K+1$ equations will not generally be enough to non-parametrically identify a distribution.[9] Given that the full distribution of cheat rates is not non-parametrically identified, the next section proceeds by developing a parametric estimator for F.[10] As we shall see in Section A.2, our overall conclusions regarding F are robust to using different flexible parametric families for F, suggesting that the data is in fact highly informative about the distribution.

A.1.5 Estimation of the Full Distribution of Dishonesty

We now turn to the construction of an estimator for the full distribution of dishonesty, F. As discussed in the previous section, we start by restricting F to belong to some flexible parametric family of distributions on $[0,1]$, parameterized by a vector $\lambda \in \mathbb{R}^v$.[11]

---

[8] Formally, when cheating behavior is assumed independent over time, Y_i summarizes all the information the data contains about F. Moreover, because Y_i is a binomial random variable conditional on θ_i in this case, the conditional probability of observing x reported wins is just $P(Y_i = x|\theta_i) = \binom{K}{x}(p^* + (1-p^*)\theta_i)^x (1 - p^* + (1-p^*)\theta_i)^{K-x}$ for all $x = 0, 1, ...K$. Integrating over the distribution of cheat rates yields the unconditional probability of observing x wins and so that we arrive at the following $K+1$ moment conditions:

$$P(Y_i = x) = \int_0^1 r_{K,x}(\theta)\, dF(\theta) \quad \text{for} \quad x = 0, 1, ..., K$$

$$r_{K,x}(\theta_i) \equiv \binom{K}{x}(p^* + (1-p^*)\theta_i)^x (1 - p^* + (1-p^*)\theta_i)^{K-x}$$

[9]To see how the moment conditions are informative about F, note that the functions $r_{K,x}$ involved in the moment conditions (see footnote 8) are all positive, single-peaked and have their peaks located in different areas along $[0,1]$ As each of the moment conditions correspond to an integral over one of these functions, each moment conditions therefore provide information on how much mass the distribution F puts in a particular region of $[0,1]$. At the same time, if two candidate distributions F' and F^* give rise to the same value of the $K+1$ integrals involved in the moment conditions, the experimental data will not allow us to distinguish which one of them (if any) is the true distribution of dishonesty.

[10]We do not have a general identification result for the specific parametric families we consider. Across all our simulations and estimations on both the actual experimental data and various bootstrap samples, however, the distribution F has been identified within the parametric families we use.

[11]We discuss the specific choice of parametric family in Section A.2.1.

Once F is assumed to belong to some parametric family, we can develop a Maximum Likelihood estimator for F. When cheating behavior is independendent over time, the number of reported wins conditional on θ_i is a binomial random variable. The probability of Y_i reported wins is therefore $\binom{K}{Y_i}(p^* + (1-p^*)\theta_i)^{Y_i}(1 - p^* + (1-p^*)\theta_i)^{K-Y_i}$. We can then integrate out θ_i to get the unconditional probability of observing Y_i correct guesses:

$$\int_0^1 \binom{K}{Y_i}(p^* + (1-p^*)\theta)^{Y_i}(1 - p^* + (1-p^*)\theta)^{K-Y_i}\, dF(\theta;\lambda)$$

Given a sample of individuals with $Y_1, Y_2, ..., Y_N$ the log likelihood function is then:

$$\log \mathcal{L}(\lambda) = \sum_{i=1}^{N} \log\left(\int_0^1 \binom{K}{Y_i}(p^* + (1-p^*)\theta)^{Y_i}(1 - p^* + (1-p^*)\theta)^{K-Y_i}\, dF(\theta;\lambda)\right)$$

Maximization of the log likelihood function with respect to the parameter vector λ yields the Maximum Likelihood estimator for F. In Section A.2 we implement this estimator on our experimental data.

A.1.6 Joint Estimation of the Distribution of Dishonesty and Job Preferences

Over the preceding sections we first considered how to estimate the relationship between cheat rates and job preferences (or other characteristics) under minimal assumptions and then focused on how to estimate the distribution of dishonesty by invoking additional assumptions on the time dependence of cheating behavior. If one is willing to impose these additional assumptions throughout, however, it is possible to combine the two estimation problems and jointly estimate the distribution of both dishonesty and job preferences. We finish our section on econometric details by extending the Maximum Likelihood estimator from the previous section to this case.

In the rest of this section, we will treat X_i as an indicator for whether individual i prefers a public service careeer. In addition, we let $m(\theta_i)$ denote the conditional probability that some individual i prefers a public service career, conditional on his cheat rate: $P(X_i = 1|\theta_i) = m(\theta_i)$. To estimate the joint distribution of dishonesty and job preferences, we will construct a joint estimator of F and m.[12]

We maintain the same parametric assumption on F as in the previous section, so that estimation of $F(\cdot;\lambda)$ is again equivalent to estimation of λ. Our approach to estimating m will depend on whether F is assumed to be discrete. When F is discrete and the population consists of a finite number of types, each with a fixed cheat rate, we take a fully non-parametric approach and estimate a different probability of preferring a public service career for each of the types. When F is continuous (possibly including one or more masspoints), this non-parametric approach is not feasible and we instead impose a functional form on m. A convenient notation that covers both cases is to write m as a function of both the cheat rate θ and a real-valued vector $\zeta \in \mathbb{R}^q$, so that estimation of m is simply equivalent to estimation of ζ:[13]

$$P(X_i = 1|\theta_i) = m(\theta;\zeta)$$

Next, to derive the likelihood function, we note that conditional on the cheat rate, θ_i, the probability of observing an individual with Y_i reported wins in the dice game and a preference for public service X_i is just :

$$\left(\binom{K}{Y_i}(p^* + (1-p^*)\theta_i)^{Y_i}(1-p^* + (1-p^*)\theta_i)^{K-Y_i}\right)\cdot\left(m(\theta;\zeta)^{X_i}(1-m(\theta;\zeta))^{1-X_i}\right)$$

As before, we can then integrate out θ_i to arrive at the corresponding unconditional probability:

$$\int_0^1 \left(\binom{K}{Y_i}(p^* + (1-p^*)\theta_i)^{Y_i}(1-p^* + (1-p^*)\theta_i)^{K-Y_i}\right) \cdot\left(m(\theta;\zeta)^{X_i}(1-m(\theta;\zeta))^{1-X_i}\right) dF(\theta;\lambda)$$

[12] Joint knowledge of both of these of course permits one to compute the joint distribution of the variables θ_i and X_i as well as any corresponding conditional or marginal distribution.

[13] When F is discrete, each elements of ζ will be the probability of preferring a public service career for one of the discrete types in the population. When F is continuous, the elements of ζ will instead be the parameters of the functional form imposed on m. In both cases, estimation of m is simply equivalent to estimation of ζ.

Given a sample of individuals with $(Y_1, X_1), (Y_2, X_2), ..., (Y_N, X_N)$ we get the following log likelihood function:

$$\log \mathcal{L}(\lambda, \zeta) = \sum_{i=1}^{N} \log \left(\int_0^1 \left(\binom{K}{Y_i} (p^* + (1-p^*)\theta_i)^{Y_i} (1 - p^* + (1-p^*)\theta_i)^{K-Y_i} \right) \right.$$
$$\left. \cdot \left(m(\theta;\zeta)^{X_i} (1 - m(\theta;\zeta))^{1-X_i} \right) dF(\theta;\lambda) \right)$$

Maximization of the log likelihood function with respect to λ and ζ yields the joint Maximum Likelihood estimator for F and m. In Section A.2.5 we implement this estimator on our experimental data.

A.2 The Distribution of Dishonesty, Results

In this section we present estimates of the full distribution of cheat rates in our student population using the Maximum Likelihood estimators discussed in the previous sections. Since the estimators requires imposing a parametric assumption on he distribution of the cheat rate, F, we start by discussing the choice of parametric family and then present and compare the estimated distribution under different parametric assumptions. We finish by showing joint estimates of the distribution of cheat rates and the distribution of job preferences.

A.2.1 The Choice of Parametric Model

In deciding on a parametric familiy for the distribution of cheat rates, the first choice we need to make is whether to model the distribution as continuous or discrete. Since there are pros and cons to both approaches we consider and compare both approaches here.

In Section A.2.2 we show results using a continuous parametric family for the distribution of cheat rates. Given that the number of wins is distributed quite smoothly between 0 and 40 in our data, a continuous distribution of cheat rates is likely to give a good fit using only relatively few parameters. When modelling the distribution of cheat rates as continuous, we primarily use a mixture of Beta distributions. This makes up a very flexible class of distributions on $[0, 1]$. In particular, the Beta distribution can both allow for most of the mass being concentrated in the interior of $[0, 1]$ or at one or both of the end points. This allows us to capture that there may be significant share of individuals who are almost always honest and/or almost always dishonest, without imposing that this is necessarily the case.

A potentially unattractive feature of assuming a continuous distribution for F, is that by construction it will put zero mass on any indvidual point, including the two endpoints, 0 and 1. This implies that the share of completely honest and dishonest individuals in the population will always equal 0. To examine the prevalence of "extreme" types, however, we instead examine what fraction of individuals are *practically* completely honest or dishonest defined as cheating less than 2 percent or more than 98 percent of the time.

In Section A.2.2 we instead model the distribution of cheat rates as a discrete distribution. This is equivalent to assuming that the population consists of some finite number of types, each of which have a different cheat rate. While a relatively large number of types is likely necessary to fit the smooth distribution of cheat rates in our data, results using a discrete distribution can be easier to interpret in some cases. Many theoretical models of dishonesty use a discrete type spaces so estimates from a discrete distribution can provide a more natural link to theory. Moreover, discrete distributions allow for a strictly positive share of individuals to have cheat rates of exactly 0 or 1. Finally, specifying and estimating the conditional distribution of job preferences conditional on a given cheat rate is also particularly straightforward when the cheat rates distribution is discrete, as one can simply estimate a separate probability of preferring public service for each of the types.

A.2.2 Results Using a Continuous Distribution

In Table 5 we present estimates of the distribution of cheat rates when modelling the distribution (primarily) as continous. We present results using three different models for the distribution. Model (1) is our preferred model. It parameterizes the distribution of cheat rates as a mixture of two Beta distributions with parameters and weights to be estimated. Parameterizing the Beta-distributions in terms of mean and variance, the first column in the table show the estimated parameters and weights for each of the two components in the mixture. The corresponding estimated distribution of cheat rates is shown in Figure 5. We see extensive heterogeneity in dishonesty: 14.0 percent of individuals are estimated to be practically completely honest and cheat less than 2 percent of the time, while 17.0 percent are practically completely dishonest and cheat more than 98 percent of the time.[14] The standard deviation of the distribution is 0.39.

The rest of Table 5 presents results using alternative parametric forms for the distribution. Model (2) in the table extends Model (1) by including an additional Beta

[14]The standard error on the estimated share of the distribution that is practically completely honest is 2.0 percentage points. The standard error on the estimated share of the distribution that is practically completely dishonest is 1.4 percentage points.

Figure 5: Estimated distribution of dishonesty using a continuous distribution

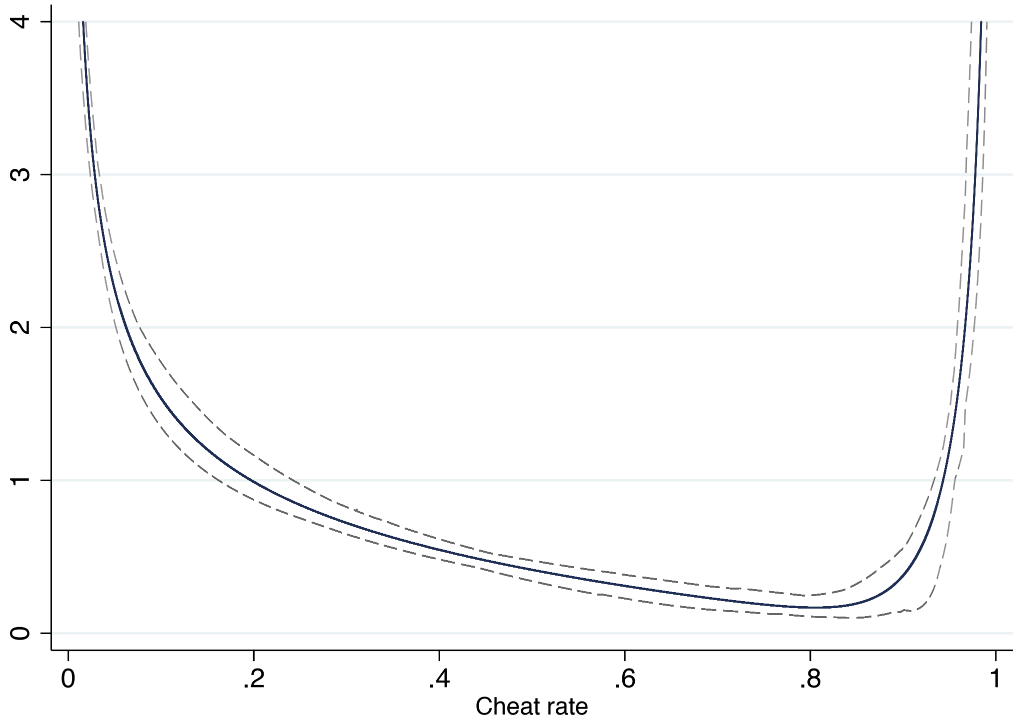

The figure shows the estimated probability density function for the distribution of cheat rates across students, based on a two component Beta-mixture. Dashed lines show pointwise 95 percent confidence intervals obtained via bootstrapping. Note that the y-axis is truncated; the function goes to infinity at the endpoints.

distribution in the mixture. The extra Beta distribution is estimated to have a weight of about 0.05, a mean of about 0.33 and a variance that is very close to zero. In practice this third estimated Beta-distribution in the mixture is thus indistinguishable from a discrete distribution with all its mass at 0.33. This motivates Model (3) in the table which instead extends Model (1) by including a mass point in addition to the continuous two component Beta-mixture. Similar to the results in Model (2), the included mass point is estimated to have a mass of about 0.05 and be located at 0.33.

