

# Sustaining Honesty in Public Service: The Role of Selection\*

Sebastian Barfort<sup>†</sup>    Nikolaj A. Harmon<sup>‡</sup>    Frederik Hjorth<sup>§</sup>  
Asmus Leth Olsen<sup>¶</sup>

This version: August 2018

First version: September 2015

## Abstract

We study the role of self-selection into public service in sustaining honesty in the public sector. Focusing on the world's least corrupt country, Denmark, we use a survey experiment to document strong self-selection of more honest individuals into public service. This result differs sharply from existing findings from more corrupt settings. Differences in pro-social vs. pecuniary motivation appear central to the observed selection pattern. Dishonest individuals are more pecuniarily motivated and self-select out of public service and into higher-paying private sector jobs. Accordingly, we find that increasing public sector wages would attract more dishonest candidates to public service in Denmark.

---

\*This research was funded by grants from the Faculty of Social Sciences at the University of Copenhagen and from the Economic Policy Research Network. We thank the editor, Matt Notowidigdo, and three anonymous referees for useful comments. We thank Andreas Gotfredsen, Ray Fisman, Rema Hanna, Marco Piovesan, Søren Serritzlew, Peter Bjerre Mortensen, David Broockman, Robert Klemmensen, Alexander Sebald, Lene Holm Pedersen, Donald Green, Jon Fiva, Edward Miguel, Ritwik Banerjee, Frederico Finan, Ernesto Dal Bó and Alex Mas for helpful input and discussions. We also thank seminar participants at the University of California Berkeley, the American Political Science Association Meetings 2015, Copenhagen Business School, the University of California Santa Barbara, Aix-Marseille School of Economics, University of Copenhagen, the Midwest Political Science Association Meetings 2015, the Public Management Research Conference 2015, Stockholm University SOFI, Aarhus University and the 2016 NGO conference at the Paris School of Economics. The paper was previously circulated under the title “Dishonesty and Selection into Public Service in Denmark: Who Runs the World’s Least Corrupt Public Sector?”.

<sup>†</sup>London School of Economics and Political Science, [s.barfort@lse.ac.uk](mailto:s.barfort@lse.ac.uk).

<sup>‡</sup>University of Copenhagen, [nikolaj.harmon@econ.ku.dk](mailto:nikolaj.harmon@econ.ku.dk) (corresponding author, mailing address: Nikolaj A. Harmon, University of Copenhagen, Øster Farimagsgade 5, building 26, 1353 Copenhagen K, Denmark).

<sup>§</sup>University of Copenhagen, [fh@ifs.ku.dk](mailto:fh@ifs.ku.dk).

<sup>¶</sup>University of Copenhagen, [ajlo@ifs.ku.dk](mailto:ajlo@ifs.ku.dk).

# 1 Introduction

Research on corruption has tended to emphasize formal differences in individual incentives for misuse of public office, emphasizing monitoring and punishment as deterrents from engaging in corrupt behavior. While this focus has been very fruitful (see Olken and Pande 2012 for a recent survey), recent evidence have suggested that individual attributes such as cultural values may also play a prominent role (Fisman et al. 2015; Fisman and Miguel 2007).

This paper explores the role of individual selection in generating an equilibrium of honesty and low corruption in public service. Using Denmark as a low-corruption case study, we ask whether potential candidates for public service jobs differ in their inherent propensity for dishonest behavior, and if so whether systematic self-selection of honest types into public service may be one channel that helps sustain a low level of corruption. To draw lessons for combatting corruption in other settings, we further ask how the observed selection pattern is related to other individual attributes, as well as the level of public sector wages.

Theory provides ambiguous predictions regarding the questions we pose. The inherent propensity for dishonesty could differ significantly across potential candidates for public service or could be relatively constant within a country. Moreover, even if dishonesty does vary across potential public service candidates, it is unclear how dishonesty should relate to preferences for entering public service. On the one hand, the relatively low level of public sector corruption in Denmark could discourage dishonest individuals from entering this sector. On the other hand, the Danish public sector is not immune to rent extraction (Amore and Bennedsen 2013), and the sheer size of public budgets in Denmark means that even small scale rent extraction in the public sector may be very lucrative for dishonest individuals. Finally, to the extent that dishonesty correlates with other individual attributes that shape job preferences, such as risk aversion or pro-social motivation, this may further complicate the observed selection pattern.

To provide empirical guidance on these questions, we conduct a survey experiment with students in the fields of law, economics, and political science at the University of

Copenhagen in Denmark. Given its consistent ranking as the least corrupt country in the world, Denmark is a useful benchmark for studying how countries can sustain low levels of corruption. For studying selection into public service, the particular population of students we focus on is very well suited. They face a very clear choice between public service and private sector careers and make up an important part of the public sector workforce.

We first examine the extent of heterogeneity in dishonesty and how this heterogeneity is related to preferences for entering public service. We adopt the experimental methodology of Hanna and Wang (2017) and subject students to a standard set of cheating tasks building on Fischbacher and Föllmi-Heusi (2013). In our implementation of the tasks, students can win money by correctly guessing the outcome of a series of dice rolls but are allowed to see the outcome of each roll before reporting their guess. Students therefore have the option of winning dishonestly by misreporting their guess, knowing that it can never be proven whether in fact they were dishonest. Comparing the distribution of successful guesses in the dice game to the expected distribution without lying, however, allows us to construct estimates of individual propensities for dishonesty.

The cheating tasks reveal extensive heterogeneity among potential candidates for public service. We estimate that 14-17 percent of respondents barely cheat at all, while 17-23 percent cheat practically all the time. The remaining respondents fall somewhere in between, resulting in a standard deviation of cheat rates across individuals of 0.39. Relating dishonesty to job preferences, we find clear evidence of positive self-selection into public service, as honest individuals in Denmark are systematically more likely to want to enter public service. Students ranking public administration as one of their top two job choices cheat 10 percentage points less than other students.

To shed some some light on why this selection pattern exists, we next examine how dishonesty and job preferences correlate with other attributes. We find no evidence that ability or risk preferences correlate with dishonesty. Differences in pro-social vs. pecuniary motivation, however, turn out to be strong predictors of both dishonesty and job preferences. Pro-social individuals who donate more in a dictator game are both more honest and

more likely to prefer a public service career. Conversely, pecuniarily motivated individuals who view the wage level as a particularly important job characteristic are less honest and less likely to prefer a public service career. This can explain a significant part of the observed selection pattern; controlling for two simple measures of pro-sociality and pecuniary motivation reduces the relationship between dishonesty and preferences for public service by 30 percent.

Finally, we examine how the observed selection pattern is related to the level of public sector wages. Public sector jobs in Denmark are characterized by relatively low wages compared to the private sector. Combined with our findings regarding pro-sociality and pecuniary motivation, this suggests that the observed selection pattern may reflect that more pecuniarily motivated dishonest individuals self-select out of the Danish public sector due to its relatively low wage level. We provide evidence in favor of this hypothesis by analyzing a set of counterfactual job preference question that ask students to choose between a job in the public and the private sector given different counterfactual relative wage levels. These show that increases in the level of public sector wages would attract more dishonest candidates to public service.

The idea that individuals may differ in their inherent propensity for dishonesty has a long tradition in the theoretical literature on corruption (Lui 1986; Cadot 1987; Andvig and Moene 1990) and is supported empirically by the fact that personality traits predict corrupt behavior (Callen et al. 2015). The role of selection on the dishonesty dimension has also received theoretical attention (Caselli and Morelli 2004; Besley 2004; Bernheim and Kartik 2014). In particular, our finding that higher public sector wages attract more dishonest candidates mirror the theoretical predictions regarding the effect of politician salaries in the influential work of Besley (2004). Our result that the observed selection pattern is related to differences in pro-social vs. pecuniary motivation also directly relates to the literature on job choice and extrinsic vs. intrinsic motivation (Benabou and Tirole 2003; Besley and Ghatak 2005; Delfgaauw and Dur 2007).

Empirically, a number of recent papers have examined selection into public service.<sup>1</sup>

---

<sup>1</sup>Dal Bó, Finan, and Rossi (2013), Ashraf, Bandiera, and Lee (2016), and Deserranno (2016) use field experiments to examine how pecuniary incentives affect selection into public service jobs in Mexico,

However, only a handful of empirical papers have focused on dishonesty. Closest to the present paper are Hanna and Wang (2017), who use the same experimental methodology to show that dishonest university students are more likely to want to enter public service in India. Similarly, Banerjee, Baul, and Rosenblat (2015) run a corruption experiment at two different Indian universities and find more dishonest behavior at the university targeting public service careers. Finally, Alatas et al. (2009) find no correlation between preferences for working in the public sector and bribing behavior in an explicit corruption game among Indonesian students.

To our knowledge, this paper is the first to examine dishonesty and selection into public service in a low-corruption environment and is also the first to document a positive selection pattern where honest individuals systematically prefer public service. Additionally, our paper provides the first evidence that the level of public sector wages may impact the selection of honest and dishonest individuals into public service.

Besides the literature on corruption and selection into public service, our paper builds on a larger experimental literature exploring the nature and causes of dishonest behavior across societies (e.g. Gino, Ayal, and Ariely 2009; Shalvi et al. 2011; Weisel and Shalvi 2015; Gächter and Schulz 2016; Houser et al. 2016). Besides Fischbacher and Föllmi-Heusi (2013) and Hanna and Wang (2017), the design of our experimental dishonesty task draws particularly on Jiang (2013). Finally, by using experimental methods to study questions specifically related to corruption, our paper also relates to the experimental literature on corruption or bribery games (see Abbink and Serra 2012 for a recent survey).

The paper proceeds as follows. In Section 2, we lay out the context and aim of the study. In Section 3 we present the survey experiment used to construct the key variables of the study. In Section 4 we present the main results regarding dishonesty and selection into public service. In Section 5 we present additional results regarding the mechanisms behind the observed selection pattern and consider robustness. Section 6 concludes.

---

Zambia and Uganda in various dimensions, including ability and pro-social preferences. Combining survey and experimental data, Kolstad and Lindkvist (2013) and Serra, Serneels, and Barr (2011) examine how pro-social preferences correlate with wanting to work in the public sector in Tanzania and with working in the non-profit sector in Ethiopia. Buurman et al. (2012) examine whether public sector employees in the Netherlands differ in risk preference or their level of altruism. Finan, Olken, and Pande (2015) provides a broader survey.

## 2 Context, Framework, and Aim of the Study

The present paper is motivated by the juxtaposition of two simple observations: The first observation is that some countries seem to be perennial high performers when it comes to maintaining a low level of corruption. While many countries struggle with high level of corruptions, some countries have been consistently able to sustain an equilibrium of honesty and low corruption in public service.

The second observation is that differences in individuals' inherent propensity for dishonesty may play a role in shaping corruption. Whether a public official engages in corrupt behavior will depend on many institutional features of the environment he faces such as monitoring and punishment schemes and the size of the potential gains. In addition to these factors, however, it may also be influenced by his own inherent propensity for engaging in dishonest behavior. For a given level of monitoring and other factors, a public sector that primarily attracts inherently honest employees may thus exhibit lower levels of fraudulent behavior than one that attracts more dishonest individuals.