Comparing the fit of the three models, the practical similarity of models (2) and (3) is evidenced by the fact that they both yield a log likelihood of -2813, whereas model (1) yields a slightly worse log likelihood of -2814. Since models (2) and (3) also include more free parameters, however, model selection based on standard information criteria (IC) suggests that Model (1) is preferred as it has a strictly smaller Bayesian IC and Akaike IC than both Model (2) and (3). Conducting Likelihood Ratio tests of Model (1) against

Figure 6: Actual vs predicted distribution of reported wins using a continuous distribution

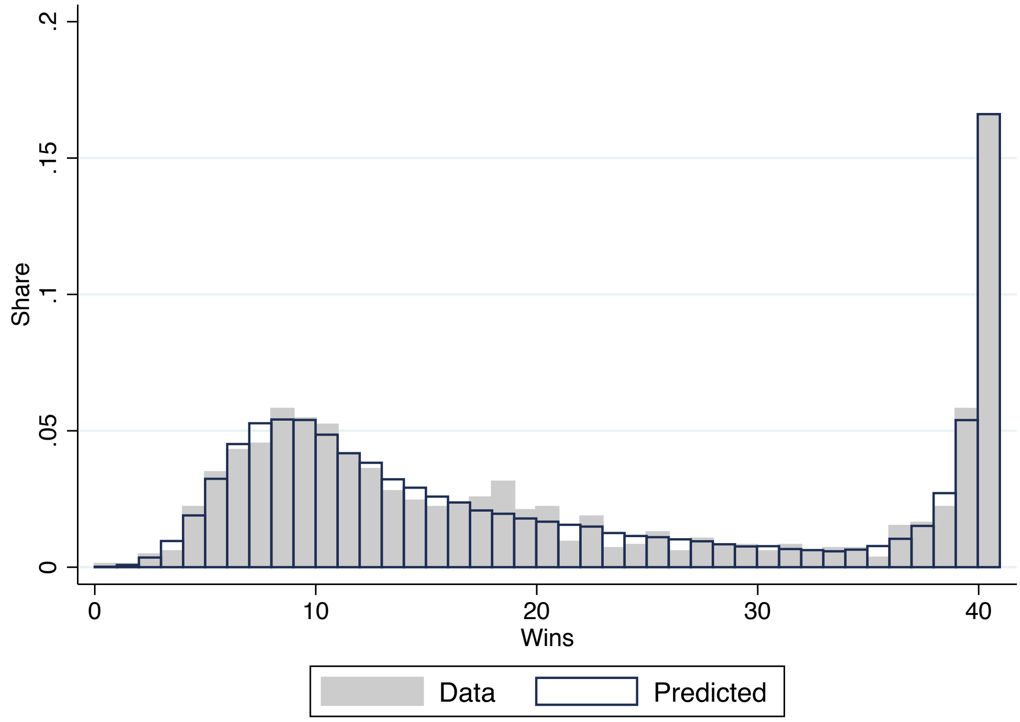

The histogram shows the observed number of correct guesses in the data as well as the predicted distribution based on the estimated distribution of cheat rates using a two component Beta-mixture (Model (1) of Table 5).

Model (2) and Model (3), we also cannot reject Model (1) at any conventional level of significance ($p = 0.15$ and $p = 0.18$ respectively).[15]

Finally, Figure 6 provides a different check on the fit of Model (1) by plotting the predicted distribution of correct guesses under the estimated distribution against the actually observed distribution of correct guesses. As the figure shows, the estimated distribution does a very good job of fitting the observed distribution.

A.2.3 Results Using a Discrete Distribution

In Table 6 we present estimates of the distribution of cheat rates, when modelling the distribution as discrete. Again we present results using three different models. Model (1) assumes that the population consists of six discrete types: a fully honest type with a

[15] Testing Model (1) against the other models implies testing whether one of the components in a mixture has zero weight. This is a non-standard testing problem. We therefore base the likelihood ratio test on McLachlan (1987)'s parametric bootstrap procedure for mixture distributions.

Table 5: Distribution of cheat rates, continuous distribution, detailed estimates

	(1)	(2)	(3)
		Beta-mixture component I:	
Weight	0.275	0.274	0.288
	(0.058)	(0.063)	(0.037)
Mean	0.975	0.975	0.975
	(0.052)	(0.038)	(0.018)
Variance	0.001	0.001	0.001
	(0.018)	(0.015)	(0.006)
		Beta-mixture component II:	
Weight	0.725	0.672	0.712
	(0.058)	(0.060)	(0.037)
Mean	0.214	0.205	0.205
	(0.018)	(0.056)	(0.038)
Variance	0.049	0.052	0.052
	(0.008)	(0.021)	(0.012)
		Beta-mixture component III:	
Weight	-	0.054	-
		(0.069)	
Mean	-	0.331	-
		(0.054)	
Variance	-	<0.001	-
		Additional mass point:	
Mass at point	-	-	0.052
			(0.057)
Mass point location	-	-	0.334
			(0.152)
Log likelihood	-2814	-2813	-2813
Akaike IC	5638	5644	5640
Bayesian IC	5662	5687	5673
p-value, LR-test, H_0 : Model (1)	-	0.149	0.178

The table shows maximum likelihood estimates for the distribution of cheat rates based on three different model specifications. Model (1) specifies the distribution to be a two-component beta-mixture. Model (2) specifies the distribution to be a three-component beta-mixture. Model (3) specifies the distribution to be mixture between a two-component beta-mixture and a mass point. For each model the estimated parameters and mixture weights are shown along with resulting Log Likelihood and Information criteria (IC). Bootstrapped standard errors are in parenthesis. The last row shows p-values of likelihood ratio tests, based on the parametric bootstrap of McLachlan (1987).

cheat rate of 0, a fully dishonest type with a cheat rate of 1, and four intermediate types with cheat rates strictly between 0 and 1 that are to be estimated from the data. The first column of the table shows the estimated population shares for each of the six types as well as the estimated cheat rates for each of the four intermediate types under Model (1). The second and third column of Table 5 shows corresponding results from alternative Models (2) and (3). These extend Model (1) by allowing for five or six intermediate types instead of only four.

Comparing the columns, we note initially that the three different models give rise to virtually indistinguishable estimated population shares and cheat rates for the fully honest type, as well as for the first two or three intermediate type. For the fully dishonest type and towards the top of the cheat rate distribution, the introduction of additional intermediate types changes estimates more however.

To asses which of the three model should be preferred, the bottom of the table presents various measures of fit. We see straight away that Model (2) dominates Model (1). Looking at standard model selection criteria, Model (2) has both a smaller Akaike IC and Bayesian IC than Model (1). In addition, a likelihood ratio test strongly rejects Model (1) against Model (2) ($p < 0.01$).[16]

The comparison of Model (2) and (3), however, is more complicated. The Akaike IC suggest that Model (1) is the preferred model, however the two Models give a similar value of the Bayesian IC. In addition, a likelihood ratio test of Model (2) against Model (3) rejects at the 5 percent level ($p = 0.037$), suggesting Model (3) as the preferred model.

In Panels A and B of Figure 7 we examine the estimated distribution of cheat rates under Models (2) and (3). Despite differences in the exact location and population shares for the more dishonest types, we note that the overall features of the two distributions are in fact very similar. In particular, the standard deviation of cheat rates is 0.39 under both distributions and the estimated share of fully honest individuals is 17.1 and 16.6 and percent respectively.

Focusing on the fully dishonest individuals, the two estimated distributions differs

[16] Again we use the parametric bootstrap of McLachlan (1987) to deal with the fact that this is a non-standard testing problem. See footnote 15.

Table 6: Distribution of cheat rates, discrete distribution, detailed estimates

	(1)	(2)	(3)
	Fully honest type (cheat rate equals 0):		
Share in population	0.177	0.171	0.166
	(0.037)	(0.040)	(0.042)
	Intermediate type I:		
Share in population	0.273	0.272	0.271
	(0.030)	(0.032)	(0.035)
Cheat rate	0.105	0.100	0.096
	(0.020)	(0.019)	(0.018)
	Intermediate type II:		
Share in population	0.196	0.195	0.193
	(0.022)	(0.023)	(0.027)
Cheat rate	0.355	0.343	0.333
	(0.021)	(0.026)	(0.031)
	Intermediate type III:		
Share in population	0.074	0.071	0.058
	(0.012)	(0.015)	(0.023)
Cheat rate	0.680	0.642	0.589
	(0.022)	(0.046)	(0.058)
	Intermediate type IV:		
Share in population	0.015	0.060	0.032
	(0.015)	(0.015)	(0.015)
Cheat rate	0.952	0.892	0.755
	(0.007)	(0.053)	(0.052)
	Intermediate type V:		
Share in population	-	0.177	0.068
		(0.040)	(0.027)
Cheat rate	-	0.986	0.922
		(0.012)	(0.017)
	Intermediate type VI:		
Share	-	-	0.212
			(0.032)
Cheat rate	-	-	0.992
			(0.003)
	Fully dishonest type (cheat rate equals 1):		
Share	0.138	0.054	<0.001
	(0.015)	(0.052)	
Log likelihood	-2841	-2833	-2832
Akaike IC	5773	5741	5751
Bayesian IC	5701	5689	5689
p-value, LR-test, H_0 : Model (1)	-	<0.01	<0.01
p-value, LR-test, H_0 : Model (2)	-	-	0.037

The table shows maximum likelihood estimates for the distribution of cheat rates based on three different model specifications. Each model assumes that there exist fully honest and dishonest individuals and then some number of intermediate types. For each model the estimated population shares and the estimated cheat rates for the intermediate types is shown along with the Log Likelihood and Information criteria (IC). Bootstrapped standard errors are in parenthesis. The last rows show p-values of likelihood ratio tests, based on the parametric bootstrap of McLachlan (1987).

more, as Model (2) estimates that 5.4 percent of individuals are fully dishonest, while Model (3) actually estimates that essentially noone in the population is fully dishonest. This difference, however, is much less stark when one notes that both models also estimate the existence of a large group of individuals who cheat *almost* all the time: Model (3) estimates that 21.2 percent of the population belong to a discrete type that cheats 99.2 percent of the time, while Model (2) estimates that 17.7 percent of the population belongs to a type that cheats 98.6 percent of the time. Accordingly, both models imply that 21-23 percent of the population cheat more than 98 percent of the time.[17]

Finally, Panels A and B of Figure 7 illustrate both the fit and similarity of Models (2) and (3) by showing the predicted distribution of reported wins for each of the models along with the actual distribution observed in the data. The predicted distribution under the two models is very similar and provides a good fit to the data.

A.2.4 Comparing Results Using Continuous and Discrete Distributions

Given the discussion in Section A.2.1 regarding the pros and cons of using a continuous vs discrete parametric family, it is instructive to compare the results we get using the two approaches. In Figure 9 we plot the estimated CDFs of the preferred models from the previous two subsections: The continuous distribution using a two component Beta-mixture, and the two discrete distributions with either 7 or 8 discrete types. We note that the estimated CDFs from the three models follow each other quite closely. Accordingly, the results from the three different models also yield quite similar conclusions regarding the heterogeneity in dishonesty. All three models imply that the standard deviation of cheat rates is 0.39.

If we focus instead on the share of people who are practically completely honest or dishonest, however, we note that the discrete models tend to imply a higher fraction of

[17]The fact that the precise distribution of individuals across the most dishonest types is sensitive to the choice parametric model but that the share of individuals cheating more that 98 percent of the time is not, illustrates the limits to what we can identify in our data. When respondents repeat the dice game 40 times, the difference in the expected number of reported wins between a fully dishonest individual and an individual with a cheat rate of 98.6 percent is less than one half win. Accordingly, the estimated shares become sensitive to the choice parametric model and the standard errors become large. If we lump the most dishonest types together however and simply ask what share of people cheat more than 0.98 percent, our data is much more informative.

Figure 7: Estimated distributions of dishonesty using discrete distributions

Panel A: Distribution of dishonesty, 7 discrete types, Model (2)

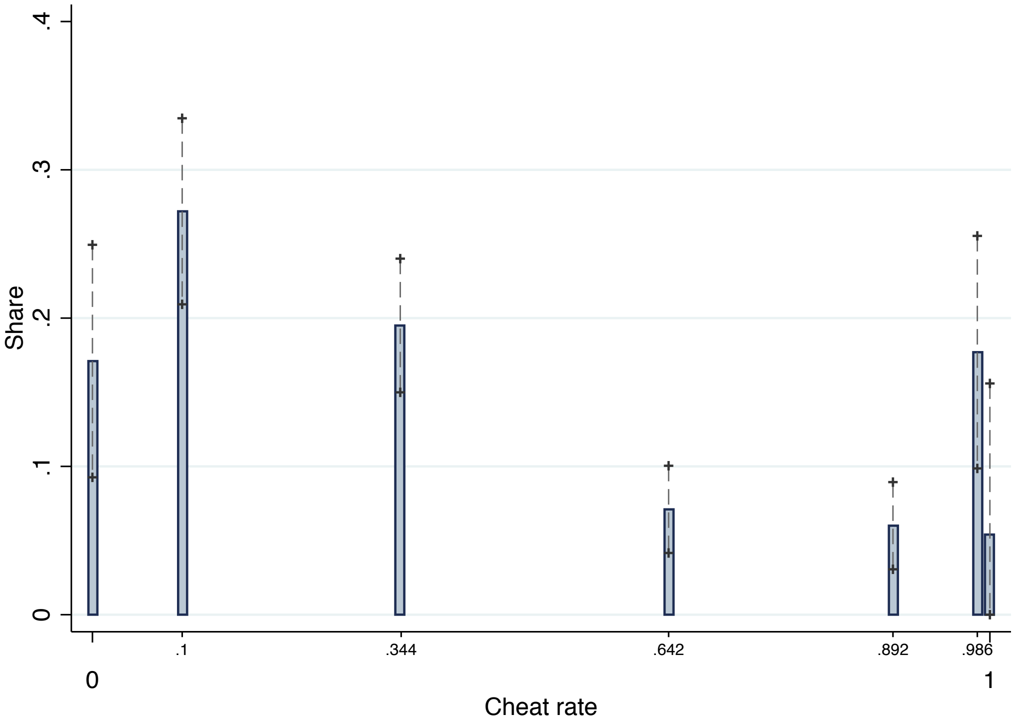

Panel B: Distribution of dishonesty, 8 discrete types, Model (3)

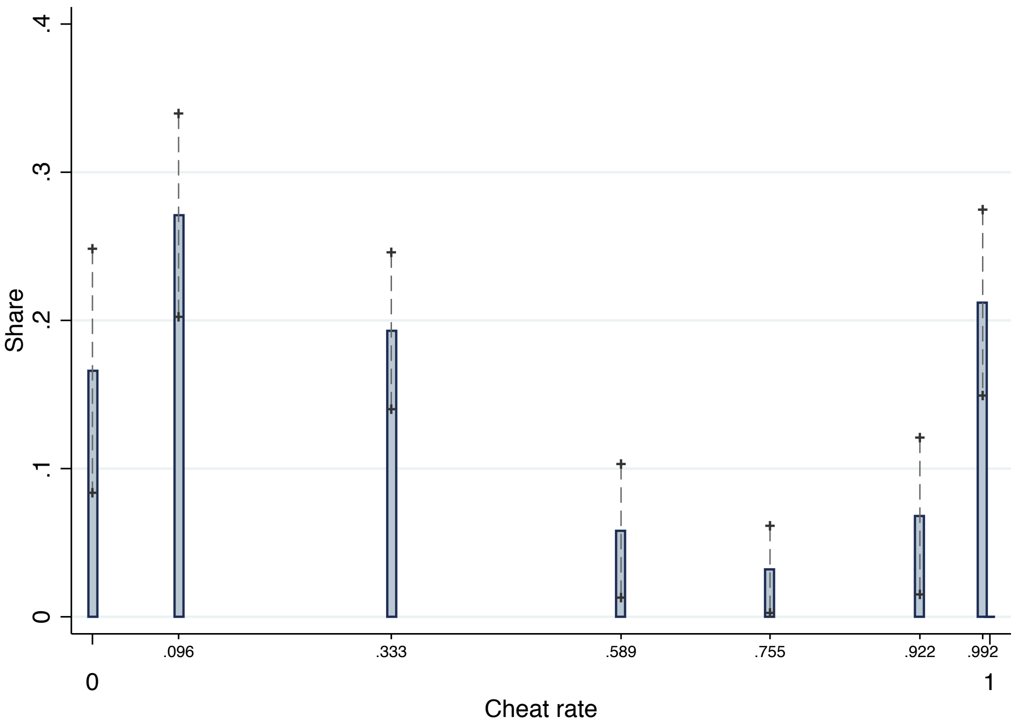

The two panels show the estimated distribution of dishonesty using different discrete distributions. The x-axis show the cheat rate for each of the discrete types, while the y-axis shows the population shares of each of the types along with 95 percent confidence intervals based on bootstrap standard errors. The confidence interval for the fully honest type is omitted in Panel B since its population share is estimated to be zero and thus be on the boundary of the parameter space.

Figure 8: Actual vs predicted distributions of reported wins using discrete distributions

Panel A: Distribution of reported wins, 7 discrete types, Model (2)

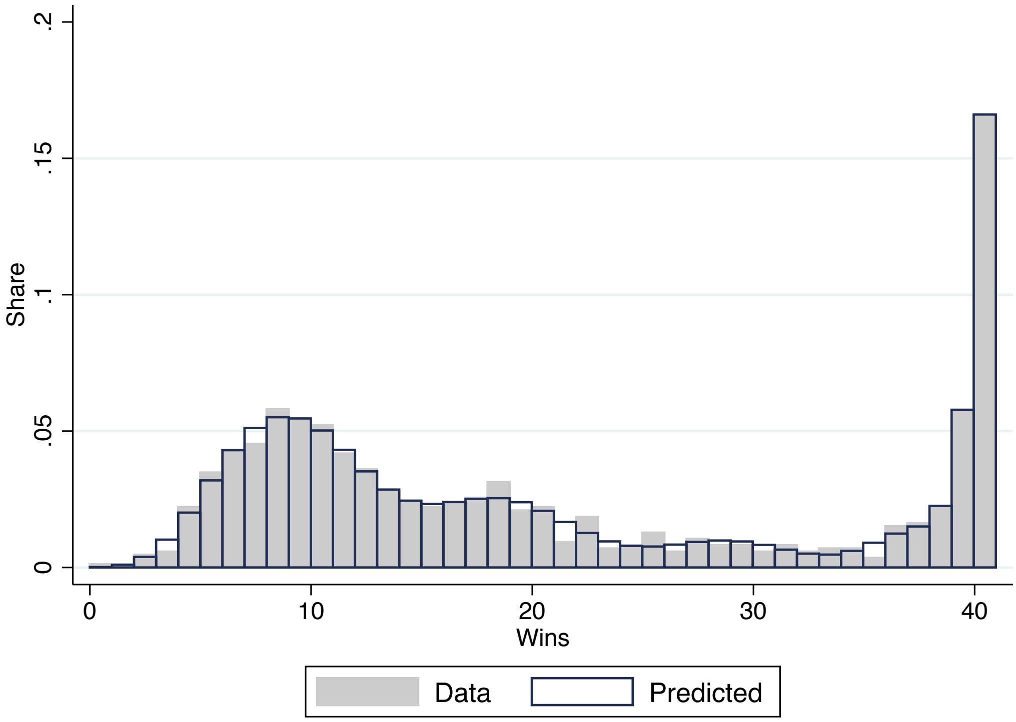

Panel B: Distribution of reported wins, 8 discrete types, Model (3)

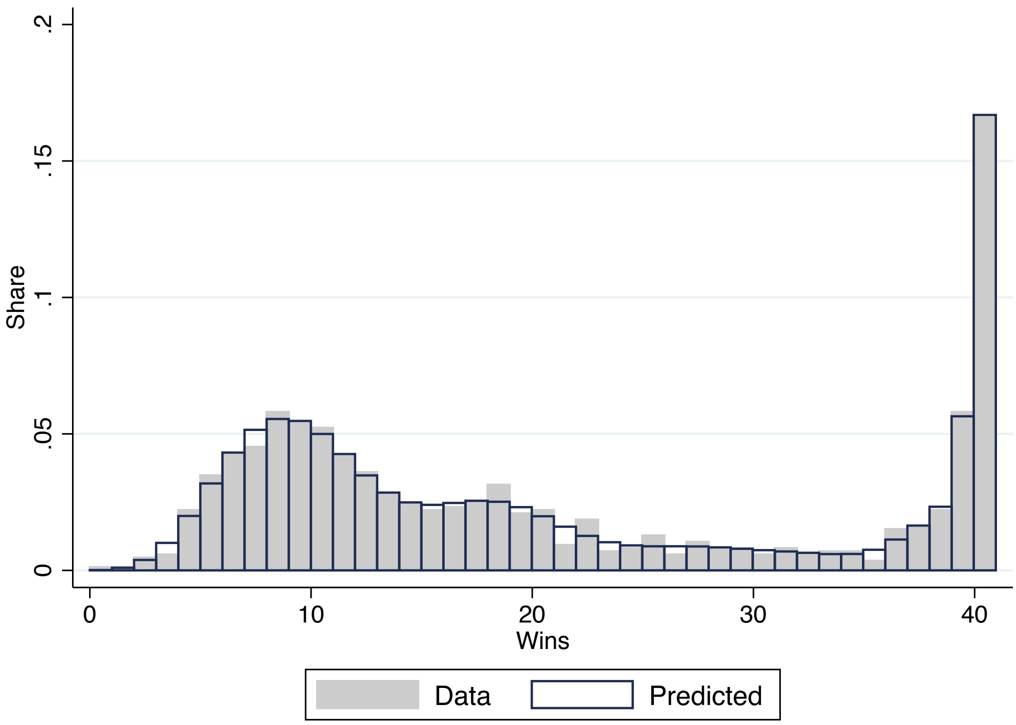

The two panels show the observed distribution of correct guesses in the data as well as the predicted distribution based on the estimated distribution of cheat rates using discrete distributions. Panel A shows results for Model (2) of Table 6, which assumes that there are 7 discrete types, while Panel B shows results for Model (3), which allows for 8 discrete types.

these "extreme" types. Using the continuous two component Beta-mixture, we estimate that the share of individuals cheating less than 2 percent of the time is 14.0 percent and the share of individuals cheating more than 98 percent of the time is 17.0 percent. When using the discrite distributions the corresponding shares are instead 16.6-17.1 percent and 21-23 percent.

Finally, we note that assuming a continuous distribution allows us to obtain a good fit to the data with a more parsimonious model. The preferred continuous model (the two component Beta-mixture) only has five free parameters, while the preferred discrete models have eleven and thirteen parameters respectively.[18] Accordingly, we see that model selection based on the Akaike or Bayesian IC that penalize models with more free parameters would imply that the continuous distribution is preferred over the discrete distribution.

A.2.5 Joint Estimation of Cheat Rate and Job Preference Distribution

We finish this section by presenting estimates of the joint distribution of dishonesty and job preferences using the Maximum Likelihood estimator presented in Section A.1.6. We again do this two different ways, treating the distribution of cheat rates as either discrete or continuous. When using a discrete distribution we focus on the model with seven discrete types, which yielded a positive estimated population shares for all types.[19] For each of the seven types we then estimate a separate probability of preferring a public service career. As in the main text, we define an individual to prefer a public service career if the individual ranked public administration in the top two of the eight job options in our survey.

For the continuous distribution we use a two component Beta-mixture. In this case, however, we need to impose a functional form on the probability of preferring public service as function of the cheat rate. To allow for a flexible relationship between job

[18]For the two component Beta-mixture, the free parameters are the mean and the variance of each of the two components in the mixture as well as the mixture weight of the first component. For the discrete models, the free parameters are the population share and cheat rate for each of the five or six intermediate types and the population share of the fully honest type.

[19]Types with a zero share obviously do not ever appear in the population. This creates an identification problem since we can not identify job preferences for types we never observe.

Figure 9: Estimated cumulative distribution functions using different models

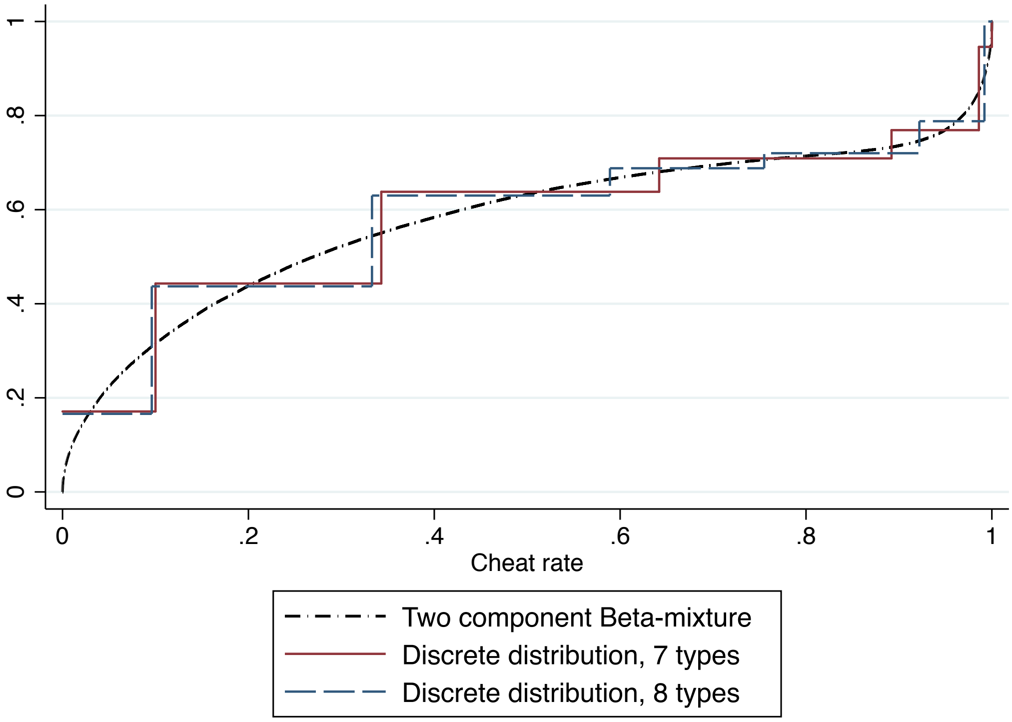

The Figure shows the estimated cumulative distribution functions for the distribution of cheat rates when using either a two component Beta-mixture, a discrete distribution with 7 types or a discrete distribution with 8 types.

prefences and dishonesty, we use a cubic polynomial in the cheat rate and apply the logistic function to restrict the probabilities to be between zero and one:

$$P(X_i = 1|\theta_i) = m(X_i = 1|\theta_i) = \frac{1}{1 + \exp\left(-\sum_{j=0}^{3} \kappa_j (\theta_i)^j\right)}$$

Panels A and B of Figure 10 shows the results using the two Models. The dots and solid line shows the conditional probability of preferring public service as a function of the cheat rate, while the dashed line and bars show the estimated distribution of cheat rates.[20]

The two panels of the figure show a very similar pattern. Among the most honest individuals, the probability of preferring public service is about 53 percent however this share falls rapidly with the cheat rate. Among individuals who cheat 35 percent of the time the share preferring public service is down to just under 40 percent. For individuals with cheat rates above 50 percent, however job preferences appear more stable. For these individuals, the share preferring public service ranges from 29 percent and 37 percent across the two panels.