Put together these two observations raise the question of whether systematic selection of honest individuals into public service may be one reason that some countries are able to maintain persistent low levels of corruption. With this hypothesis in mind, the main aim of our analysis is to examine whether there is systematic self-selection of honest (or dishonest) individuals into public service in one of the world's least corrupt countries, Denmark. Additionally, we aim to shed some light on possible mechanisms that explain the observed selection pattern.

In the rest of this section, we first lay out the conceptual framework for our analysis and then describe the particular empirical context we examine.

### 2.1 Dishonesty: Conceptual Framework

At the heart of this paper is the idea that when faced with the same exact situation, some individuals may be more likely to be dishonest than others. The way we conceptualize individual dishonesty is a simple application of this idea: Throughout the paper when we

say that some individual A is inherently more dishonest (or have a higher propensity for dishonesty) than some individual B, we simply mean that in a given situation, individual A has a higher likelihood of engaging in dishonest behavior than B.

There are several things to note about this definition: First, the definition above does not take a stand on *why* some individuals are more dishonest. For example, individual A may be more likely to engage in dishonest behavior because dishonesty is an important fundamental personality trait for him, or because his particular risk preferences, his level of pro-sociality or his other characteristics makes dishonesty more attractive or palatable. In either case we may worry that individual A would engage in more corrupt behavior during a potential public service career. Accordingly, for the main part of the paper, we examine the unconditional relationship between dishonesty and job preferences without controlling for any other variables that may explain or correlate with dishonest behavior. When we turn to shedding light on the mechanisms behind the observed selection pattern, however, we do examine specifically how dishonesty correlates with other individual variables. As we shall see, the evidence suggest that dishonesty is particularly strongly correlated with differences in pro-social vs pecuniary motivation.

Second, the definition above implicitly assumes that individual dishonesty is relatively stable across settings. If the ranking of individuals in terms of dishonesty is wildly different in different settings, it makes little sense to talk of differences in overall dishonesty. As we return to in Section 3, however, previous work has documented that individuals who behave more dishonestly in our experimental dishonesty task also tend to behave more dishonestly in a range of other settings. This suggests that individual dishonesty does indeed exhibit stability across settings.

Finally, there are a number of different theories for why individual dishonesty (as defined above) may be positively or negatively correlated with preferences for a public service career. We return to these in Section 5, after we have presented our main result regarding the overall selection pattern.

## 2.2 Study Setting: Denmark

Most research on corruption focuses on high-corruption settings and aims to understand why corruption is so prevalent or how it is affected by policy changes. This paper takes the opposite approach. We focus on a benchmark country that has been successful in maintaining a consistently low level of corruption and try to shed light on how this low corruption equilibrium is being sustained.

Accordingly, the setting of the study is Denmark. For studying how to sustain honesty in public service, Denmark is a natural benchmark country given its consistent ranking among the very least corrupt countries in the world. Figure 1 shows the levels of corruption in different countries 1996-2014 as measured by the commonly used Transparency International's Corruption Perceptions Index (CPI), with Denmark highlighted (Transparency International 2016). Since 2007 Denmark has ranked as the least corrupt country in the CPI every year but two, and in the history of the CPI Denmark has never ranked lower than fourth.<sup>2</sup>

## 2.3 Study Population

Within Denmark, the population we study consists of university students in the fields of law, economics, and political science. Higher education is highly specialized in Denmark and practically all university students complete both a bachelor and a master's degree in their chosen field of study, which in turn strongly influences the jobs available to them upon graduation.

For the purpose of our study, students in law, economics, and political science are ideal for two reasons: First, these students face a very clear choice between entering the private sector or going into public service. For current employees with a background in economics, law, or political science, around 46 percent work in the public sector, mostly in public administration, and 54 percent in the private sector, typically in finance, law firms, and lobbying organizations. Second, this population is large and influential enough to

---

<sup>2</sup>This pattern is not exclusive to the CPI. For example, the World Bank Governance Indicator "Control of Corruption", detailed in Kaufmann, Kraay, and Mastruzzi (2010), has ranked Denmark as the least corrupt country in the world every year since 2007 and never ranked Denmark lower than second.



actually affect the corruption level of the public sector. About 60 percent of all state-level employees in administrative functions have a background in one of the three fields we study or a closely related field. They are also dominant at the top level of the public sector: 100 percent of current deputy secretaries and about 40 percent of members of parliament hold a degree in one of the three fields.

## **2.4 Possibilities for Corruption**

Given the motivation for our study, it is worth considering what types of corrupt behavior our student population can undertake in their public service careers. For those entering public administration at the local level, many of them will be engaged in direct administrative work that affect individual citizens and businesses and offers the opportunity of bribe-taking or other forms of misbehavior. Amore and Bennedsen (2013), for example, have documented rent extraction in the procurement of local public services in Denmark. In other instances, potential corruption could take a more indirect form. Some graduates work in offices which help develop and prepare legislative input to elected officials. These may be influenced into serving the interests of private companies or other organizations. Additionally, our student population may also engage in smaller scale dishonesty at their workplace such as fraudulent absenteeism or the abuse of expense accounts.

As discussed in the previous section, it is by all accounts rare that any Danish public employee engages in corrupt behavior. However, this does not imply that there is no corruption in Denmark or that it is impossible for public sector employees to take advantage of the corruption opportunities described above. A recent corruption scandal involving the IT company Atea showcases how dishonesty, in the form of bribe-taking, can be found in rare but serious cases throughout the Danish administrative system. Since 2015 Danish authorities and state prosecution have indicted multiple state, regional, and local public employees for accepting bribes from Atea. The bribes have been given with the aim of providing Atea with a more favorable bidding position in public procurement auctions. The indictments have up until this point lead to more than 10 convictions of higher-level public employees from many branches of government including the Ministry of Foreign

Affairs, The Region of Southern Denmark, the Municipality of Copenhagen, and the Danish state railways. The bribes were mostly given in the form of gifts, such as paid trips or discounts on personal IT equipment, and ranged in value from 1,000 USDs to 50,000 USDs. The indicted have for the most part been employed in positions and organizations that are dominated by employees with the same educational background as the students in our sample.

In sum, graduates from our student population face serious opportunities for corrupt behavior in their public service careers. As a reflection of the sustained low corruption level in Denmark, however, they rarely take advantage of these opportunities. One explanation for this is likely the monitoring schemes and institutional setup of the Danish public sector. Another, complementary explanation may be however, that graduates who actually choose a public service career tend to be systematically very honest individuals who often refrain from corrupt behavior even in cases where they could get away with it. The next section explains the data we use to examine dishonesty and selection into public service within our student population.

### 3 Data and Experimental Design

Our empirical analysis is based on an online survey experiment conducted at the University of Copenhagen during December 2014. The university administration provided us with complete lists of everyone who enrolled as undergraduates in law, economics, and political science, including student e-mail addresses. From these lists random samples of 1,000 students who enrolled over the years 2009-2011 and 2013-2014<sup>3</sup> were drawn from each of the three fields and were invited to participate in the survey experiment.<sup>4</sup> The invitation to participate was sent as an e-mail with a link to the survey along with a username and password.<sup>5</sup> Participants were told that the survey dealt with their attitudes to various

---

<sup>3</sup>Students who enrolled prior to 2009 were not invited as many of them have graduated and therefore may no longer use their student e-mail. Pilot studies were run on students enrolling in 2012, so these were not invited so as to not contaminate the subject pool.

<sup>4</sup>A translated version of the invitation mail can be found in Appendix A.10.

<sup>5</sup>The experiment was run using a software called “ILab” developed by Andreas Gotfredsen and Alexander Sebald.

topics and “how they acted in situations characterized by uncertainty.” The latter referred to the various incentivized games which they would encounter in the survey and which will be outlined in detail below. Participants were also told that they would be paid to participate. In accordance with the actual outcomes, participants were informed that the average participant would earn no less than 50 DKK (8 USD), that the maximum payoff was above 300 DKK (50 USD), and that the survey would take approximately 20 minutes to complete.<sup>6</sup> For comparison, the student population in question would in a typical student job usually receive a union-defined hourly wage of about 110 DKK (18 USD), corresponding to 37 DKK (6 USD) per 20 minutes.

In total 863 students completed the survey. From these we drop one individual who experienced technical difficulties during the main dishonesty experiment in the survey, leaving us with a base sample of 862 participants. In terms of representativeness, our sampling scheme by definition implies that the pool of invitees is statistically representative within each field of study. At the end of Section 5 and in Appendix A.9, we examine potential issues related to selective non-participation by exploiting the availability of administrative university data for non-participants.

### 3.1 Measuring Dishonesty: Experimental Dice Task

The first main purpose of our survey experiment is to measure respondents’ inherent propensities for dishonesty. We follow Hanna and Wang (2017) and measure dishonesty using a repeated variation of the *dice-under-cup* game approach from Fischbacher and Föllmi-Heusi (2013). Behavior in these types of games have become a widely used measure of dishonesty (see for example Shalvi, Eldar, and Bereby-Meyer 2012; Gneezy, Rockenbach, and Serra-Garcia 2013; Hilbig and Hessler 2013; Cohn, Fehr, and Maréchal 2014; Ariely et al. 2014) and has been shown to predict real-world dishonest behavior and rule breaking in a range of settings (Cohn, Maréchal, and Noll 2015; Cohn and Maréchal 2015). Given that the present paper is motivated by understanding public sector corruption, we note

---

<sup>6</sup>These announcements were based on observed outcomes in pilot studies and ended up closely mirroring actual outcomes. The realized maximum payoff for a participant was 315 DKK (53 USD) and the average payoff was 80 DKK (13 USD). The median time from opening the survey to completion was 25 minutes. At the time of the survey experiment 1 USD equaled about 6 DKK.

in particular that dishonesty in dice under cup games has been shown to predict actual fraudulent behavior among public sector employees (Hanna and Wang 2017).

For our specific implementation of the dice under cup approach, we build on the computer-based variation of Jiang (2013).<sup>7</sup> Our implementation works as follows:<sup>8</sup> At four different points in the survey experiment, participants were asked to play ten rounds of a dice guessing game. Students were told that the game was intended to test how they “guess in situations characterized by randomness” and that they could win money in the game by correctly guessing the outcome of a dice roll. In each round of the dice game respondents were first asked to think of a number between 1 and 6 that they expected the dice to show after the dice roll. Students then clicked “next” while keeping their guess in mind. A dice was rolled on screen and the outcome of the dice roll was reported. The participants were then asked to report their guess while the actual outcome of the dice roll was still displayed. On the following screen the payoff from the round was reported. Reporting a correct guess yielded a gain of 2 DKK (0.33 USD) relative to an incorrect one.<sup>9</sup>

The point of the dice guessing game is that in each round, respondents have the option of winning dishonestly by reporting the actual outcome of the dice roll regardless of what their initial guess was. Moreover, a strength of the design is that respondents are not explicitly primed to think about dishonesty and respondents know in each round that it can never be revealed whether in fact they reported their guess honestly.<sup>10</sup> Because an honest individual always has a one-in-six chance of correctly guessing each dice roll, however, we can make statistical statements about individual dishonesty after observing more

---

<sup>7</sup>Our motivation for using the computer-based implementation is that it can be conducted online. Using an online implementation allowed us to systematically sample and invite participations directly from the university e-mail database, while simultaneously keeping participation costs low and ensuring as high a participation rate as possible. Both of these features are important for dealing with the issue of sample representativeness, which is particularly critical given that our focus is on estimating the relationship between job preferences and dishonesty in the underlying population.