Overall, we see that the systematic self-selection of honest individuals into public service, is driven by particularly strong preferences for public service among the most honest individuals, while the preferences for public service jobs differ less across the moderately to very dishonest.

[20]Since we are estimating the job preferences and the distribution of cheat rates jointly, the estimated distributions will not necessarily be the same as the ones in Sections A.2.2 and A.2.3. As the Figures show, however, we see only very slight differences in the distributions.

Figure 10: Jointly estimated cheat rate and job preference distributions

Panel A: Estimates using two component Beta-mixture

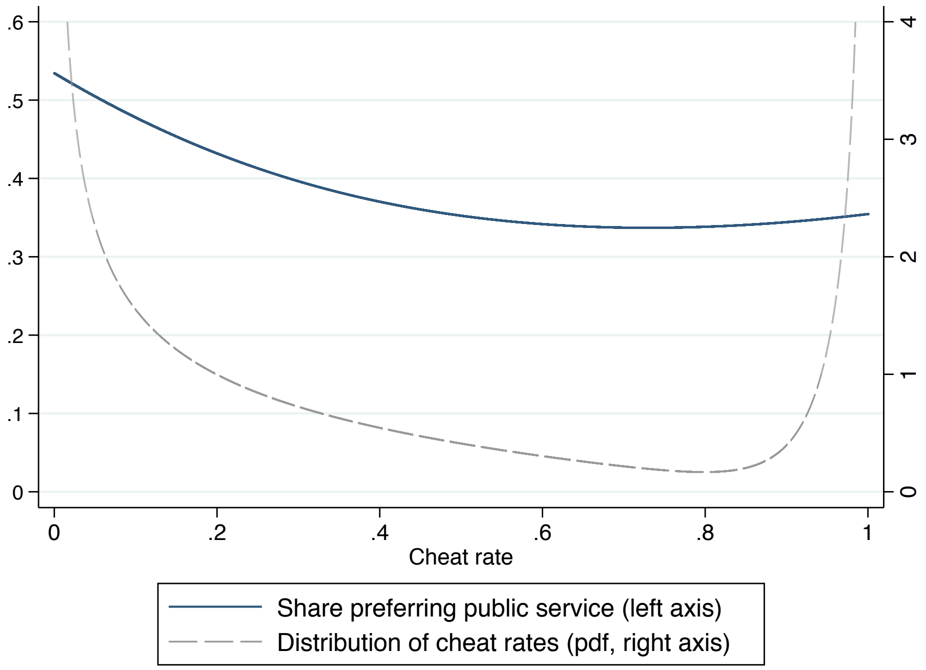

Panel B: Estimates using discrete distribution with 7 types

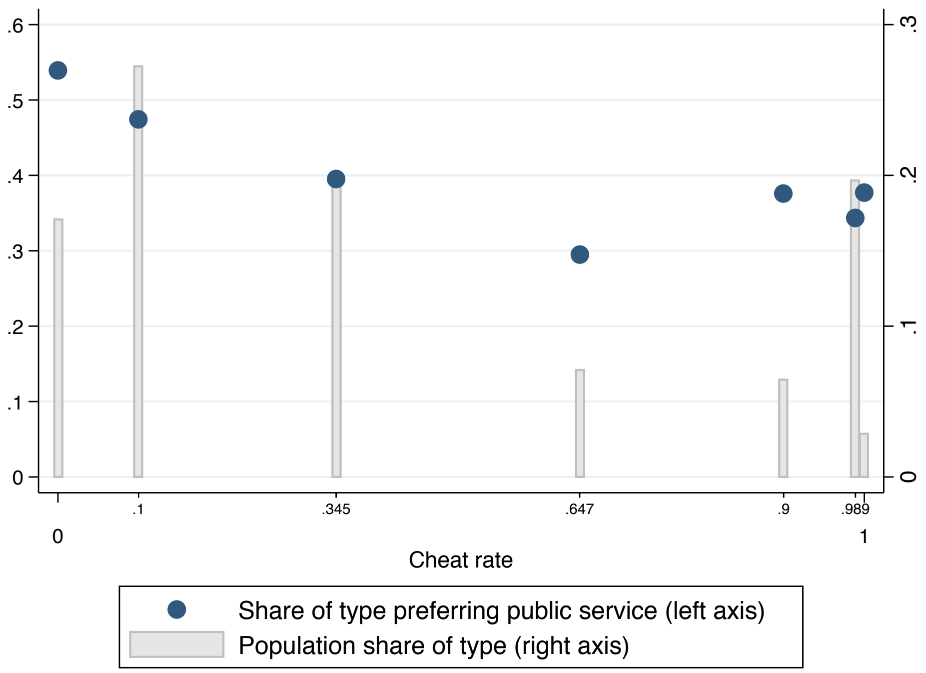

The panels show joint estimates of the distribution of cheat rates and job preferences. Panel A shows results using a two component Beta-mixture for the distribution of cheat rates and the logistic function of a polynomial for the job preference probabilities. The solid line shows the share of individuals who prefer public service as a function of the cheat rate (left axis). The dashed line shows the probability density function for the distribution of cheat rates (right axis). The right axis is truncated as the function goes to infinity at the endpoints. Panel B shows results when modelling the distribution of cheat rates using a discrete distribution with 7 types and allowing for a different conditional job preferencer probability for each type. The x-axis shows the cheat rates for each of the discrete types and the dots show the share of each type preferring a public service career (left axis). The bars show the population shares of the types (right axis).

A.3 Job Choices of Dishonest Students

In this section we look at which job categories dishonest students are particularly likely to prefer. To this end, Table 7 splits the sample into an *honest* and a *dishonest* half based on the estimated cheat rate and then compares how many students in each group rank the eight different job categories as their most preferred. The last row of the table thus restates the paper's main results by showing that public administration is ranked as the top job much more often for honest students than dishonest students: 26 percent of the honest half of students rank public administration as their preferred job, while only 17 percent of the dishonest half do so.

Looking at which jobs the dishonest half of students rank in the top instead of public administration, we see that by far the most important category is the financial sector. 19 percent of dishonest students rank the financial sector at the top versus only 8.6 percent among honest students: a bigger gap than we observe for any other job category. While jobs in the various listed categories may differ in many different dimensions, financial sector jobs particularly stand out as by far the best paid jobs for our student population. The popularity of financial sector jobs among dishonest students thus dovetails our findings regarding pro-social vs. pecuniary motivations. Dishonest individuals self-select out of the public sector jobs and into high-paying private sector jobs in part because they are more pecuniarily motivated.

Finally, two other job categories show statistically significant differences in how many honest vs. dishonest students rank them at the top. Dishonest students are 5.5 percentage points more likely to rank a central bank job at the top, while honest students are 5.3 percentage points more likely to rank a job in a political party or lobby organization at the top. These differences likely reflect that central bank jobs often serve as stepping stones for other financial sector jobs and that jobs in political parties and lobby organizations serve as stepping stones for running for political office.

Table 7: Top ranked job categories among less and more dishonest

Top ranked job	Est. cheat rate $<$ median	Est. cheat rate $\geq$ median	Difference	p-value
Financial sector	8.62	18.94	10.31	$<$0.01
Central bank	4.66	10.16	5.50	$<$0.01
Other private	19.11	20.79	1.67	0.60
Law firm	11.89	11.55	-0.34	0.96
Other public	3.96	3.23	-0.73	0.69
Public relations	6.76	4.16	-2.60	0.13
Political party or lobby org.	19.11	13.86	-5.26	0.05
Public administration	25.87	17.32	-8.55	$<$0.01

The table examines top ranked job categories among more dishonest vs. less dishonest students. Each row corresponds to a different job category. The first numerical column shows the fraction of students ranking each job category as the preferred one among students with an estimated cheat rates below the median. The second numerical column shows the fraction of students ranking each job category as the preferred one among students with an estimated cheat rates above the median. The last two columns shows the difference in these fractions for each of job category as well as the p-value for testing whether the difference is zero.

A.4 Comparing Dishonesty Measure with Previous Work

As discussed in the main text, our experimental approach to measuring dishonesty has been widely used in the literature and behavior in this type of experiment has been shown to predict fraudulent behavior among public sector employees by Hanna and Wang (2017). As always however, differences in stakes, framing and other implementation decisions may be a concern when comparing results to existing paper or relying on past validations of the experimental measures. In particular, since we draw on the variation in Jiang (2013), our computer-based *dice guessing* game differs from the canonical *dice under cup* game of Fischbacher and Föllmi-Heusi (2013) by asking to participants report (and possibly lie about) their own previous *guess* about a dice roll instead of reporting on the *outcome* of the dice roll.

To assess whether and how our specific implementation may have affected our experimental measures of inherent propensity for dishonest, this section compares the data from our survey experiment in Denmark with data on Indian students from the closely related experiment of Hanna and Wang (2017). In Hanna and Wang (2017) individual dishonesty was measured by asking each student to perform and report the outcome of 42 dice rolls, while paying 0.5 Indian Rupees (INR) for each eye rolled across the 42 dice rolls.

In Table 8 we examine the amount of dishonesty observed in the two data sets by comparing observed individual winnings in the two dishonesty games to the predicted distribution of winnings under full honesty.[21] In both data sets dishonesty is pervasive. 89 percent of students have winnings above the median in both cases. Overall, however, dishonest behavior appears somewhat higher in the Danish sample, especially towards the top of the distribution. 60 percent of students in our samples have winnings that are above the 99th percentile, compared to 33 percent of students in the Indian sample. This difference in the level of observed dishonesty is consistent with the conclusions of Jiang (2013), who finds that our *dice guessing* variation of the game leads to higher levels of dishonesty.

[21] In our experiment, the expected distribution of winnings under full honesty is distributed as a binomial random variable with 40 trials and a success probability of $\frac{1}{6}$ multiplied by 2 DKK. In Hanna and Wang (2017) the distribution of points under full honesty is simply the sum of 42 discrete uniform variables on 1, 2, 3, 4, 5, 6 multiplied by 0.5 INR.

Since the focus of the present paper is not on the *level* of dishonesty but on the *correlation* between inherent dishonesty and job preferences, differences in the measured level of dishonesty caused by the experimental design is less problematic. A much more severe concern is that our experimental design may bias the correlation between dishonesty and other variables. To assess this concern, we can look at the correlations between measured dishonesty and other respondent attributes in our data relative to the data from Hanna and Wang (2017). Three of the respondent attributes from our main analysis was also used in the analysis of Hanna and Wang (2017): GPA, donation in a dictator game and gender. In addition, our survey included a measure of External Locus of Control, which is used extensively in Hanna and Wang (2017).

In Table 9 we examine raw and partial correlations between dishonesty and these four attributes in the two data sets. To deal with differences in answer scales, currency denominations and grading scales, we standardize all the continuous variables within the two data sets and use standardized total winnings as our measure of individual dishonesty.[22] Looking at Columns 1 and 2, we see that the pattern of raw correlations is the same across the two samples. Dishonesty exhibits a statistically significant negative correlation with donations in the dictator game and a statistically significant positive correlation with being male. There is no significant correlation between dishonesty and GPA or External Locus of Control. Looking at the size of the observed correlations in the two samples, they are also very similar and are never significantly different from each other at conventional levels. For example, the correlation between donations in the dictator game in Hanna and Wang (2017)'s sample is -0.20, while it is -0.27 in our sample.

In Columns 3 and 4 we instead examine partial correlations by regressing individual dishonesty on the four other attributes simultaneously. A very similar picture emerges here. In both samples, the relationship between dishonesty and gender is no longer statistically significant once the other attributes are controlled for but dictator donations remains a statistically significant predictor in both samples. The actual size of the estimated

[22] The estimated cheat rate measure that we examine in the main text is a linear transformation of total winnings in our experiment. After standardization, we would thus get numerically the same results if we used our estimated cheat rates instead of total winnings. The main dishonesty measure in Hanna and Wang (2017) is total points in the dice game which is also a linear transformation of total winnings.

coefficients are again very similar across the two samples and never significantly different from each other. Finally, we see that the four attributes explain a similar fraction of the variation in dishonesty in the two samples. The R^2 from the linear regressions are 0.05 and 0.08 respectively.

To the extent that the true correlation between dishonesty and these four attributes is stable across Denmark and India, the very similar estimated correlations in Table 9 suggests that our experimental measure of dishonesty is comparable to the one used in Hanna and Wang (2017) despite any differences in the implementation of the experiment.

Table 8: Comparing the level of dice game cheating with previous literature

	Hanna and Wang (2017)	Danish sample
Share above 50th percentile of honest distribution	0.89	0.89
Share above 75th percentile of honest distribution	0.74	0.84
Share above 90th percentile of honest distribution	0.59	0.79
Share above 99th percentile of honest distribution	0.33	0.60

The table compares the amount of dice game cheating in the present paper's sample of Danish students with the amount of cheating among Indian students in the related experiment conducted by Hanna and Wang (2017). The rows of the table refer to different percentiles of the distribution of winnings that is expected under full honesty. The columns show how many participants had winnings above those percentiles in the two experiments.

Table 9: Comparing correlates of dice game cheating with previous literature

	Raw Correlations		Linear Regression	
	Hanna and Wang (2017)	Danish sample	Hanna and Wang (2017)	Danish sample
	(1)	(2)	(3)	(4)
GPA, standardized	0.046	0.014	0.050	0.032
	(0.051)	(0.035)	(0.052)	(0.033)
Dictator donation, standardized	-0.192**	-0.269**	-0.189**	-0.271**
	(0.054)	(0.032)	(0.052)	(0.033)
Male	0.125*	0.076*	0.170	0.055
	(0.058)	(0.034)	(0.111)	(0.046)
Locus of control, standardized	0.015	-0.045	0.005	-0.058
	(0.036)	(0.035)	(0.036)	(0.034)
N	614	862	614	862
R^2			0.050	0.078

The table compares correlates of dishonesty between the present paper's sample of Danish students and the sample of Indian students in the related experiment conducted by Hanna and Wang (2017). Columns (1) and (2) show raw correlations between total winnings and other characteristics and experimental measures in the two different samples. Columns (3) and (4) regresses standardized total winnings on other attributes and experimental measures in the two different samples. The attributes and measures used are standardized GPA, standardized donations in the dictator game, an indicator for being male and a standardized measure of external locus of control. Standard errors are in parenthesis. Robust standard errors are reported for the Danish sample, while standard errors clustered at the session level are reported for the Indian sample (see Hanna and Wang (2017) for details).$^*p<0.05$; $^{**}p<0.01$.

A.5 Validating Job Preference Measures

One potential concern with the data from our survey experiment is that students' stated job preferences may be poor measures of actual job preferences at the time when students finish their degrees and enter the labor market. While most of the other experimental measurements in the survey are incentivized using monetary stakes, the job preference questions are not. This may raise questions about the validity of the given answers. Moreover, it is possible that students' job preferences may change between the time of the survey and the time of graduation.

In this section we attempt to validate our job preference measures against actual job outcomes after graduation using administrative data. In doing so we exploit the fact that students in our survey experiment have been recruited from university registers that contain the unique Danish person identifier (the so-called "Central Person Registry" number). Because of this we are able to link individual student responses from the survey experiment to centrally-collected administrative data on completed degrees and actual employment outcomes.[23]

Because of long degree completion times and because of time lags in the availability of administrative data, a comprehensive examination of the actual labor market outcomes for the students in our experiment is not possible.[24] For a subset of the oldest students in the data, however, we are able to examine how their stated job preferences correlate with actual job outcomes after graduation. From the most recent administrative data available to us, we are able to construct information on whether students in our sample had completed their degree by October 2016, whether they held a job as of January 2017 and if so whether their job was in public administration.

After matching our sample of 862 students to the administrative data, we find that 155 individuals (18 percent) had both completed their degree by October 2016 and had a

[23]We note that actual job outcomes may also be an imperfect measure of job preferences. Actual job outcomes conflate the job preferences of individuals with the screening and sorting that occurs when people are hired into public service jobs. The questions in our survey experiment can only aim to measure the job preferences of respondents.

[24]It will be several years until such an examination is possible. Historically, a substantial fraction of students in our study population have taken up to a full seven years to finish their degree and enter the labor market. The youngest students in our sample (those starting at the university in 2014) thus will not be fully in the labor market until 2021.

job in January 2017. Beyond limiting statistical power, having job outcomes only for this modest subset of individuals raises obvious concerns about selection. Nonetheless, the data allows us to get some sense of how stated job preferences in the survey experiment predicts actual post-graduation job outcomes.

Table 10 shows regression results from the linked data. In Column 1 we regress an indicator for having a public administration job in January 2017 on the main measure of job preferences used in the paper: an indicator for ranking public administration among the top two jobs in the survey experiment. We see that the stated preferences in the survey are highly predictive of the actual job outcome. Individuals who ranked public administration among the top two jobs in the survey experiment are 48 percentage points more likely to be in a public administration job in January 2017 and this difference in highly significant. In Columns 2 to 5, we repeat this specification for the other measures of public service job preferences used in the paper. With the exception of the public service motivation score, all of the measures show a positive and highly significant correlation with the actual job outcome. For public service motivation, the coefficient is also positive but not significantly different from zero ($p = 0.15$). With the caveat that we are only able to examine a modest subset of our respondents, we overall conclude that the stated job preferences seem to correlate very well with observed job outcomes.

Finally, given that we have linked the experimental data linked to actual job outcomes, it is natural to also ask how our experimental measure of dishonesty relates to actual job outcomes in the administrative data. We examine this in the rest of Table 10. We first check to what extent our main results on dishonesty and (stated) job preferences hold in the subsample of respondents where we have linked data on post-graduation job outcomes. Focusing only on this subsample, Column 6 replicates the main specification from the paper by regressing the estimated cheat rate on an indicator for ranking public administration among the top two jobs. We exactly replicate the point estimate of 0.1 from the main paper. With the smaller number of observations and resulting larger standard errors, however, we can no longer rule out a zero coefficient at conventional levels of significance ($p = 0.14$). Next, in Column 7 we regress the estimated cheat rate

Table 10: Validating survey experiment against administrative data

	Has public administration job in January 2017					Estimated cheat rate	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Public administation ranked ≤ 2	0.479**					-0.101	
	(0.0714)					(0.0671)	
Higher ranking of public administration		0.104**					
		(0.0182)					
Public service motivation score			0.123				
			(0.0845)				
Public sector picked at current wage				0.235**			
				(0.0846)			
Probability of public administration					1.465**		
					(0.176)		
Has public administration job in Jan 2017							-0.0852
							(0.0665)
Constant	0.235**	0.785**	0.144	0.377**	0.105*	0.489**	0.482**
	(0.0463)	(0.0668)	(0.211)	(0.0474)	(0.0525)	(0.0433)	(0.0456)
N	155	155	155	155	155	155	155

The table shows regressions for the sample of subjects which had completed their degree by October 2016 and held some job in January 2017. In Columns 1 to 5, the outcome variable is an indicator for whether the respondents' job in January 2017 was in public administration. In Columns 6 and 7 the outcome variable is the subjects' estimated cheat rate. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap, the subjective probability of ending up in public administration. Robust standard errors in parenthesis. $^{*}p<0.05$; $^{**}p<0.01$.

on the indicator for having a job in public administration in January 2017. We find a similar negative estimate: respondents with a public administration job in January 2017 on cheated 9 percent less on average than those who held a different job in January 2017. As in Column 6, however, this difference is not statistically significant ($p = 0.20$). With the important caveats that estimates are imprecise and that we are looking at a selected sample of respondents, these results are at least indicative that the self-selection pattern in the survey experiment carries over to actual job outcomes.

A.6 Win Rates Across Dice Rolls

To examine how dishonest behavior evolves over the repetitions of our dice game, Figure 11 shows the average win rate for each of the 40 repetitions of the dice game. With the possible exception of the very first roll, we see that the win rate is quite stable across rounds.

Figure 11: Win rates across dice rolls

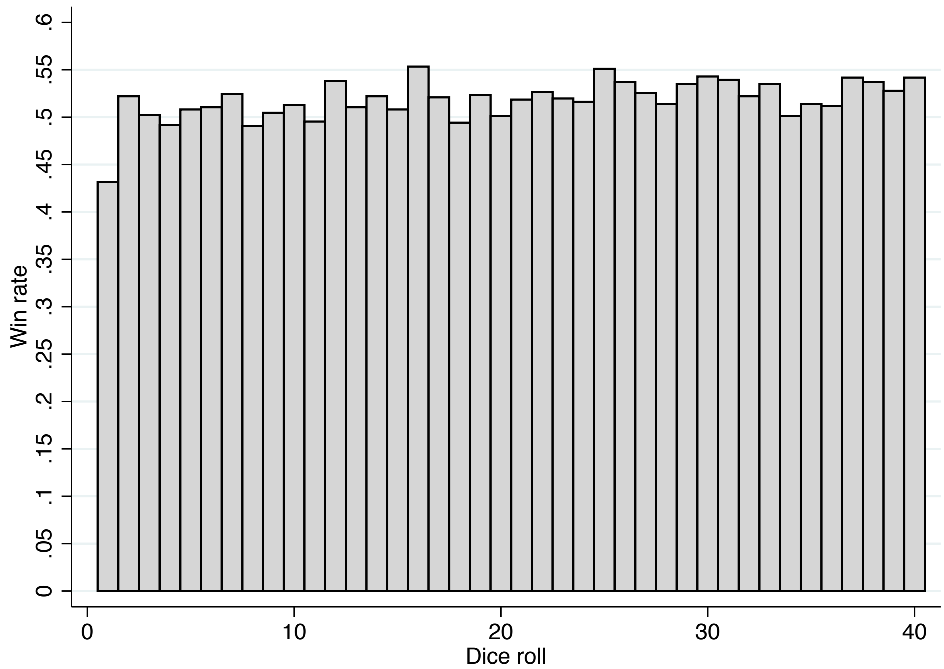

The figure shows the win rate across individuals separately for each of the 40 repetitions of the dice game.

A.7 Robustness Checks

In this section, we present a series of robustness checks to shore up various concerns with our empirical analysis:

First, our implementation of the dice-under-cup approach differs from many previous implementations in that we ask respondents to play many rounds of the game. This repetition may raise concerns that respondents become fatigued or otherwise change their game perception or behavior. As a robustness check, Tables 11-15 and 18-22 therefore reexamines the correlation between dishonesty, job preferences and other attributes using different subsets of the 40 dice rolls in our data. In particular, we consider using only the very first dice roll, dice rolls 1-10, dice rolls 11-20, dice rolls 21-30 or dice rolls 31-40.

Second, given the student population we focus on, another concern is that the behavior of some respondents may be affected by knowledge of the existing academic literature on dishonesty and its relation to our experimental tasks. At the end of the survey experiment, we asked respondents whether they had prior familiarity with any of its elements. Independent coding of the responses show that 40 respondents expressed awareness of either dice-under-cup games, similar experimental games (e.g. coin flipping), or explicitly mentioned the potential for cheating. Table 16 and 23 reexamines the correlation between dishonesty, job preferences and other attributes after excluding these respondents.

Third, our sample includes 143 respondents who cheat on all dice rolls and report the maximum number of correct guesses in our dice-under-cup games. As an additional robustness check, Table 17 and 24 reexamines the correlation between dishonesty, job preferences and other attributes without these respondents.

As Tables 11 to 24 show, the papers conclusions are robust to all three alternative sample restrictions. Besides the obvious loss of precision when dropping observations, the alternative specifications lead to very similar results as the ones presented in the main text.

Table 11: Estimated cheat rates and public service job preferences using only the first dice game

	Estimated cheat rates for first dice roll				
	(1)	(2)	(3)	(4)	(5)
Public administation ranked ≤ 2	-0.097^{*}				
	(0.041)				
Higher ranking of public administration		-0.018			
		(0.010)			
Public service motivation score			-0.134^{**}		
			(0.039)		
Public sector picked at current wage				-0.072	
				(0.045)	
Probability of public administration					-0.103
					(0.162)
Constant	0.359^{**}	0.256^{**}	0.645^{**}	0.338^{**}	0.339^{**}
	(0.027)	(0.039)	(0.097)	(0.024)	(0.039)
N	862	862	860	862	858

The table shows regressions of students' estimated cheat rates on various measures of public service job preferences, where the cheat rate estimated is based only on the first dice game. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap and the subjective probability of ending up in public administration. Robust standard errors in parentheses. *p<0.05; **p<0.01.

Table 12: Estimated cheat rate and public service job preferences using only dice rolls 1-10

	Estimated cheat rate for dice rolls 1-10				
	(1)	(2)	(3)	(4)	(5)
Public administation ranked $\leq$ 2	-0.088^{***}				
	(0.027)				
Higher ranking of public administration		-0.021^{***}			
		(0.006)			
Public service motivation score			-0.146^{***}		
			(0.026)		
Public sector picked at current wage				-0.076^{***}	
				(0.029)	
Probability of public administration					-0.254^{**}
					(0.103)
Constant	0.437^{***}	0.327^{***}	0.756^{***}	0.421^{***}	0.453^{***}
	(0.018)	(0.026)	(0.066)	(0.016)	(0.025)
N	862	862	860	862	858

The table shows regressions of subjects' estimated cheat rate on various measures of public service job preferences, where the cheat rate estimated is based only on the dice rolls 1-10. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap, the subjective probability of ending up in public administration. Robust standard errors in parenthesis. *p<0.1; **p<0.05; ***p<0.01.