<sup>8</sup>Screen caps of the game as viewed by the respondents, including exact translations of all instructions for the game, are presented in Appendix A.11.

<sup>9</sup>In our pilot studies, we explicitly tested whether behavior depended on the level of payoffs or gains and found no evidence of stakes-dependency in our setting.

<sup>10</sup>One may still worry that upon realizing that they can lie undetected in the game, students implicitly feel that being dishonest is the point of the game. In an attempt to mitigate this type of experimental demand, we concluded the introduction screen by stating that: “it is important that you are careful about remembering and reporting the exact number on which you guessed prior to rolling the die”.

repetitions of the game. The next section formalizes this and presents the individual-level estimate of dishonesty that we use in our analysis.

### 3.2 Measuring Dishonesty: Econometrics

The data from our dice under cup task consists of a sample of  $N$  respondents, which we index by  $i$ . Each respondent participates in a series of  $K$  rounds of a dice game, which we index by  $k$ . As described above, our experiment has  $N = 862$  and  $K = 40$ . In each round the respondent can either win or lose. The rounds are independent of each other with a constant probability of winning of  $p^*$ . In our experiment the probability of (truthfully) guessing a dice roll is one in six so  $p^* = \frac{1}{6}$  in our case. We do not directly observe whether respondents win or lose however. For each round and each respondent, we instead observe a self-reported measure of whether the respondent won or not, where respondents are free to report dishonestly. We let  $y_{ik}$  be an indicator variable for whether respondent  $i$  reported winning in round  $k$ . In the context of our implementation of the dice guessing game,  $y_{ik}$  is simply an indicator for whether the reported guess matches the actual dice roll. We let  $Y_i = \sum_{k=1}^K y_{ik}$  denote the total number of wins (total number of correct guesses) reported by respondent  $i$ .

We introduce heterogeneity in the propensity for being dishonest by assuming that when reporting the outcome of a given round, respondent  $i$  reports dishonestly some fraction  $\theta_i \in [0, 1]$  of the time.<sup>11</sup> We further make the assumption that if reporting dishonestly, a respondent reports a win for sure in that round. Otherwise he or she reports the truth. The individual-specific  $\theta_i$  therefore captures respondent  $i$ 's propensity for dishonesty and we refer to it as respondent  $i$ 's *cheat rate*.

In order to examine the relationship between cheat rates, job preferences and other attributes, we will construct individual measures of each respondent's cheat rate. The probability of observing a win for a respondent with a given cheat rate,  $\theta_i$ , is

---

<sup>11</sup>Note that here we do not take stance on whether cheating behavior exhibits dependence over time. In the construction of our cheat rate estimator, we only require weak stationarity so that the (unconditional) probability of cheating (and thus of reporting a win) is the same in any given round. In the appendix, however, when deriving the variance of our cheat rate estimator and developing our estimator of the full distribution of dishonesty, we do it under the assumption that cheating behavior is independent over time for a given individual.

$P(y_{ik} = 1 | \theta_i) = p^* + (1 - p^*)\theta_i$ . From this we can construct an unbiased Method of Moments estimator of  $i$ 's cheat rate by replacing population moments with empirical moments and rearranging:<sup>12</sup>

$$\widehat{CheatRate}_i = \frac{6}{5} \cdot \left( \frac{Y_i}{40} - \frac{1}{6} \right)$$

We refer to this as individual  $i$ 's *estimated cheat rate* and use this as our main measure of individuals' propensity for dishonesty.

Three properties of the measure are worth highlighting here:

First, the inherent randomness in whether dice rolls match respondents' (honest) guesses implies that the estimated cheat rate suffers from measurement error. As we show in Section A.1 in the appendix, however, this measurement error is classical. Using the estimated cheat rate as the outcome variable in a regression therefore provides consistent estimates of the relationship between the true cheat rate and the included covariates under the usual conditions.

Second, despite not affecting the consistency of our estimates, the measurement error in the estimated cheat rate does lower the precision of our estimates. In Section A.1 of the appendix we invoke simple assumptions on the time dependence of dishonest behavior to show that the amount of measurement error is decreasing in the number of dice guessing rounds that each person plays and increasing in the true probability of a correct guess. This motivates our chosen implementation of the dice guessing game which has many rounds and a low win probability in each round.

Third, because our estimated cheat rate is a linear transformation of the total number of reported wins, alternative approaches that use the raw number of correct guesses (or total winnings) as measures of dishonesty would simply rescale the regression results we present later.

Given the properties above, the main analysis of the paper involves regressing respon-

---

<sup>12</sup>Unbiasedness is easily seen from  $E(\widehat{CheatRate}_i | \theta_i) = \frac{1}{1-p^*} \frac{1}{K} \sum_{k=1}^K P(y_{ik} = 1 | \theta_i) - \frac{p^*}{1-p^*} = \theta_i$ . It is worth noting that the estimated cheat rate 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 not be unbiased and are therefore unattractive for our purposes (see section A.1 in the appendix).

dents' estimated cheat rates on other characteristics, in particular preferences for entering public service. We note, however, that one can also use the data from our experimental dice task to estimate the full distribution of dishonesty in the population, as well as how this distribution relates to job preferences. In Sections A.1 and A.2 of the appendix we discuss, develop and implement such estimators in detail.

### 3.3 Measuring Job Preferences

The second key variable in the empirical analysis is the students' preferences for public service jobs. For our main measure of job preferences, we asked students to imagine that they have obtained their academic degree and are now free to choose between jobs. In this scenario they were then asked to rank eight categories based on the most common jobs held by graduates from our student population: public administration, private sector job in the financial sector, private sector job in a political party or lobby organization, private sector job within public relations, private sector job in a law firm, a job in the Danish Central Bank, other public sector job, or other private sector job. These particular categories were chosen to match the eight most common industry categories for our study population in the official Danish employment statistics. As noted, public administration is by far the most important public service career for our population. For our main measure of students' preferences for entering public service, we therefore focus on the rank given to public administration.

For robustness, we also elicited additional measures of job preferences. In one question, we asked students to report the likelihood of them ending up in each of the eight job categories described above.<sup>13</sup> In addition, we administered a standard 16-item questionnaire measuring Public Service Motivation (PSM), which is often used as an indication of respondents' dispositional preferences for working in the public sector (Perry and Wise 1990).<sup>14</sup>

---

<sup>13</sup>To ease students' way through the survey, we did not require that the reported probabilities sum to a hundred. In the empirical analysis we rescale them appropriately.

<sup>14</sup>Given the setup of our survey, we note that all our job preference measures are likely to center on students preferences for the first job out of university. They may therefore miss potential differences in long-term career plans.

Finally, for examining the role played by the level of public sector wages, we subjected all respondents to nine different counterfactual wage scenarios. In each scenario respondents were asked to choose between their preferred job in the private sector and their preferred job in the public sector given a particular wage gap between the two jobs ranging from the private sector job paying 20,000 DKK (3,300 USD) more per month to the public sector job paying 20,000 DKK more per month.

### 3.4 Additional Measures

To examine how dishonesty and self-selection into public service are related to other student attributes, we included a range of other standard experimental tasks and questions in the survey experiment. At the beginning of the survey experiment, we asked respondents to play a simple dictator game. Respondents were given a gift of 15 DKK (2.5 USD).<sup>15</sup> They were then offered to get the money transferred to their account when the survey was finished or donate some or all of the money to one of five charities of their choice. Furthermore, as they increased their own donation we matched their donation amount with up to 4 DKK (0.75 USD) using a concave matching schedule.

We also included an incentivized measure of risk aversion at the beginning of the survey. Students were told that one in ten of them would be randomly selected to enter into a coin flip lottery at the end of the survey. They were then asked to choose between five different such lotteries with varying risk profiles.<sup>16</sup>

As a proxy of ability, we asked students to report their high school GPA. High school exams are standardized nationally in Denmark and provide a good measure of ability for the population we study. In the empirical analysis, we standardize GPAs across field to avoid mechanical correlations stemming from the admissions cut-offs for the different fields.<sup>17</sup>

---

<sup>15</sup>In pilot studies we experimented with the placement of the dictator game but found no evidence that the timing of the dictator game mattered for dishonesty behavior or public sector preferences.

<sup>16</sup>The lotteries were designed based on the range of constant relative risk aversion coefficients reported for the Danish population in Harrison, Lau, and Rutström (2007). The most risky coin lottery involved a gain of 200 DKK (33 USD) in case of heads and 0 DKK for tails. The least risky lottery involved a gain of 80 DKK (16 USD) regardless of the coin flip.

<sup>17</sup>Admission to different fields in Danish higher education is based high school GPA, with the necessary GPA varying widely across different fields. This introduces strong mechanical differences in student GPAs



To get direct measures of students' preferences regarding job characteristics, we asked them to rank the following five job characteristics in order of importance: wage level, work hours and other terms of work, importance, entertainment value, and job security.

Finally, we use data on the students' gender. Table 1 provides summary statistics. As the table shows, a few of the observations lack information about some variables. These are caused by erroneous reporting and a few students experiencing technical issues during parts of the survey experiment.

## 4 Main Results: Dishonesty and Self-selection

We start our empirical investigation by examining the variation in dishonest behavior in our dice guessing experiment. Figure 2 shows the distribution of the observed number of correct guesses across students in our dice guessing experiment, along with the distribution of correct guesses that would be expected if students report their guesses completely honestly. Comparing the expected honest distribution with the actual outcomes in the experiment, we see evidence of extensive dishonest behavior. For example, the probability of an honest respondent having 10 or more correct guesses is about 12 percent, yet 73 percent of respondents report 10 or more correct guesses in our sample.

The figure also suggests that the amount of dishonest behavior differs very significantly across individuals. While many students' report a number of correct guesses well above the 99th percentile of the honest distribution, other students actually report fewer correct guesses than what would be expected under full honesty.<sup>18</sup> In Sections A.1 and A.2 of the appendix, we show how the data in Figure 2 can be used to construct estimators of the full distribution of dishonesty across individuals if one imposes simple assumptions on the time dependence of dishonesty. Applying such estimators we find that 14-17 percent of individuals in our data are practically completely honest and cheat less than 2 percent of the time, while 17-23 percent are practically completely dishonest and cheat more than 98 percent of the time. The remaining respondents are spread fairly evenly in between,

---

across fields, which are unrelated to their own career preferences.

<sup>18</sup>Under full honesty the 99th percentile is 13 correct correct guesses, while the expected number of correct guesses is 6.7.

and the standard deviation of cheat rates across individuals is 0.39.<sup>19</sup> Despite facing the same opportunities and incentives to behave dishonestly in the survey experiment, we thus see extensive heterogeneity in dishonesty within our pool of potential candidates for public service.