Table 13: Estimated cheat rate and public service job preferences using only dice rolls 11-20

	Estimated cheat rate for dice rolls 11-20				
	(1)	(2)	(3)	(4)	(5)
Public administation ranked $\leq$ 2	-0.096^{***}				
	(0.028)				
Higher ranking of public administration		-0.022^{***}			
		(0.007)			
Public service motivation score			-0.136^{***}		
			(0.028)		
Public sector picked at current wage				-0.074^{**}	
				(0.031)	
Probability of public administration					-0.285^{***}
					(0.111)
Constant	0.460^{***}	0.345^{***}	0.753^{***}	0.441^{***}	0.479^{***}
	(0.019)	(0.027)	(0.070)	(0.017)	(0.027)
N	862	862	860	862	858

The table shows regressions of subjects' estimated cheat rate on various measures of public service job preferences, where the cheat rate estimated is based only on the dice rolls 11-20. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap, the subjective probability of ending up in public administration. Robust standard errors in parenthesis. *p<0.1; **p<0.05; ***p<0.01.

Table 14: Estimated cheat rate and public service job preferences using only dice rolls 21-30

	Estimated cheat rate for dice rolls 21-30				
	(1)	(2)	(3)	(4)	(5)
Public administation ranked $\leq$ 2	-0.113^{***}				
	(0.029)				
Higher ranking of public administration		-0.024^{***}			
		(0.007)			
Public service motivation score			-0.165^{***}		
			(0.027)		
Public sector picked at current wage				-0.100^{***}	
				(0.032)	
Probability of public administration					-0.339^{***}
					(0.114)
Constant	0.482^{***}	0.353^{***}	0.837^{***}	0.462^{***}	0.505^{***}
	(0.019)	(0.027)	(0.069)	(0.017)	(0.027)
N	862	862	860	862	858

The table shows regressions of subjects' estimated cheat rate on various measures of public service job preferences, where the cheat rate estimated is based only on the dice rolls 21-30. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap, the subjective probability of ending up in public administration. Robust standard errors in parenthesis. *p<0.1; **p<0.05; ***p<0.01.

Table 15: Estimated cheat rate and public service job preferences using only dice rolls 31-40

	Estimated cheat rate for dice rolls 31-40				
	(1)	(2)	(3)	(4)	(5)
Public administation ranked $\leq$ 2	-0.111^{***}				
	(0.029)				
Higher ranking of public administration		-0.023^{***}			
		(0.007)			
Public service motivation score			-0.161^{***}		
			(0.028)		
Public sector picked at current wage				-0.109^{***}	
				(0.032)	
Probability of public administration					-0.261^{**}
					(0.112)
Constant	0.479^{***}	0.355^{***}	0.825^{***}	0.463^{***}	0.487^{***}
	(0.019)	(0.027)	(0.070)	(0.017)	(0.027)
N	862	862	860	862	858

The table shows regressions of subjects' estimated cheat rate on various measures of public service job preferences, where the cheat rate estimated is based only on the dice rolls 31-40. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap, the subjective probability of ending up in public administration. Robust standard errors in parenthesis. *p<0.1; **p<0.05; ***p<0.01.

Table 16: Estimated cheat rates and job preferences excluding students with dice game experience

	Estimated cheat rate				
	(1)	(2)	(3)	(4)	(5)
Public administration rank ≤ 2	-0.103^{**}				
	(0.027)				
Higher ranking of public administration		-0.022^{**}			
		(0.007)			
Public service motivation score			-0.144^{**}		
			(0.026)		
Public sector picked at current wage				-0.093^{**}	
				(0.030)	
Probability of public administration					-0.295^{**}
					(0.109)
Constant	0.453^{**}	0.332^{**}	0.762^{**}	0.435^{**}	0.470^{**}
	(0.018)	(0.026)	(0.067)	(0.016)	(0.026)
N	822	822	820	822	818

The table shows regressions of students' estimated cheat rates on various measures of public service job preferences, excluding students that explicitly indicated that they were cheating or had prior knowledge of the dice task. The exclusion was based on students responses in an open-ended text box in which they were asked about their impression of the survey and whether they had prior familiarity with any of its elements. The exclusion is based on an independent coding of the responses. It indicated that 40 students expressed awareness of either dice-under-cup games, similar experimental games (e.g. coin flipping), or explicitly mentioned the potential for cheating. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual ranked given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap and the subjective probability of ending up in public administration. Robust standard errors in parentheses. $^{*}p<0.05$; $^{**}p<0.01$.

Table 17: Estimated cheat rates and job preferences excluding students with 100% win rate

	Estimated cheat rate				
	(1)	(2)	(3)	(4)	(5)
Public administation ranked ≤ 2	-0.080^{**}				
	(0.024)				
Higher ranking of public administration		-0.017^{**}			
		(0.006)			
Public service motivation score			-0.093^{**}		
			(0.025)		
Public sector picked at current wage				-0.053^{*}	
				(0.027)	
Probability of public administration					-0.209^{*}
					(0.094)
Constant	0.342^{**}	0.249^{**}	0.537^{**}	0.322^{**}	0.351^{**}
	(0.017)	(0.023)	(0.065)	(0.015)	(0.023)
N	719	719	717	719	716

The table shows regressions of students' estimated cheat rates on various measures of public service job preferences, excluding students who reported a correct guess for all dice rolls. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual ranked given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap and the subjective probability of ending up in public administration. Robust standard errors in parentheses. *p<0.05; **p<0.01.

Table 18: Estimated cheat rates and other attributes using only the first dice game

	Estimated cheat rates for first dice roll						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.019						0.028
	(0.020)						(0.020)
Picks risky lottery		0.053					0.025
		(0.041)					(0.042)
Job security rank ≤ 2			0.007				0.009
			(0.063)				(0.062)
Donation				-0.011^{**}			-0.011^{**}
				(0.003)			(0.003)
Wage rank ≤ 2					0.034		-0.006
					(0.045)		(0.045)
Male						0.155^{**}	0.143^{**}
						(0.040)	(0.042)
Constant	0.319^{**}	0.291^{**}	0.317^{**}	0.396^{**}	0.308^{**}	0.235^{**}	0.305^{**}
	(0.020)	(0.029)	(0.022)	(0.030)	(0.024)	(0.029)	(0.043)
N	861	862	862	862	862	862	861

The table shows regressions of students' estimated cheat rates on various measures of other student attributes, where the cheat rate estimated is based only on the first dice game. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics *p<0.05; **p<0.01.

Table 19: Estimated cheat rates and other attributes using only dice rolls 1-10

	Estimated cheat rate for dice rolls 1-10						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.013						0.019
	(0.014)						(0.013)
Picks risky lottery		0.030					0.030
		(0.027)					(0.028)
Job security ranked ≤ 2			0.004				0.001
			(0.040)				(0.039)
Donation				-0.015^{***}			-0.014^{***}
				(0.002)			(0.002)
Wage ranked ≤ 2					0.069^{**}		0.037
					(0.029)		(0.029)
Male						0.057^{**}	0.035
						(0.027)	(0.028)
Constant	0.401^{***}	0.385^{***}	0.399^{***}	0.500^{***}	0.380^{***}	0.369^{***}	0.454^{***}
	(0.014)	(0.019)	(0.015)	(0.019)	(0.016)	(0.019)	(0.028)
N	861	862	862	862	862	862	861

The table shows regressions of students' estimated cheat rates on various measures of other student attributes, where the cheat rate estimated is based only on dice rolls 1-10. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics *p<0.05; **p<0.01.

Table 20: Estimated cheat rates and other attributes using only dice rolls 11-20

	Estimated cheat rate for dice rolls 11-20						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.004						0.010
	(0.015)						(0.014)
Picks risky lottery		0.036					0.036
		(0.028)					(0.029)
Job security ranked ≤ 2			-0.023				-0.027
			(0.043)				(0.041)
Donation				-0.017^{***}			-0.016^{***}
				(0.002)			(0.002)
Wage ranked ≤ 2					0.081^{***}		0.045
					(0.030)		(0.030)
Male						0.060^{**}	0.032
						(0.028)	(0.029)
Constant	0.421^{***}	0.402^{***}	0.423^{***}	0.533^{***}	0.397^{***}	0.388^{***}	0.487^{***}
	(0.014)	(0.019)	(0.015)	(0.020)	(0.017)	(0.020)	(0.029)
N	861	862	862	862	862	862	861

The table shows regressions of students' estimated cheat rates on various measures of other student attributes, where the cheat rate estimated is based only on dice rolls 11-20. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics *p<0.05; **p<0.01.

Table 21: Estimated cheat rates and other attributes using only dice rolls 21-30

	Estimated cheat rate for dice rolls 21-30						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.010						0.017
	(0.015)						(0.015)
Picks risky lottery		0.038					0.040
		(0.029)					(0.029)
Job security ranked ≤ 2			0.020				0.017
			(0.042)				(0.040)
Donation				-0.017^{***}			-0.017^{***}
				(0.002)			(0.002)
Wage ranked ≤ 2					0.084^{***}		0.048
					(0.031)		(0.031)
Male						0.065^{**}	0.036
						(0.029)	(0.029)
Constant	0.435^{***}	0.415^{***}	0.432^{***}	0.552^{***}	0.410^{***}	0.400^{***}	0.494^{***}
	(0.014)	(0.020)	(0.015)	(0.020)	(0.017)	(0.020)	(0.030)
N	861	862	862	862	862	862	861

The table shows regressions of students' estimated cheat rates on various measures of other student attributes, where the cheat rate estimated is based only on the dice rolls 21-30. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics *p<0.05; **p<0.01.

Table 22: Estimated cheat rates and other attributes using only dice rolls 31-40

	Estimated cheat rate for dice rolls 31-40						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.003						0.009
	(0.016)						(0.015)
Picks risky lottery		0.038					0.039
		(0.029)					(0.030)
Job security ranked ≤ 2			0.005				0.003
			(0.044)				(0.042)
Donation				-0.017^{***}			-0.016^{***}
				(0.002)			(0.002)
Wage ranked ≤ 2					0.097^{***}		0.063^{**}
					(0.032)		(0.032)
Male						0.061^{**}	0.031
						(0.029)	(0.030)
Constant	0.433^{***}	0.414^{***}	0.432^{***}	0.547^{***}	0.405^{***}	0.400^{***}	0.490^{***}
	(0.015)	(0.020)	(0.016)	(0.021)	(0.017)	(0.021)	(0.030)
N	861	862	862	862	862	862	861

The table shows regressions of students' estimated cheat rates on various measures of other student attributes, where the cheat rate estimated is based only on dice rolls 31-40. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics *p<0.05; **p<0.01.

Table 23: Estimated cheat rates and other attributes while excluding students with dice game experience

	Estimated cheat rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.011						0.016
	(0.014)						(0.014)
Picks risky lottery		0.024					0.028
		(0.027)					(0.028)
Job security rank ≤ 2			0.016				0.013
			(0.040)				(0.038)
Donation				-0.016^{**}			-0.016^{**}
				(0.002)			(0.002)
Wage rank ≤ 2					0.089^{**}		0.057
					(0.030)		(0.029)
Male						0.045	0.020
						(0.027)	(0.028)
Constant	0.409^{**}	0.397^{**}	0.407^{**}	0.517^{**}	0.383^{**}	0.385^{**}	0.472^{**}
	(0.014)	(0.019)	(0.015)	(0.019)	(0.016)	(0.019)	(0.029)
N	821	822	822	822	822	822	821

The table shows regressions of students' estimated cheat rates on various measures of other student attributes while excluding students that explicitly indicated that they were cheating or had prior knowledge of the dice task. The exclusion was based on students responses in an open-ended text box in which they were asked about their impression of the survey and whether they had prior familiarity with any of its elements. The exlcusion is based on an independent coding of the responses. It indicated that 40 students expressed awareness of either dice-under-cup games, similar experimental games (e.g. coin flipping), or explicitly mentioned the potential for cheating. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics. Robust standard errors in parentheses. *p<0.05; **p<0.01.

Table 24: Estimated cheat rates and other attributes while excluding students with 100% win rate

	Estimated cheat rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.002						0.006
	(0.013)						(0.012)
Picks risky lottery		-0.005					0.004
		(0.025)					(0.024)
Job security rank ≤ 2			0.030				0.030
			(0.036)				(0.035)
Donation				-0.014^{**}			-0.013^{**}
				(0.002)			(0.002)
Wage rank ≤ 2					0.100^{**}		0.076^{**}
					(0.027)		(0.027)
Male						0.013	-0.001
						(0.025)	(0.024)
Constant	0.307^{**}	0.309^{**}	0.303^{**}	0.406^{**}	0.278^{**}	0.300^{**}	0.375^{**}
	(0.012)	(0.017)	(0.013)	(0.019)	(0.014)	(0.017)	(0.026)
N	718	719	719	719	719	719	718

The table shows regressions of students' estimated cheat rates on various measures of other student attributes while excluding students who reported a correct guess for all dice rolls. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics *p<0.05; **p<0.01.

A.8 Self-selection Conditional on Attributes, Other Measures

In the main text we analyzed self-selection conditional on other attributes by regressing the estimated cheat rate on our main job preferences measure while including various controls. In particular, we found that the relationship between dishonesty and job preferences drops by 30 percent if we control for dictator game donation and an indicator for whether the wage was ranked as one of the two most important job characteristics. In this section, we examine how these results change if we use alternative job preferences measures.

In Table 25, we regress the estimated cheat rate on various measure of job preferences while controlling for dictator game donations and whether the wage was ranked as important job characteristic (corresponding to Column 8 of Table 5 in the main text). Comparing the estimated coefficients on the job preferences measures to the corresponding estimates without controls (Columns 2-6 of Table 2 in the main text), we see a very similar pattern to the one we found for our main job preference variable. For the first three individual job measures as well as the principal component-based measure, the coefficient drops by between 32 and 36 percent when the controls are added. For the last individual measure, the drop is a bit larger (44 percent) so that the coefficient is no longer significantly different from zero once controls are added ($p = 0.11$).

Table 25: Conditional results using other job measures

	Estimated cheat rate				
	(1)	(2)	(3)	(4)	(5)
Higher ranking of public administration	-0.015^{**}				
	(0.006)				
Public service motivation score		-0.103^{***}			
		(0.027)			
Public sector picked at current wage			-0.058^{**}		
			(0.029)		
Probability of public administration				-0.160	
				(0.101)	
Principal component of all five measures					-0.015^{**}
					(0.006)
Donation	-0.015^{***}	-0.014^{***}	-0.015^{***}	-0.015^{***}	-0.015^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Wage ranked ≤ 2	0.043	0.038	0.048^{*}	0.049^{*}	0.043
	(0.029)	(0.029)	(0.029)	(0.029)	(0.029)
Constant	0.464^{***}	0.758^{***}	0.529^{***}	0.545^{***}	0.513^{***}
	(0.031)	(0.069)	(0.024)	(0.030)	(0.022)
N	862	860	862	858	856

The table shows regressions of students' estimated cheat rates on preference for public service, while controlling for various measures of other student attributes. The job preference measures are the flipped actual rank given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap and the subjective probability of ending up in public administration. The measures of other attributes are the amount donated in the dictator game and an indicators for whether the wage was ranked in the top two of the five job characteristics. Robust standard errors in parentheses. $^{*}p<0.05$; $^{**}p<0.01$.

A.9 Analyzing Representativeness and Selective Non-participation

This section examines potential issues with selective nonparticipation among students invited for participation in our survey experiment. The concern is that students self-select into participation based on particular traits which creates selection bias in our estimates. In our experiment, 862 students ended up participating. Relative to the 3,000 e-mail invitations that was sent out, this yields a response rate of 29 percent.

One strength of our experimental design is that since we sample and invite students from the university registers, we have data also on the characteristics of those who do not participate. Table 26 compares participants to nonparticipants in terms of the available characteristics: field of study, age, gender and study experience as measured by the number of earned ECTS point (European Credit Transfer System). We see clear difference in the participation rate across fields and some moderate systematic differences in other characteristics, with participants being on average younger and more likely to be male than the average nonparticipant. There are no mean differences between the two groups on study experience, although we find evidence of systematic differences in the distribution of the study experience variable.

Table 26: Comparing participants to invited non-participants

	mean participant	mean nonparticipant	diff	t test p value	KS p value
Age	24.128	25.176	-1.049	0.000	0.000
Female	0.466	0.503	-0.037	0.067	-
Study experience (ECTS points)	45.112	44.482	0.630	0.754	0.066
Field: Law	0.182	0.390	-0.207	<0.001	-
Field: Economics	0.445	0.294	0.152	<0.001	-
Field: Politial Science	0.369	0.312	0.057	0.003	-

The table compares the sample of participants in the survey experiment with the sample of invited non-participants using the available data from university records. The available variables are student age, an indicator for the student being female, the students study experience as measured by the earned number of ECTS points (European Credit Transfer System), as well as indicators for field of study. Each row corresponds to a different variable. The first numerical columns shows the variable mean among participants, while the second column shows the mean among non-participants. The third and fourth columns show the difference in means between the groups and the p-value for a t-test that the means are the same. The last column shows the p-values for a Kolmogorov-Smirnoff test that the distributions of the variable is the same across the two groups.

To asses whether our results are driven by selective nonparticipation, we implement a correction based on inverse probability weighting. We estimate a logit model for participation in the experiment across all invitees. We use the six variables in Table 26 as explanatory variables in the logit model. This generates, for each student, a predicted probability of participating in the experiment. We then weight each observation with the inverse of this probability in our regression. To obtain standard errors, we use a bootstrap procedure that resamples the full set of invitees.

Tables 27 through 29 show the results. Throughout, the point estimates are close to those of the unweighted regressions in the main text. Although we can never rule out that there is selection on unobservables, there is little evidence that the results presented in the main text are affected by selective non-participation.

Table 27: Estimated cheat rates and public service preferences with reweighting to correct for non-participation

	(1)	(2)	(3)	(4)	(5)
Public administation ranked ≤ 2	-0.084^{*}				
	(0.034)				
Higher ranking of public administration		-0.018^{*}			
		(0.008)			
Public service motivation score			-0.103^{*}		
			(0.042)		
Public sector picked at current wage				-0.069^{*}	
				(0.034)	
Probability of public administration					-0.173
					(0.153)
Constant	0.418^{**}	0.322^{**}	0.633^{**}	0.403^{**}	0.419^{**}
	(0.027)	(0.025)	(0.114)	(0.023)	(0.043)

The table shows weighted regressions of students' estimated cheat rates on various measures of public service job preferences. The applied weights are the inverse of the predicted participation probability from a logit-model that includes age, an indicator variable for being male, study experience as measured by earned number of ECTS points and indicators for field of study. The job preference measures are an indicator for whether public administration was ranked in the top two of the eight job categories, the flipped actual ranked given to public administration (so that a higher value means a stronger preference for public administration), the public service motivation score, an indicator for whether the public sector was picked in the wage scenario corresponding to the current wage gap and the subjective probability of ending up in public administration. Bootstrapped standard errors are in parentheses. $^{*}p < 0.05$; $^{**}p < 0.01$.

Table 28: Estimated cheat rates and student characteristics with reweighting to correct for non-participation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.035						0.041
	(0.024)						(0.024)
Picks risky lottery		0.045					0.044
		(0.035)					(0.030)
Job security ranked ≤ 2			0.007				0.013
			(0.044)				(0.041)
Donation				-0.014^{**}			-0.014^{**}
				(0.003)			(0.002)
Wage ranked ≤ 2					0.082^{*}		0.046
					(0.035)		(0.032)
Male						0.077^{*}	0.046
						(0.034)	(0.031)
Constant	0.390^{**}	0.363^{**}	0.384^{**}	0.474^{**}	0.361^{**}	0.344^{**}	0.423^{**}
	(0.015)	(0.030)	(0.019)	(0.032)	(0.023)	(0.027)	(0.039)

The table shows weighted regressions of students' estimated cheat rates on various measures of other student attributes. The applied weights are the inverse of the predicted participation probability from a logit-model that includes age, an indicator variable for being male, study experience as measured by earned number of ECTS points and indicators for field of study. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics. Bootstrapped standard errors are in parentheses. $^{*}p < 0.05$; $^{**}p < 0.01$.

Table 29: Preference for public service and student characteristics with reweighting to correct for non-participation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.027						0.029
	(0.025)						(0.028)
Picks risky lottery		-0.043					-0.036
		(0.046)					(0.042)
Job security ranked ≤ 2			-0.007				-0.032
			(0.060)				(0.058)
Donation				0.011^{**}			0.009^{**}
				(0.003)			(0.003)
Wage ranked ≤ 2					-0.178^{**}		-0.163^{**}
					(0.046)		(0.042)
Male						-0.117^{*}	-0.085
						(0.048)	(0.046)
Constant	0.402^{**}	0.419^{**}	0.399^{**}	0.330^{**}	0.449^{**}	0.459^{**}	0.458^{**}
	(0.021)	(0.037)	(0.024)	(0.032)	(0.029)	(0.036)	(0.048)

The table shows weighted regressions of an indicator for students ranking public administration in the top two of the eight job categories on various measures of other student attributes. The applied weights are the inverse of the predicted participation probability from a logit-model that includes age, an indicator variable for being male, study experience as measured by earned number of ECTS points and indicators for field of study. The measures of other attributes are GPA standardized by field, an indicator for choosing the one of the two most risky lotteries, the amount donated in the dictator game, an indicator for being male and indicators for whether job security and wage was ranked in the top two of the five job characteristics. Bootstrapped standard errors are in parentheses. $^{*}p < 0.05$; $^{**}p < 0.01$.

A.10 Translation of the Initial Survey Invitation Mail

Figure 12 shows the translated text from the invitation e-mail for the experiment.

Figure 12: Invitation e-mail for the experiment

```
From: Survey research, University of Copenhagen
Subject: You are invited to a scientific survey from the University of
Copenhagen

Dear [salutation],

We would very much like you to participate in a scientific survey. You will be
paid for participating in the survey.

The survey deals with attitudes towards the world and how you act in situations
characterized by uncertainty. It will only take 20 minutes to complete.

It is very important to us that many people participate so we will be very
grateful if you take the time to participate. For the same reason we have also
chosen to give a high remuneration: The average participant will earn at least
50 DKK but the maximum reward exceeds 300 DKK.

Press this link to participate in the survey: [link]

You will have to provide a user name and password to participate.

Your username is: [username]

Your password is: [password]

If you have any questions or comment regarding the survey then you are welcome
to respond back to this email address ([e-mail]).

Thank for your participation.

Kind regards,

Department of Economics
University of Copenhagen
Oster Farimagsgade 5
Building 26
1353 Copenhagen K
```

A.11 Screencaps and Instructions from Survey Experiment

Below we present translated instructions and screencaps from the survey experiment:

A.11.1 Dice Guessing Game, Instructions and Screenshots

You have now responded to the first series of questions.

The purpose now is to see how you guess in situations marked by randomness. You will play 10 rounds of a dice game in which you can win money in each round.

You have to guess what the die will show. The more die rolls you guess, the more money you will win.

Each round of the game proceeds like this:

1. First, you will have to guess a number of dots from 1 to 6. When you have made your guess you can press the continue bottom.

2. Hereafter a die will roll and you will be asked to report the number of dots which you guessed earlier.

3. The next screen will show the result of the round. If your guess matched the number of dots on the die then you will win 3 DKK, else you will win 1 DKK.

You should avoid using the back bottom during the dice games as it might delete your total winnings.

Note: it is important that you are careful about remembering and reporting the exact number of dots which you guessed prior to rolling the die.

Figure 13: Intro screen to first set of dice games as viewed by the respondent

UNDERSØGELSE BLANDT STUDERENDE VED KØBENHAVNS UNIVERSITET

Du har nu svaret på den første serie af spørgsmål.

Formålet er nu at undersøge hvordan du gætter i situationer præget af tilfældighed. Du skal derfor spille 10 runder af et terningespil, hvor du i hver runde kan vinde pengebeløb.

Du skal gætte, hvad terningen slår. Jo mere rigtigt du gætter, desto større beløb vinder du.

Hver runde af terningespillet foregår således:

1. Først skal du gætte et antal øjne fra 1 til 6. Når du har lagt dig fast på et gæt, trykker du på fortsæt-knappen
2. Der vil herefter blive slået med en terning, og du vil blive bedt om at indtaste det antal øjne, du gættede på tidligere
3. Den næste skærm viser resultatet af runden. Hvis dit tal matcher terningens øjne, vinder du **3 kr**., ellers vinder du **1 kr**.

Under terningsspillene skal du lade være at bruge din browsers “back”/“tilbage” knap da dette kan komme til at slette noget af din gevinst.

Bemærk: det er vigtigt, at du er grundig med at huske og angive netop det tal, du tænkte på før du slog med terningen.

Figure 14: Dice game test screen as viewed by the respondent


Før vi går i gang med terningspillet vil vi gerne høre om reglerne er klare. Vi vil derfor bede dig om at rapportere hvad en person vinder i denne situation

Screenshot

Terningeslaget blev en 2'er.

Hvilket tal tænkte du på? Indtast i feltet:

2

Submit

Rapportér dit svar her

Vi vil også gerne bede dig rapportere hvad en person vinder i denne situation

Screenshot

Terningeslaget blev en 1'er.

Hvilket tal tænkte du på? Indtast i feltet:

4

Submit

Rapportér dit svar og tryk derefter på fortsæt knappen for at gå i gang med terningespillene.

Fortsæt

Figure 15: Intro screen: *Guess a number between 1 and 6. Hereafter, press the bottom below in order to throw with the digital die.*

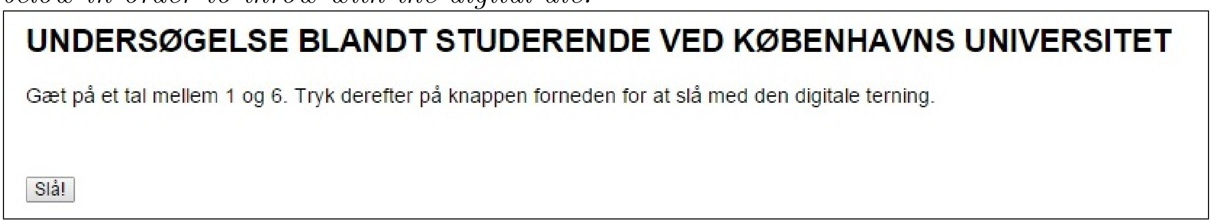
UNDERSØGELSE BLANDT STUDERENDE VED KØBENHAVNS UNIVERSITET

Gæt på et tal mellem 1 og 6. Tryk derefter på knappen forneden for at slå med den digitale terning.

Slå!