## 4.1 Dishonesty and Self-selection into Public Service

Next we turn to the main focus of the paper and examine whether the observed differences in the propensity for dishonesty are correlated with preferences for a public service career. We do this in the context of a simple regression that relates individual  $i$ 's estimated cheat rate to an indicator for whether individual  $i$  prefers a public service career,  $PublicService_i$ . Here and throughout the paper, we use the estimated cheat rate as the outcome variable in the regression to deal with the measurement error in our cheat rate estimate:<sup>20</sup>

$$\widehat{CheatRate}_i = \beta_0 + \beta_1 PublicService_i + \varepsilon_i$$

The estimate of  $\beta_1$  in this regression estimates the average gap in individual cheat rates between individuals that have a preference for a public service career and those who do not.

Table 2 shows estimates of this regression using different measures of job preferences. Column 1 focuses on our main measure of job preferences: whether students rank public administration in the top two of the eight job categories described in Section 3. The estimated coefficient on the indicator for job preferences is -0.10 and is highly significant, suggesting that more honest individuals are systematically more likely to prefer a public service career. In Figure 3 we provide a transparent illustration of this main result by comparing the average cheat rate between students who rank public administration among the top two jobs and those who do not. Students ranking public administration in the

<sup>19</sup>This is relative to a mean of 0.42. The result that many respondents cheat a little bit but not the full amount is a standard finding in these types of dice games (Fischbacher and Föllmi-Heusi 2013; Hilbig and Hessler 2013; Shalvi, Handgraaf, and De Dreu 2011).

<sup>20</sup> $\widehat{CheatRate}_i$  is equal to individual  $i$ 's true cheat rate plus classical measurement error. If we place in on the left hand side of the regression, we can therefore ignore the fact that we are using the estimated cheat rate instead of the actual cheat rate. See Section 3 and Section A.1 in the appendix.

top two cheat 36 percent of the time on average, while other students cheat 46 percent of the time.

Columns 2 to 5 of Table 2 examine the robustness of the result to using other measures of job preferences. In Column 2 we replace the indicator variable from Column 1 with the flipped actual rank given to public administration (so a higher value means a stronger preference for public service). In Column 3 we use the measured PSM score. In Column 4 we use data from our counterfactual wage question, focusing on whether the student would choose the public sector over the private sector if faced with a sectoral wage gap of 5,000 DKK (833 USD) per month, corresponding to the typical gap in starting wages between the two sectors. Finally, in Column 5 we include the students' reported probability of entering public administration. Across all these measures we see a negative and highly significant correlation between cheat rates and expressing a preference for entering public service. At the end of Section 5 and in the appendix, we show that this result is robust to a wide range of other checks on the specification.

Next, we examine whether the magnitude of the observed correlation depends on the exact measure of job preference we use. The bottom row of Table 2 rescales each of the estimated coefficients so that they reflect the change in cheat rate associated with a one standard deviation increase in the relevant job preference measure. We see that the relationship between dishonesty and job preferences is relatively stable across the different measures. Across all job preference measures, a one standard deviation increase in preferences for public service is associated with a decrease in the estimated cheat rate of between 4 and 8 percentage points.

Finally, we can examine whether the observed selection pattern is driven by particularly strong job preferences among students with a certain level of dishonesty or whether it reflects differences in job preferences across throughout the distribution of dishonesty. To answer this question, Sections A.1 and A.2 of the appendix construct estimators for the entire joint distribution of dishonesty and job preferences. Applying such estimators we find that the observed selection pattern reflects particularly strong job preferences for

public service among the very least dishonest students.<sup>21</sup>

In sum, we find a clear pattern of positive self-selection into public service in Denmark, as more honest individuals systematically tend to prefer public service jobs. In our baseline specification find that students ranking public administration among the top two job options cheat 10 percentage points less than other students. Relative to the mean cheat rate of 0.42 this represents a 24 percent relative gap in cheat rates.

## 5 Additional Results: Mechanism and Robustness

In the previous section, we saw that there is systematic self-selection of more honest individuals into public service in Denmark. Next, we try to shed light on why this selection pattern exists.

Many different factors could contribute to the observed selection pattern. If more honest individuals tend to stand out in terms of other attributes or preferences, the particular job characteristics offered in public service may systematically attract these individuals. For example, we might imagine that being dishonest is correlated with lower levels of risk aversion and that people with different risk preferences tend to be attracted to systematically different types of jobs.

Alternatively, if dishonest individuals are attracted by the opportunity to profit from corrupt behavior, the currently perceived scope for public sector dishonesty may influence selection patterns.<sup>22</sup> This could imply that negative or positive selection patterns may be self-reinforcing. It could also imply that the observed selection patterns respond to the level of monitoring and punishment in the public sector as dishonest individuals opt out of public service when the opportunities for public sector corruption diminishes.

Given that our survey experiment does not yield any variation in the perceived corruption level or the level of monitoring and punishment, we are unable to explore the

---

<sup>21</sup>Among students who virtually never cheat, 53 percent prefer a public service career. Moving up the dishonesty distribution, this figure drops rapidly to just under 40 percent among students who cheat about a third of the time. In the upper part of the distribution (from students who cheat 50 percent of the time and up), however, the share preferring public service remains fairly constant around 35 percent.

<sup>22</sup>Corbacho et al. (2016) have shown that such a self-reinforcing effect is present when individuals are deciding whether to engage in corruption.

role played by these factors. Since our survey experiment measured a range of additional attributes, we can however examine whether honest individuals stand out in terms of other attributes and whether this appears to play a role in shaping job preferences. We can also examine the role played by public sector wages using the counterfactual wage questions.

## 5.1 Correlates of Dishonesty and Job Preferences

We start our investigation by asking whether dishonest individuals and/or individuals preferring public service careers stand out in terms of other attributes. In doing so, we focus on four key attributes which ex ante seem likely to correlate both with dishonesty and job preferences: ability, risk aversion, pro-social vs. pecuniary motivation, and gender.

We again use a simple regression framework to examine the correlation between dishonesty, job preferences, and these additional attributes. For each attribute, we regress the individual estimated cheat rate and the indicator for preferring a public service career on a measure of the attribute in question,  $Attribute_i$ :

$$\widehat{CheatRate}_i = \gamma_0 + \gamma_1 Attribute_i + e_i$$

$$PublicService_i = \eta_0 + \eta_1 Attribute_i + u_i$$

In these regressions,  $\gamma_1$  shows how the attribute correlates with dishonesty, while  $\eta_1$  shows how it correlates with job preferences. Table 3 shows the results. Panel A of the tables shows the regressions using the estimated cheat rate as the outcome variable, while Panel B shows the regressions using the indicator for preferring a public service career as the outcome variable.

Column 1 of the table examines how dishonesty and job preferences correlate with ability, as measured by GPA. We see no evidence that ability correlates with dishonesty or job preferences in our data. In both panels, the estimated coefficient on GPA is close to zero and statistically insignificant.

In Column 2 and 3, we examine risk aversion. In Column 2 we focus on our incentivized

risk aversion measure and include an indicator for whether the student chose one of the two most risky lotteries offered.<sup>23</sup> In Column 3 we instead include an indicator for whether the student ranked job security among the two most important job characteristics. Panel A shows no statistically significant correlation between dishonesty and either of the two risk preference measures. In Panel B, there is some indications that risk averse individuals prefer public service. In Column 2, we see that individuals who chose one of the risky lotteries are 5.9 percentage points less likely to prefer public service and this difference is marginally significant ( $p = 0.09$ ), however the estimated coefficient on valuing job security in Column 3 points in the opposite direction and is insignificant.<sup>24</sup>

Columns 4 and 5 look at differences in pro-sociality and pecuniary motivation, as measured by donations in the dictator game and whether individuals ranked wage among the two most important job characteristics. These turn out to be strong predictors for both dishonesty and job preferences. In Column 4, we see that each additional DKK donated in the dictator game is associated with a 1.6 percentage decrease in the cheat rate and a 0.9 percentage point increase in the likelihood of preferring public service.<sup>25</sup> Conversely, in Column 5, we see that individuals who rank the wage as an important job characteristics cheat 8.3 percentage points more and are 20 percentage points less likely to prefer public service. All of these differences are highly statistically significant.

Column 6 looks at gender. We see that men cheat 6.1 percentage points more than women and are 13 percentage points less likely to prefer public service. Both differences are statistically significant.

Finally in Column 7 we include all six measures in the regressions simultaneously. In Panel B, dictator game donation, gender, and the importance of the wage level remain

---

<sup>23</sup>50 percent of students in our sample chose one of the two most risky lotteries so the simple indicator measure summarizes most of the variation in risk aversion in our sample. We get similar results if we use the risk rank of the chosen lottery, the implied coefficient of relative risk aversion or dummies for each of the lotteries.

<sup>24</sup>If risk averse individuals rank job security as important and prefer public service we would expect the coefficient in Column 3 to show a *positive* association between preferences for public service and ranking job security as important.

<sup>25</sup>The highest possible donation was 15 DKK so the estimates imply that an individual who donates the maximum amounts cheats 24 percentage points less and is 14 percentage points more likely to prefer public service than an individual who donates nothing. 33 percent of students choose the maximum possible donation of 15 DKK, while 40 percent choose to donate nothing.

statistically significant predictors of job preferences, and the estimates are very similar to those reported in the previous columns. In Panel A, dictator game donation continues to show up as a strong predictor of dishonesty, while the coefficient on ranking wage among the two most important job characteristics drops a bit and becomes only marginally significant ( $p = 0.10$ ). The coefficient on gender drops even more, however, and becomes insignificant. We interpret this as evidence that the relationship between gender and dishonesty is working mostly through gender differences in pro-sociality and pecuniary motivation.

## 5.2 Self-selection Conditional on Attributes

The results in the preceding section suggest that systematic selection of honest individuals into public service may be related to differences in pro-social vs. pecuniary motivation. Pro-social individuals who make large donations in the dictator game are systematically more honest and more likely to prefer a public service career. Conversely, pecuniarily motivated individuals that rank the wage level as an important job characteristic are systematically less honest and less likely to prefer a public service career.

These patterns offer an intuitive explanation for the observed selection pattern: As we return to further below, public sector wages in Denmark tend to be systematically lower than in the private sector, suggesting that non-pecuniary motivations are important for entering public service in Denmark. On the other hand, a main motivation for dishonest behavior is - in our lab experiment as well as in real world settings - that it offers material gains. This in turn suggests that dishonesty should be more prevalent among pecuniarily motivated individuals.<sup>26</sup>

To further explore this explanation we can examine selection into public service conditional on the different attributes in our data. In particular, we include the different measures from Table 3 as controls in the regression of estimated cheat rate on job preferences:

$$\widehat{CheatRate}_i = \pi_0 + \pi_1 PublicService_i + \pi_2 X_i + \nu_i$$

---

<sup>26</sup>Indeed, some studies in personality psychology even find that dishonesty and greediness can be treated as part of the same fundamental personality trait (see for example Ashton and Lee (2007)).

In this regression,  $\pi_1$  captures the relationship between job preferences and dishonesty after conditioning on the attribute  $X_i$ . To the extent that the observed self-selection of honest individuals into public service is driven by one or more other attributes, the estimate of  $\pi_1$  should decrease when the attribute(s) are added as controls.

Table 4 shows the results. Across Columns 1 to 3 we add the different measures for ability and risk aversion as controls. As would be expected given the results in Table 3, none of these controls affect the estimated relationship between dishonesty and preferences for public service. The coefficient on the indicator for ranking public service among the two most attractive jobs is -0.10 in all three columns. This is the same as in the specification without controls in Table 2,

Columns 4 and 5 add the measures of pro-sociality and pecuniary motivation to the regression. This reduces the estimated coefficient on the indicator for preferring a public service career. Controlling for donations in the dictator game lowers the estimated coefficient on job preferences to -0.08, while controlling for whether individuals ranked wage as an important job characteristic reduces the coefficient to -0.09. As shown at the bottom of the table these differences in the estimated coefficients are statistically significant. Column 6 adds gender as a control. This also lowers the estimated coefficient on job preferences slightly, although this difference is only marginally significant ( $p = 0.10$ ).

Finally in Column 7, we control for all the different measures simultaneously. After conditioning on all the measures, the coefficient on job preferences is -0.07. As shown in Column 8, this change is driven entirely by controlling for dictator game donations and the indicator for ranking wage as an important job characteristic.<sup>27</sup>

Overall, we conclude that systematic differences in pro-social vs. pecuniary motivation can explain a significant part of the observed selection pattern. Conditioning only on our two simple measures of pro-sociality and pecuniary motivation reduces the cheat rate gap between students with a preference for public service and other students by 30 percent.<sup>28</sup>

---

<sup>27</sup>In the appendix we reexamine this result using the alternative measures of job preferences in our data. Across all measures we see that the coefficient on preferences for public service drops when dictator game donations and the indicator for ranking wage as an important job characteristic are added as controls. The drop in the coefficient is slightly larger when we use the alternative job preference measures (the drop is between 32 and 44 percent for these other measures).

<sup>28</sup>The fact that there is still a significant cheat rate gap after conditioning on these measures is suggestive



Conversely, none of the other attributes we examine appear important for the observed selection pattern.

### 5.3 The Role of Public Sector Wages

Next we focus on the role played by the level of public sector wages in shaping selection into public service. The results in the preceding sections have interesting implications for the effect of public sector wages on selection. If dishonest individuals tend to be motivated by pecuniary incentives while honest individuals tend to be motivated more by pro-social concerns, we might expect high public sector wages to affect selection by systematically attracting more dishonest individuals to public service.

Two features of our institutional setting lend support to this idea: First, public sector wages in Denmark tend to be systematically lower than private sector wages. For the population we study, entry level wages in the private sector are typically around 5,000 DKK (833 USD) higher per month. This stands in stark contrast to the considerable public sector wage premiums that are typical in many developing countries struggling with corruption (Finan, Olken, and Pande 2015). Second, in Section A.3 in the appendix we analyze the job preferences of dishonest individuals in our sample and document that dishonest individuals are particularly likely to want a job in the financial sector, which for our student population stand out as by far the best paid job category. This is indicative that the positive selection pattern we observe is driven in part by dishonest individuals being more pecuniarily motivated and self-selecting out of public service jobs and into higher-paid private sector jobs.<sup>29</sup>

---

that other factors also play a role in shaping the observed selection pattern. It could however also reflect the simplicity of our measures of pro-sociality and pecuniary motivation. Our measures are based on behavior in a single dictator game and a single question regarding the ranking of job preferences, which may not perfectly capture all underlying differences in pro-sociality and pecuniary motivation.

<sup>29</sup>Public sector jobs may stand out from private sector jobs in other dimensions than the wage level. In the Danish context, another salient difference between public and private sector jobs is that public sector contracts tend to be more family-friendly. For example, while all parents are entitled to 11 months parental leave at partial pay, public sector employees typically receive full pay for a larger fraction of the leave period. Public sector jobs are also traditionally viewed as offering a lower unemployment risk, although it is somewhat unclear whether this is a salient difference, especially for the population we study. Formal employment protection is fairly similar between public sector and private sector contracts in Denmark and our student population already face a very low unemployment risk upon entering the labor market.

To provide a more direct test of how public sector wages relate to the observed selection pattern, we use data from our set of counterfactual wage gap questions. As described in Section 3, these questions ask students to choose between their preferred private and public sector jobs conditional on the two jobs having different wage gaps. From the answers to these questions and the individuals' estimated cheat rates, we can examine how changes in the public-private wage gap would affect the selection of honest and dishonest individuals into public service.

Panel A of Figure 4 shows the results from the counterfactual wage gap questions. Each pair of lines in the panel correspond to a different hypothetical wage gap scenario, ranging from the private sector paying 20,000 DKK less per month (3,333 USD) to the private sector paying 20,000 DKK more. For each wage gap, the height of the lines shows the average estimated cheat rate among those who would prefer the public and private sector at the given wage gap.

Furthest to the right, in the scenario where the private sector job pays 20,000 DKK more, the average estimated cheat rate among students preferring the private sector is 0.43 as opposed to only 0.31 among students preferring the public sector, a gap of 12 percentage points. Moving left to scenarios where the public sector wage is relatively higher, the average cheat rates in the two groups begin to converge. In the scenario where the private sector pays 5,000 DKK, roughly the current level of the public-private wage gap for our student population, the gap in cheat rates is down to 9.0 percentage points. Moving further left, the pattern continues. As the relative public sector wage is increased, the average cheat rate increases among public sector candidates and the public-private gap in dishonesty narrows. It eventually flips in scenarios where the public sector wage is 10,000 DKK or more above the private sector wage.

To understand the quantitative relationship between the private-public wage gap and the cheat rate gap, Panel B plots the value of the implied private-public cheat rate gap in the different wage scenarios against the value of the private-public wage gap in of these each scenarios (error bars show 95 percent confidence intervals for the cheat rate gaps in each scenario). Mirroring the conclusion from Panel A, we see a clear upward

sloping relationship between the points in the figure. The solid line in the figure adds the corresponding regression line. Its slope implies that a relative increase in private sector wages of 1,000 DKK per month (166 USD) would increase the average gap in cheat rates between candidates for the private vs public sector by 0.4 percentage points (with a standard error of 0.1).

Looking at the plot and regression line, we note that the leftmost and rightmost points in the plot correspond to quite extreme wage scenarios and also involve very imprecisely estimated cheat rate gaps.<sup>30</sup> To check that these are not driving the results, the dotted line therefore shows how the regression line changes if we exclude these two points. We see that the line actually becomes slightly steeper; a 1,000 DKK relative increase in private sector wages is now estimated to increase the gap in cheat rates by 0.6 percentage points (with a standard error of 0.2).

To the extent that students' answers in the hypothetical wage scenarios reflect actual preferences, these results suggest that higher public sector wages would lead to a more dishonest pool of candidates for public service jobs. This supports the notion that the relatively low level of public sector wages in Denmark is important for the observed selection pattern.

## 5.4 Robustness Checks

We finish this section of the paper by summarizing some additional results and robustness checks that are presented at length in the appendix.

Our experimental approach to measuring dishonesty has been widely used in the literature and has been validated to predict fraudulent behavior among public sector employees by Hanna and Wang (2017). As always however, differences in the exact experimental implementation may be a concern when comparing results to existing papers or relying on past validations. In Section A.4 in the appendix we compare the data from our survey experiment with data from the closely related experiment of Hanna and Wang

---

<sup>30</sup>The imprecisely estimated cheat rate gaps in these scenarios reflect that when the wage gap becomes very large, relatively few students select the lower paying sector. This implies that the average cheat rate for those preferring the lower paying sector is imprecisely estimated.

(2017). We see remarkably similar correlations between dishonesty and other inherent attributes across the two data sets, suggesting that the differences in experimental implementation do not affect the measurement of dishonesty. This mirrors previous conclusions in Hilbig and Zettler (2015).<sup>31</sup>

Another potential concern with our empirical analysis is whether our survey based job preference measures successfully capture the actual job preferences that students express upon graduating and entering the labor market. In Section A.5 in the appendix, we examine this concern using administrative data on actual post-graduation job outcomes for a subset of our sample. Using the most recent administrative data available to us, we can examine the 155 students from our sample that had graduated and entered the labor market by the end of 2017. The job preferences measures from our survey turn out to be strong (and statistically significant) predictors of actual job outcomes. If we regress the estimated cheat rate on actual job outcomes as opposed to stated job preference measures we also get virtually the same estimated coefficient as in our main analysis, although in the much smaller sample, none of these are statistically significant. Overall, the administrative data suggest that the job preference measures from our survey capture actual job preferences well.

In Section A.7 of the appendix, we also conduct a series of robustness checks to shore up various concerns with our empirical analysis: To address concerns that the many repetitions in our dice game have made respondents fatigued or otherwise affected their behavior, we try only using data from different subsets of the dice rolls in our experiment, including using only the first dice roll for each respondent.<sup>32</sup> To assess concerns that some respondents may be affected by knowledge of the existing academic literature on dishonesty and its relation to our experimental tasks, we try dropping respondents that indicate awareness of experimental dishonesty games. To assess concerns that our results are driven only by extreme cheaters, we try excluding respondents who guess correctly in

---

<sup>31</sup>Hilbig and Zettler (2015) compare a survey-based measure of dishonesty with behavior in a range of different variations of the basic *dice-under-cup* game, including a dice-guessing similar to the one used here. They find very similar correlations across all the different implementations of the game.

<sup>32</sup>Our estimated cheat rate remains an unbiased (although quite noisy) estimate of the true cheat rate even when we only use data on a single roll for each respondent. As a result, we can still use it to estimate differences in dishonesty between different groups of respondents (Houser, Vetter, and Winter 2012).

every round of our dice guessing game. Our results are robust to all the above mentioned alternative sample restrictions.

Finally, as usual when analyzing survey or experimental data, representativeness and selective non-participation is a concern. In Section A.9 in the appendix, we examine issues of non-participation by exploiting that the administrative university data contains information on enrollment year, field, completed classes, and gender for everyone invited to our survey experiment. Although our participation rate of 29 percent is reasonably high, our participant population does differ somewhat from invited non-participants. Participants are a bit younger, more likely to study economics, and slightly more likely to be male. Although selection in terms of unobservables can never be ruled out, we apply reweighing procedures that correct non-participation and find no evidence that selective non-response affects our results.

## 6 Conclusion

We study the role of self-selection into public service in sustaining an equilibrium of low corruption and low public sector dishonesty. Focusing on the world's least corrupt country, Denmark, we conduct a survey experiment among a relevant student population to obtain individual measures of dishonesty, preferences for entering public service, and other relevant attributes.

We document extensive heterogeneity in dishonesty among potential candidates for public service and a clear pattern of positive self-selection: Students expressing a preference for entering public service cheat 10 percentage point less in a standard experimental dishonesty task. This result stands in sharp contrast to previous results from more corrupt countries.