Figure 16: Intro screen: Guess report screen (following a three second animation of spinning die): *The die throw was six. Which number did you guess? Please report in the field:*

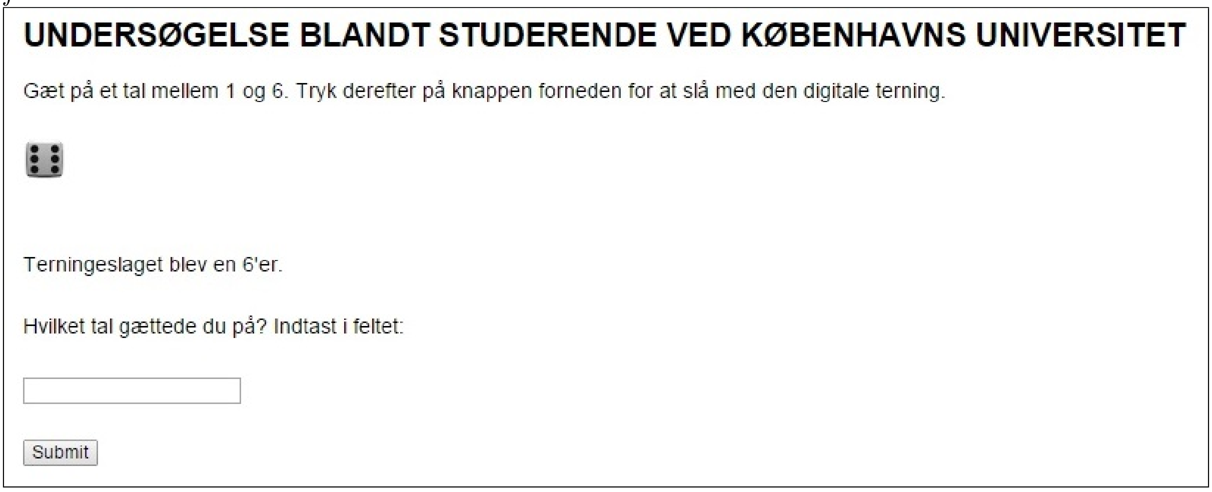

UNDERSØGELSE BLANDT STUDERENDE VED KØBENHAVNS UNIVERSITET

Gæt på et tal mellem 1 og 6. Tryk derefter på knappen forneden for at slå med den digitale terning.

Terningeslaget blev en 6'er.

Hvilket tal gættede du på? Indtast i feltet:

Submit

Figure 17: Intro screen: Payoff screen (in case of wrong guess): *Your guess did not match the die. You win 1 DKK. Your combined winnings in the survey amounts to 16 DKK.*

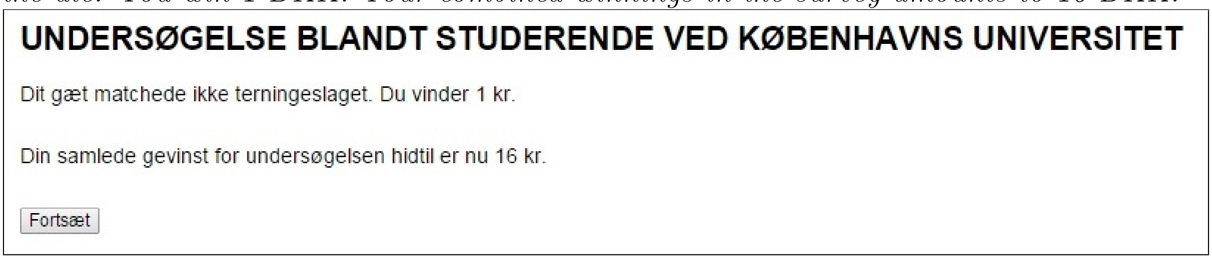

UNDERSØGELSE BLANDT STUDERENDE VED KØBENHAVNS UNIVERSITET

Dit gæt matchede ikke terningeslaget. Du vinder 1 kr.

Din samlede gevinst for undersøgelsen hidtil er nu 16 kr.

Fortsæt

A.11.2 Dictator Game, Instructions and Screenshots

Welcome to the study. Before we proceed, you are given a gift of 15 DKK (2.75 USD) as an appreciation of the time you spend on the survey.

After the survey you will have the option to get this sum automatically transferred to your bank account together with the additional rewards you collect in the survey. But you can also choose to donate some of the money to one of the following charities:

- The Danish Cancer Society (Kræftens Bekæmpelse)
- DanChurchAid (Folkekirkens Nødhjælp)
- Save the Children (Red Barnet)
- Amnesty International
- Red Cross (Røde Kors)

Depending on how much you choose to donate we will additionally donate the amount provided in the below schema of donation options:

	Your donation	Our donation	Total donation
Option A	0 DKK	0 DKK	0 DKK
Option B	5 DKK	3 DKK	8 DKK
Option C	10 DKK	4 DKK	14 DKK
Option D	15 DKK	4 DKK	19 DKK

Which of the donation options do you choose?

- Option A
- Option B
- Option C
- Option D

Figure 18: Donation screen as viewed by the respondent

UNDERSØGELSE BLANDT STUDERENDE VED KØBENHAVNS UNIVERSITET

Velkommen til undersøgelsen. Inden vi går videre modtager du allerede nu en gave på **15 kr.** som tak for at du tager dig tid til at deltage.

Efter undersøgelsen har du mulighed for at få udbetalt denne sum helt automatisk til din NemKonto sammen med de yderligere belønninger, du optjener i løbet af undersøgelsen. Men du kan også vælge at donere nogle af pengene til en af følgende velgørenhedsorganisationer:

- Kræftens Bekæmpelse
- Folkekirkens Nødhjælp
- Red Barnet
- Amnesty International
- Røde Kors

Afhængig af hvor meget du vælger at donere vil vi lægge en yderligere donation oveni som angivet i følgende skema over donationsmuligheder.

	Din donation	Vores donation	Samlet donation
Mulighed A	0 DKK	0 DKK	**0 DKK**
Mulighed B	5 DKK	3 DKK	**8 DKK**
Mulighed C	10 DKK	4 DKK	**14 DKK**
Mulighed D	15 DKK	4 DKK	**19 DKK**

Hvilken donationsmulighed vælger du?

- Mulighed A
- Mulighed B
- Mulighed C
- Mulighed D

Submit

A.11.3 Lottery Choice, Instructions and Screenshots

The survey does, as already mentioned, among other things, deal with your decisions in situations marked by randomness. Among the participants in the study we draw a subset which participate in a simple coin-flip lottery. About one in ten participants will be selected to participate.

If you are selected to participate in the lottery a virtual coin will be flipped and you will win an amount of money depending on if the coin shows heads or tails. You can choose how the reward depends on the coin flip from the list of possible options below:

	Payoff if heads	Payoff if tails
Option A	200 DKK	0 DKK
Option B	160 DKK	30 DKK
Option C	140 DKK	40 DKK
Option D	120 DKK	50 DKK
Option E	80 DKK	80 DKK

Which of the donation options do you choose?

- Option A
- Option B
- Option C
- Option D
- Option E

Please press forward when you have made your choice. You will be informed about if you have been selected to participate in the lottery by the end of the survey.

Figure 19: Lottery screen as viewed by the respondent

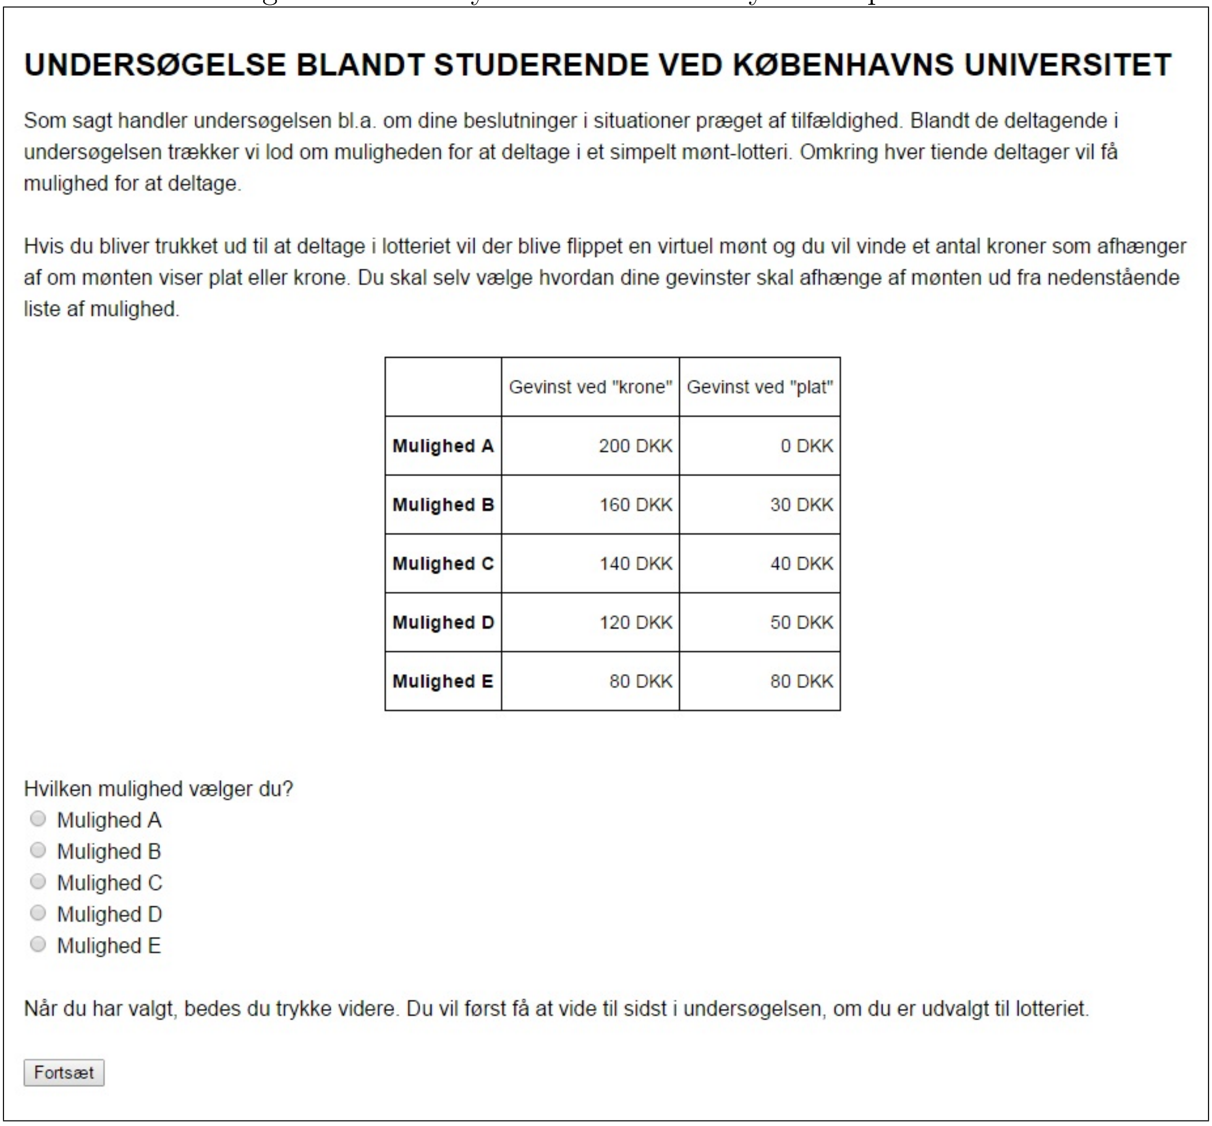

UNDERSØGELSE BLANDT STUDERENDE VED KØBENHAVNS UNIVERSITET

Som sagt handler undersøgelsen bl.a. om dine beslutninger i situationer præget af tilfældighed. Blandt de deltagende i undersøgelsen trækker vi lod om muligheden for at deltage i et simpelt mønt-lotteri. Omkring hver tiende deltager vil få mulighed for at deltage.

Hvis du bliver trukket ud til at deltage i lotteriet vil der blive flippet en virtuel mønt og du vil vinde et antal kroner som afhænger af om mønten viser plat eller krone. Du skal selv vælge hvordan dine gevinster skal afhænge af mønten ud fra nedenstående liste af mulighed.

	Gevinst ved "krone"	Gevinst ved "plat"
Mulighed A	200 DKK	0 DKK
Mulighed B	160 DKK	30 DKK
Mulighed C	140 DKK	40 DKK
Mulighed D	120 DKK	50 DKK
Mulighed E	80 DKK	80 DKK

Hvilken mulighed vælger du?

- Mulighed A
- Mulighed B
- Mulighed C
- Mulighed D
- Mulighed E

Når du har valgt, bedes du trykke videre. Du vil først få at vide til sidst i undersøgelsen, om du er udvalgt til lotteriet.

Fortsæt

Appendix References

Fischbacher, Urs, and Franziska Föllmi-Heusi. 2013. "Lies in Disguise: An Experimental Study on Cheating." *Journal of the European Economic Association* 11 (3): 525–47.

Hanna, Rema, and Shing-Yi Wang. 2017. "Dishonesty and Selection into Public Service: Evidence from India." *American Economic Journal: Economic Policy.*

Jiang, Ting. 2013. "Cheating in Mind Games: The Subtlety of Rules Matters." *Journal of Economic Behavior & Organization* 93: 328–36.

McLachlan, G J. 1987. "On Bootstrapping the Likelihood Ratio Test Stastistic for the Number of Components in a Normal Mixture." *Applied Statistics* 36 (3): 318.