To shed some light on the mechanisms behind the observed selection pattern, we examine whether dishonesty and job preferences correlate systematically with other attributes. Differences in pro-social vs. pecuniary motivation turn out to be strong predictors of both dishonesty and job preferences. Pro-socially motivated students who

make large donations in a dictator game are systematically more honest and more likely to prefer a public service career. On the other hand, pecuniarily motivated students that rank the wage level as an important job characteristic are systematically less honest and less likely to express a preference for a public service career. We find that this pattern can explain a significant part of the association between honesty and preferences for public service.

Finally, we examine the role of public sector wages in shaping the observed selection pattern based on a set of counterfactual job preference question that vary the wage gap between the public and private sector. Consistent with the results regarding pro-social vs. pecuniary motivation, we find that higher public sector wages would attract more dishonest candidates to public service in Denmark.

Overall, our results confirm that systematic selection of honest individuals into public service may be part of the reason that Denmark is able to maintain its low levels of corruption and public sector dishonesty. To the extent that current levels of public sector dishonesty affect the future career choices of honest and dishonest individuals, this suggests that Denmark may be benefitting from a virtuous cycle where low levels of corruption and the self-selection of honest individuals into public service are mutually reinforcing. Such virtuous cycles can explain why some countries are consistently able to sustain an honest public sector, while many other countries struggle with high levels of corruption.

At the same time however, our results regarding public sector wages suggest that it is possible to change the observed selection pattern by changing policy. In fact, our results suggest that that the standard policy recommendation of combatting corruption by increasing public sector wages may have unintended negative effects on selection. The implication of this is not that countries struggling with high levels of corruption should simply start lowering public sector wages; the effect of changes in public sector wages is not necessarily the same across high and low corruption settings, and high public sector wages may still have large beneficial effects on the incentives for corruption if they raise the cost of being fired for corruption or if they are necessary to keep public employees' incomes above subsistence levels. At the same time, however, our results do suggest that

high public sector wages is *not* the reason Denmark has been and continues to be among the world's least corrupt countries. Understanding the factors and policy choices that affect selection into public service should thus be a key objective for future research.

## References

- Abbink, Klaus, and Danila Serra. 2012. "Anticorruption Policies: Lessons from the Lab." *Working Paper*, 77–115.
- Alatas, Vivi, Lisa Cameron, Ananish Chaudhuri, Nisvan Erkal, and Lata Gangadharan. 2009. "Subject Pool Effects in a Corruption Experiment: A Comparison of Indonesian Public Servants and Indonesian Students." *Experimental Economics* 12 (1): 113–32.
- Amore, Mario Daniele, and Morten Bennedsen. 2013. "The Value of Local Political Connections in a Low-Corruption Environment." *Journal of Financial Economics* 110 (2): 387–402.
- Andvig, Jens Chr, and Karl Ove Moene. 1990. "How Corruption May Corrupt." *Journal of Economic Behavior & Organization* 13 (1): 63–76.
- Ariely, Dan, Ximena Garcia-Rada, Lars Hornuf, and Heather Mann. 2014. "The (True) Legacy of Two Really Existing Economic Systems." *Munich Discussion Paper*.
- Ashraf, Nava, Oriana Bandiera, and Scott S Lee. 2016. *Do-Gooders and Go-Getters: Career Incentives, Selection, and Performance in Public Service Delivery*. Working paper.
- Ashton, Michael C., and Kibeom Lee. 2007. "Empirical, Theoretical, and Practical Advantages of the HEXACO Model of Personality Structure." *Personality and Social Psychology Review* 11 (2): 150–66.
- Banerjee, Ritwik, Tushi Baul, and Tanya Rosenblat. 2015. "On Self Selection of the Corrupt into the Public Sector." *Economics Letters* 127: 43–46.
- Benabou, Roland, and Jean Tirole. 2003. "Intrinsic and Extrinsic Motivation." *The Review of Economic Studies* 70 (3): 489–520.
- Bernheim, B. Douglas, and Navin Kartik. 2014. "Candidates, Character, and Corruption." *American Economic Journal: Microeconomics* 5 (2): 205–46.
- Besley, Timothy. 2004. "Joseph Schumpeter Lecture: Paying Politicians: Theory and Evidence." *Journal of the European Economic Association* 2: 193–215.
- Besley, Timothy, and Maitreesh Ghatak. 2005. "Competition and Incentives with Motivated Agents." *The American Economic Review* 95 (3): 616–36.
- Buurman, Margaretha, Josse Delfgaauw, Robert Dur, and Seth Van den Bossche. 2012.



“Public Sector Employees: Risk Averse and Altruistic?” *Journal of Economic Behavior & Organization* 83 (3): 279–91.

Cadot, Olivier. 1987. “Corruption as a Gamble.” *Journal of Public Economics* 33 (2): 223–44.

Callen, Michael, Saad Gulzar, Ali Hasanain, Yasir Khan, and Arman Rezaee. 2015. *Personalities and Public Sector Performance: Evidence from a Health Experiment in Pakistan*. HKS Faculty Research Working Paper Series.

Caselli, Francesco, and Massimo Morelli. 2004. “Bad Politicians.” *Journal of Public Economics* 88 (3): 759–82.

Cohn, Alain, and Michel André Maréchal. 2015. *Laboratory Measure of Cheating Predicts Misbehavior at School*. University of Zurich, Department of Economics, Working Paper Series.

Cohn, Alain, Ernst Fehr, and Michel André Maréchal. 2014. “Business Culture and Dishonesty in the Banking Industry.” *Nature* 516 (7529): 86–89.

Cohn, Alain, Michel André Maréchal, and Thomas Noll. 2015. “Bad Boys: How Criminal Identity Salience Affects Rule Violation.” *The Review of Economic Studies* 82 (4): 1289–1308.

Corbacho, Ana, Daniel W Gingerich, Virginia Oliveros, and Mauricio Ruiz-Vega. 2016. “Corruption as a Self-Fulfilling Prophecy: Evidence from a Survey Experiment in Costa Rica.” *American Journal of Political Science* 60 (4): 1077–92.

Dal Bó, Ernesto, Frederico Finan, and Martín A Rossi. 2013. “Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service.” *The Quarterly Journal of Economics* 128 (3): 1169–1218.

Delfgaauw, Josse, and Robert Dur. 2007. “Incentives and Workers’ Motivation in the Public Sector\*.” *The Economic Journal* 118 (525): 171–91.

Deserranno, Erika. 2016. *Financial Incentives as Signals: Experimental Evidence from the Recruitment of Health Promoters in Uganda*. Working paper.

Finan, Frederico, Benjamin A Olken, and Rohini Pande. 2015. “The Personnel Economics of the State.” *NBER Working Paper*.

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.

Fisman, Raymond, and Edward Miguel. 2007. "Corruption, norms, and legal enforcement: Evidence from diplomatic parking tickets." *Journal of Political Economy* 115 (6): 1020–48.

Fisman, Raymond, Nikolaj A Harmon, Emir Kamenica, and Inger Munk. 2015. "Labor Supply of Politicians." *Journal of the European Economic Association* 13 (5): 871–905.

Gächter, Simon, and Jonathan F Schulz. 2016. "Intrinsic Honesty and the Prevalence of Rule Violations Across Societies." *Nature* 531: 496–9.

Gino, Francesca, Shahar Ayal, and Dan Ariely. 2009. "Contagion and Differentiation in Unethical Behavior the Effect of One Bad Apple on the Barrel." *Psychological Science* 20 (3): 393–98.

Gneezy, Uri, Bettina Rockenbach, and Marta Serra-Garcia. 2013. "Measuring Lying Aversion." *Journal of Economic Behavior & Organization* 93: 293–300.

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

Harrison, Glenn W, Morten I Lau, and E Elisabet Rutström. 2007. "Estimating Risk Attitudes in Denmark: A Field Experiment." *Scandinavian Journal of Economics* 109 (2): 341–68.

Hilbig, Benjamin E, and Corinna M Hessler. 2013. "What Lies Beneath: How the Distance Between Truth and Lie Drives Dishonesty." *Journal of Experimental Social Psychology* 49 (2): 263–66.

Hilbig, Benjamin E, and Ingo Zettler. 2015. "When the Cat's Away, Some Mice Will Play: A Basic Trait Account of Dishonest Behavior." *Journal of Research in Personality* 57: 72–88.

Houser, Daniel, John A List, Marco Piovesan, Anya Samek, and Joachim Winter. 2016. "Dishonesty: From parents to children." *European Economic Review* 82: 242–54.

Houser, Daniel, Stefan Vetter, and Joachim Winter. 2012. "Fairness and Cheating."

*European Economic Review* 56 (8): 1645–55.

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

Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi. 2010. “The Worldwide Governance Indicators: A Summary of Methodology.” *Data and Analytical Issues, World Bank Policy Research Working Paper*, no. 5430.

Kolstad, Julie Riise, and Ida Lindkvist. 2013. “Pro-Social Preferences and Self-Selection into the Public Health Sector: Evidence from an Economic Experiment.” *Health Policy and Planning* 28 (3): 320–27.

Lui, Francis T. 1986. “A Dynamic Model of Corruption Deterrence.” *Journal of Public Economics* 31 (2): 215–36.

Olken, Benjamin A, and Rohini Pande. 2012. “Corruption in Developing Countries.” *Annu. Rev. Econ.* 4 (1): 479–509.

Perry, James L, and Lois Recascino Wise. 1990. “The Motivational Bases of Public Service.” *Public Administration Review* 50 (3): 367–73.

Serra, Danila, Pieter Serneels, and Abigail Barr. 2011. “Intrinsic Motivations and the Non-Profit Health Sector: Evidence from Ethiopia.” *Personality and Individual Differences* 51 (3): 309–14.

Shalvi, Shaul, Jason Dana, Michel JJ Handgraaf, and Carsten KW De Dreu. 2011. “Justified Ethicality: Observing Desired Counterfactuals Modifies Ethical Perceptions and Behavior.” *Organizational Behavior and Human Decision Processes* 115 (2): 181–90.

Shalvi, Shaul, Ori Eldar, and Yoella Bereby-Meyer. 2012. “Honesty Requires Time (and Lack of Justifications).” *Psychological Science* 23 (10): 1264–70.

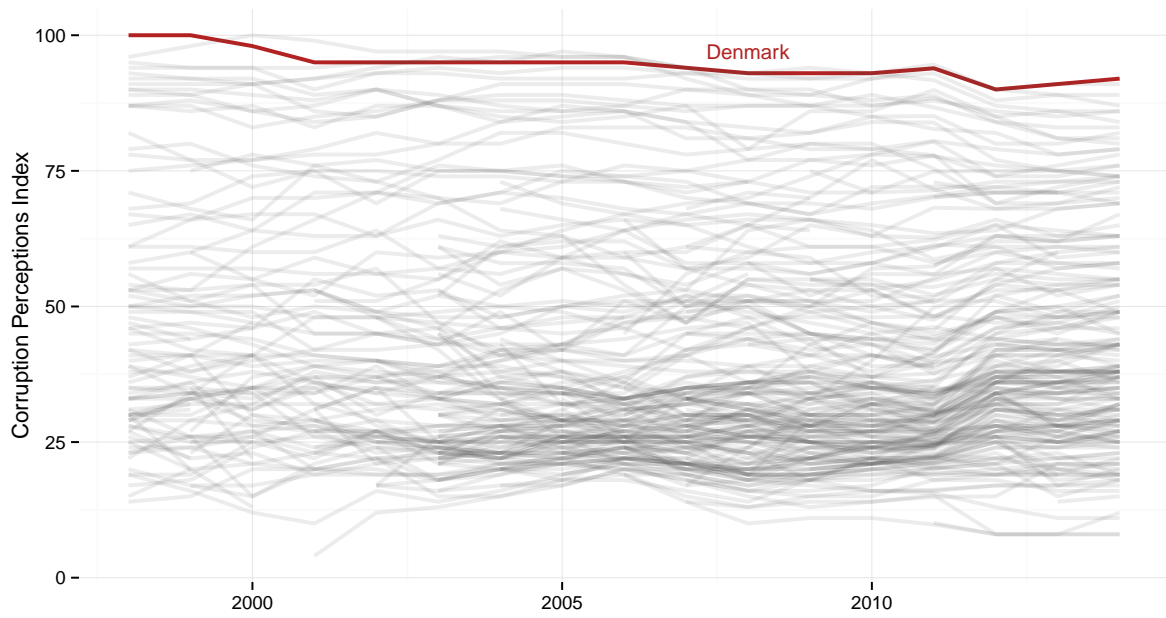
Shalvi, Shaul, Michel JJ Handgraaf, and Carsten KW De Dreu. 2011. “Ethical Manoeuvring: Why People Avoid Both Major and Minor Lies.” *British Journal of Management* 22 (s1): S16–27.

Transparency International. 2016. *Corruption Perceptions Index 2015*. Transparency International. Retrieved from: <https://www.transparency.org/cpi2015/results>.

Weisel, Ori, and Shaul Shalvi. 2015. “The Collaborative Roots of Corruption.”

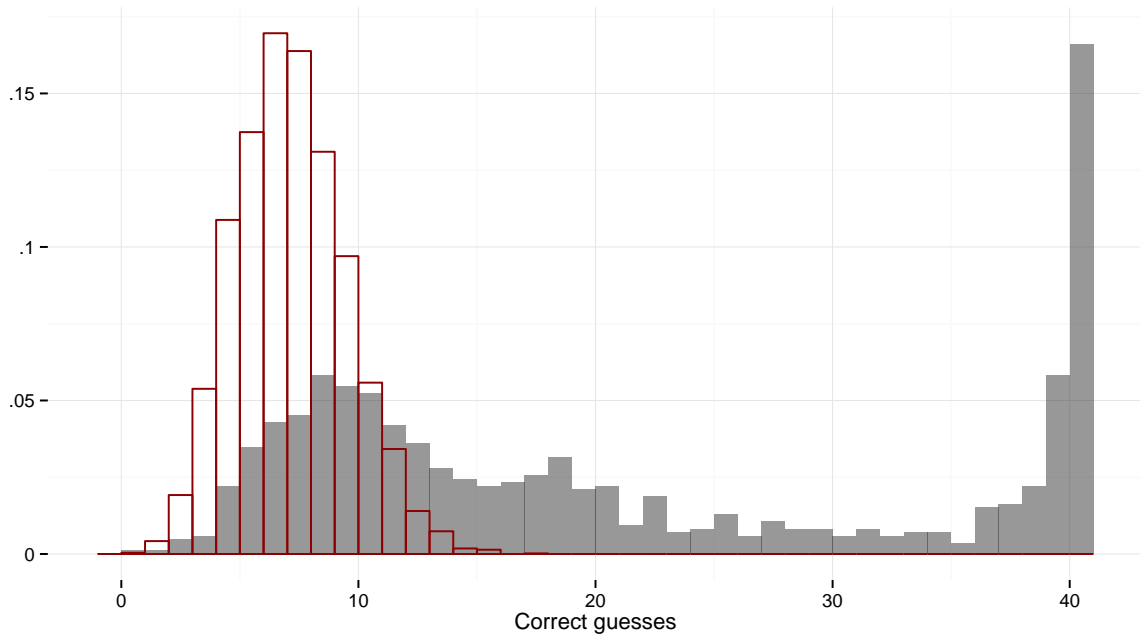
*Proceedings of the National Academy of Sciences* 112 (34): 10651–56.

Figure 1: Corruption across countries 1996-2014



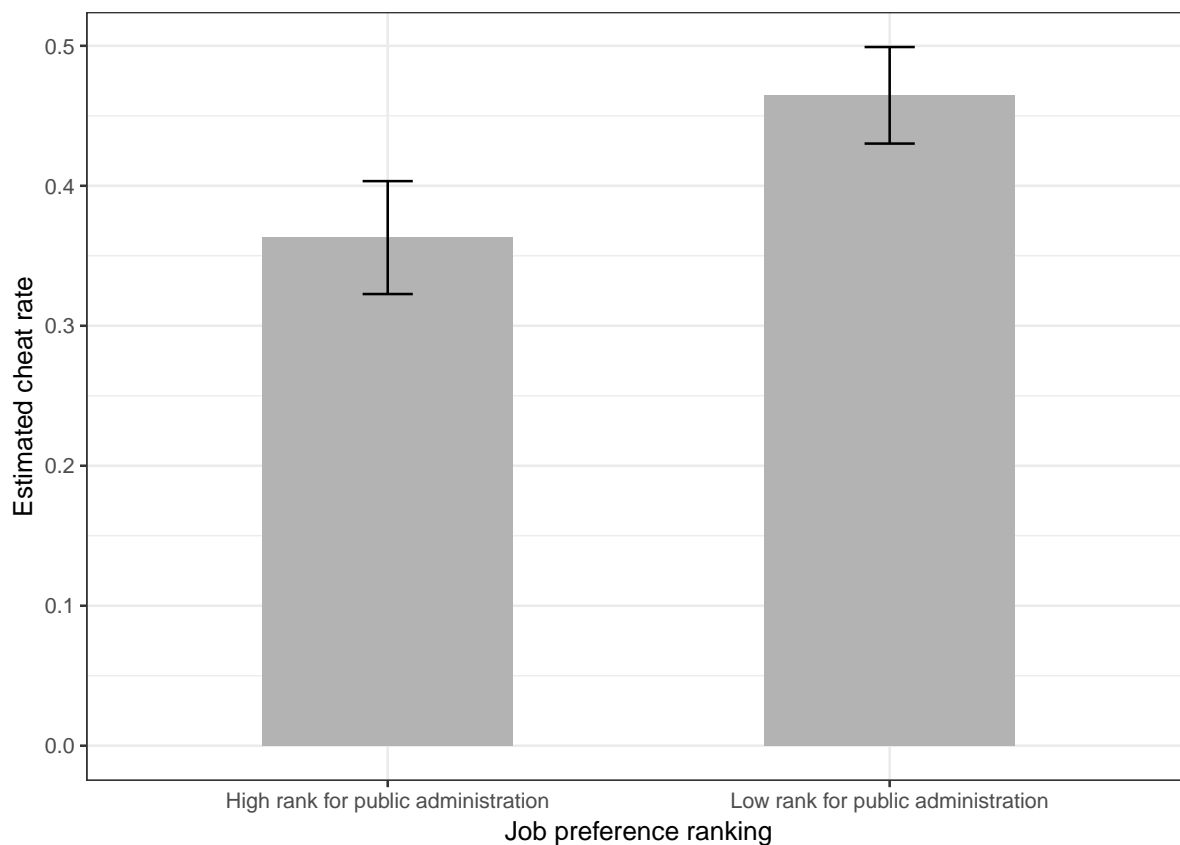
The figure shows the Corruption Perceptions Index (CPI) 1996-2014 for all available countries. The highlighted red line is Denmark, while the grey lines show the series for other countries.

Figure 2: Distribution of correct guesses and predicted distribution under full honesty



The histogram shows the observed number of correct guesses across students in our dice experiment (solid bars) and the predicted distribution under full honesty (outlined bars).

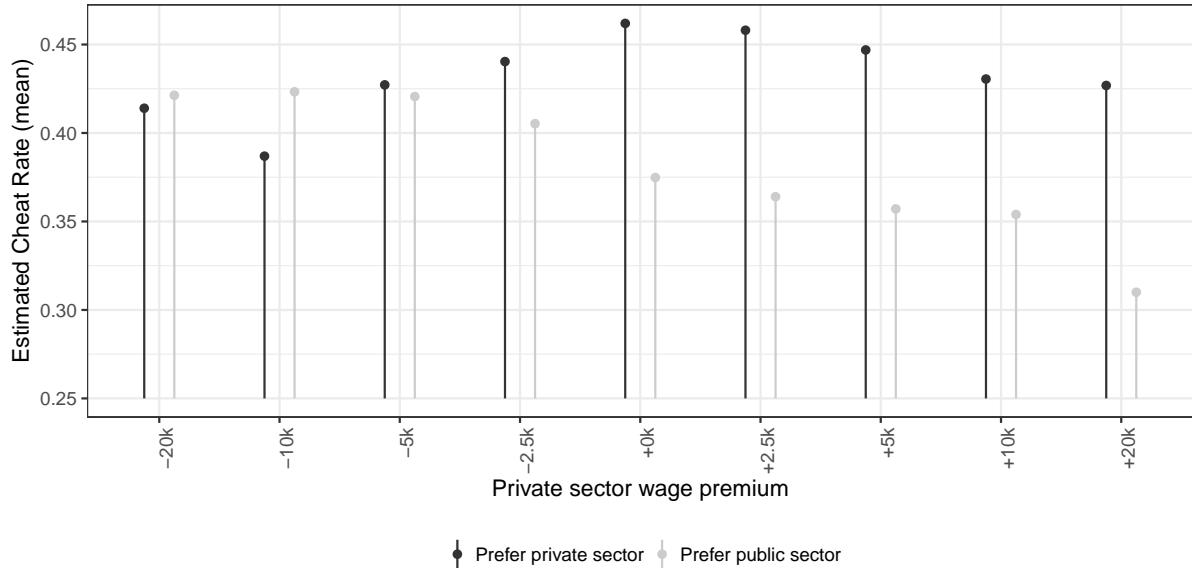
Figure 3: Mean estimated cheat rates by ranking of public administration



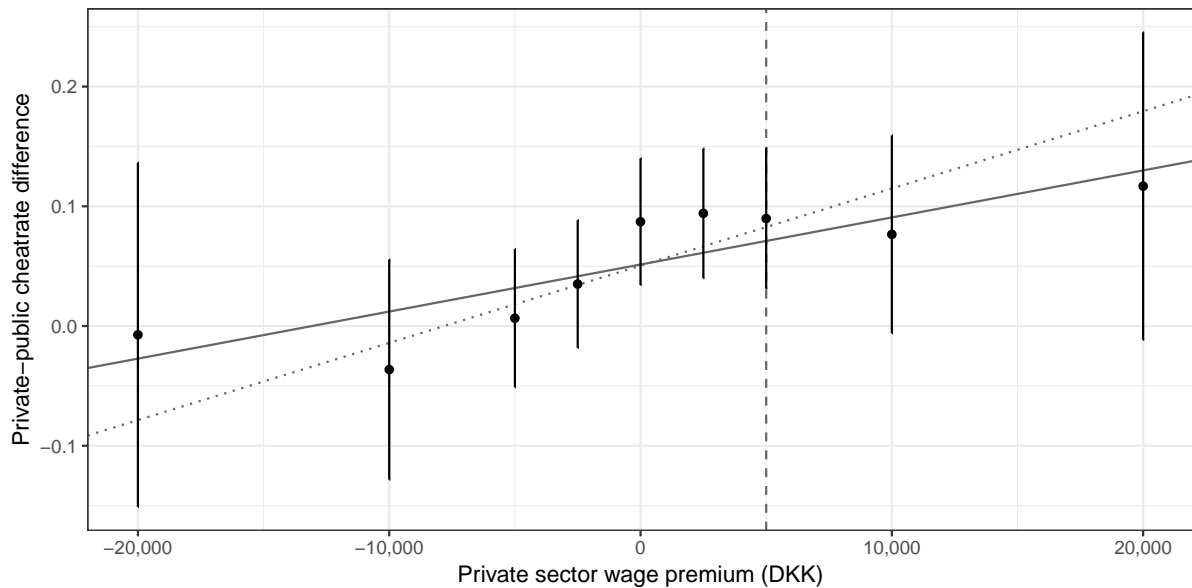
Mean estimated cheat rates for individuals ranking public administration in the top two (left) and individuals ranking public administration lower (right). The difference between the groups corresponds to the coefficient in Column 1 of Table 2. Error bars represent 95 pct. confidence intervals.

Figure 4: The private-public wage gap and the private-public dishonesty gap

Panel A: Average cheat rate for those preferring public and private sector in each wage scenario



Panel B: Private-public difference in average cheat rate by size of private-public wage gap



Panel A shows the averages estimated cheat rate among students preferring public and private sector in different counterfactual wage scenarios that vary the private sector wage premium. Each pair of one black and one grey line correspond to a different wage scenario. Black lines show the estimated cheat rates of those choosing the private sector sector in the wage scenario. Grey lines shows the estimated cheat rates for those choosing the public sector. Panel shows plots differences in estimated cheat rates between students preferring private and public sector against the public private wage gap. Error bars show 95 pct. confidence intervals. The dashed vertical line shows the approximate current private sector wage premium of +5,000 DKK.



Table 1: Summary statistics

Statistic	N	Mean	St. Dev.	Min	Max
Number of correct guesses	862	20.724	13.186	0	40
Estimated cheat rates	862	0.422	0.396	-0.200	1
Public administration rank $\leq 2$	862	0.422	0.494	0	1
Higher ranking of public administration	862	-3.414	2.079	-8	-1
Public service motivation score	860	2.440	0.521	0.250	3.950
Public sector picked at current wage	862	0.281	0.450	0	1
Probability of public administration	858	0.207	0.130	0	0.900
GPA (standardized)	861	-0.002	0.998	-5.914	2.332
Picks risky lottery	862	0.501	0.500	0	1
Job security rank $\leq 2$	862	0.119	0.325	0	1
Donation	862	6.798	6.521	0	15
Wage rank $\leq 2$	862	0.288	0.453	0	1
Male	862	0.536	0.499	0	1

The table shows summary statistics for the participants in the survey experiment. The variables are the number of reported correct guesses across the 40 dice games, the estimated cheat rate, 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, GPA standardized by field (the non-zero mean is due to the one excluded participant), an indicator for choosing one of the two most risky lotteries, the amount donated in the dictator game, the student's gender and indicators for whether job security and wage was ranked in the top two of the five job characteristics

Table 2: Dishonesty and public service job preferences

	Estimated cheat rate				
	(1)	(2)	(3)	(4)	(5)
Public administration rank $\leq 2$	-0.102** (0.027)				
Higher ranking of public administration		-0.022** (0.006)			
Public service motivation score			-0.152** (0.026)		
Public sector picked at current wage				-0.090** (0.029)	
Probability of public administration					-0.285** (0.105)
Constant	0.465** (0.018)	0.345** (0.025)	0.793** (0.066)	0.447** (0.016)	0.481** (0.026)
$N$	862	862	860	862	858
$R^2$	0.016	0.014	0.040	0.010	0.009
Implied change in cheat rate following a one std. dev. increase in preferences for public service:	-0.050** (0.013)	-0.047** (0.013)	-0.079** (0.014)	-0.040** (0.013)	-0.037** (0.014)

The table shows regressions of students' estimated cheat rates on various measures of public service job preferences. 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. The last row of the table rescales the estimated coefficients on the different job preferences measures so they reflect the implied change in cheat rates when the preference for public service is increased by one standard deviation. Robust standard errors in parenthesis. \* $p < 0.05$ ; \*\* $p < 0.01$ .

Table 3: Correlates of dishonesty and job preferences

<i>Panel A:</i>							
	Estimated cheat rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.007 (0.014)						0.014 (0.014)
Picks risky lottery		0.035 (0.027)					0.036 (0.027)
Job security rank $\leq 2$			0.002 (0.039)				-0.002 (0.038)
Donation				-0.016** (0.002)			-0.016** (0.002)
Wage rank $\leq 2$					0.083** (0.029)		0.048 (0.029)
Male						0.061* (0.027)	0.034 (0.027)
Constant	0.422** (0.013)	0.404** (0.019)	0.422** (0.014)	0.533** (0.019)	0.398** (0.016)	0.389** (0.019)	0.481** (0.028)
R <sup>2</sup>	0.000	0.002	0.000	0.073	0.009	0.006	0.082
<i>Panel B:</i>							
	Public administration rank $\leq 2$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GPA (standardized)	0.001 (0.017)						-0.002 (0.017)
Picks risky lottery		-0.058 (0.034)					-0.042 (0.034)
Job security rank $\leq 2$			-0.072 (0.051)				-0.093 (0.049)
Donation				0.009** (0.003)			0.006* (0.003)
Wage rank $\leq 2$					-0.202** (0.035)		-0.184** (0.036)
Male						-0.126** (0.034)	-0.092** (0.035)
Constant	0.423** (0.017)	0.451** (0.024)	0.431** (0.018)	0.364** (0.024)	0.480** (0.020)	0.490** (0.025)	0.513** (0.036)
R <sup>2</sup>	0.000	0.003	0.002	0.013	0.034	0.016	0.058
N	861	862	862	862	862	862	861

Panel A of the table shows regressions of students' estimated cheat rates on various measures of other student attributes. Panel B of the table shows the same regressions but replacing the outcome variable with an indicator for whether students ranked public administration in the top two of the eight job categories. 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 parenthesis. \*p<0.05; \*\*p<0.01.

Table 4: Dishonesty and job preferences conditional on other attributes

	Estimated cheat rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Public administration rank $\leq 2$	-0.103** (0.027)	-0.100** (0.027)	-0.102** (0.027)	-0.078** (0.026)	-0.091** (0.027)	-0.095** (0.027)	-0.068* (0.027)	-0.072** (0.026)
GPA (standardized)	0.007 (0.014)						0.014 (0.014)	
Picks risky lottery		0.030 (0.027)					0.033 (0.027)	
Job security ranked $\leq 2$			-0.005 (0.039)				-0.008 (0.037)	
Donation				-0.016** (0.002)			-0.016** (0.002)	-0.015** (0.002)
Wage ranked $\leq 2$					0.064* (0.029)		0.036 (0.029)	0.041 (0.029)
Male						0.048 (0.027)	0.027 (0.027)	
Constant	0.466** (0.018)	0.449** (0.023)	0.465** (0.019)	0.561** (0.021)	0.441** (0.021)	0.436** (0.023)	0.516** (0.032)	
$N$	861	862	862	862	862	862	861	861
$p$ -value, test of whether the estimate in the first row is the same as in Table 2	0.85	0.35	0.89	<0.01**	0.04*	0.10	<0.01**	<0.01**

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 measure is an indicator for whether public administration was ranked in the top two of the eight job categories. The measures of other attributes are GPA standardized by field, an indicator for choosing the most risky lottery, 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. \* $p < 0.05$ ; \*\* $p < 0.01$ .

# A Appendix (FOR ONLINE PUBLICATION)

## 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  $\theta_i$  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  $\psi_i$  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  $\psi_i$  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  $\theta_i$ . Accordingly, the probability of observing a win for individual  $i$ , conditional on his cheat rate is:

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

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

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

Replacing the expectation  $E(y_{ik} | \theta_i)$  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:

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

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:<sup>1</sup>

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

The object of interest here is the parameter  $\beta_1$ , which captures the relationship between individual  $i$ 's cheat rate,  $\theta_i$  and their job preferences,  $X_i$ . In the particular case where  $X_i$  is an indicator for whether  $i$  prefers a public service career,  $\beta_1$  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,  $\theta_i$ . 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  $\xi_i$  that is simply defined as  $\xi_i \equiv \hat{\theta}_i - \theta_i$ :

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

---

<sup>1</sup>More formally, we let  $\beta_0$  and  $\beta_1$  be defined in the usual (implicit) way by imposing  $E(\varepsilon_i) = 0$  and  $Cov(X_i, \varepsilon_i) = 0$  in (1).

Now, because  $\hat{\theta}_i$  is unbiased for  $\theta_i$  and because the measurement error in  $\hat{\theta}_i$  stems solely from randomness in whether individuals' win truthfully or cheat in any specific round of the dice game, it follows that the measurement error  $\xi_i$  has mean zero and is mean independent of both  $\theta_i$  and  $X_i$ :<sup>2</sup>

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

It follows that the measurement error  $\xi_i$  is classical.<sup>3</sup> 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  $\beta_1$  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.

---

<sup>2</sup>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  $\theta_i$ . The conditional probability of reporting a win in the dice game is therefore  $E(y_{ik} | \theta_i, X_i) = p^* + (1 - p^*)\theta_i$ . The same derivation that showed unbiasedness in Section A.1.1 therefore shows that  $E(\hat{\theta}_i | \theta_i, X_i) = \theta_i$ , which further implies  $E(\xi_i | \theta_i, X_i) = 0$ .

<sup>3</sup>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.



### 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,  $\xi_i$  is absorbed into the composite error term in (2), however, the measurement error does 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.<sup>4</sup> In this case, for an individual with cheat rate  $\theta_i$ , 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  $\theta_i$ ,  $Y_i$  therefore follows a binomial distribution:

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

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):

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

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

---

<sup>4</sup>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.

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 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.<sup>5</sup>

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 individuals have the same cheat rate.<sup>6</sup>

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 the expected share of wins for this individual,  $\frac{Y_i}{K} \xrightarrow{P} \psi_i$  when  $K \rightarrow \infty$ . 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,

---

<sup>5</sup>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.

<sup>6</sup>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  $P(y_{i1} = 1) = p^* + (1 - p^*)\bar{\theta} = x$ .

$F$ .<sup>7</sup>

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 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,  $\dots$ ,  $K$  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.<sup>8</sup>

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.<sup>9</sup> Given that the full distribution of cheat rates is not non-parametrically identified, the next section proceeds by developing a parametric estimator for  $F$ .<sup>10</sup> 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

---

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

<sup>8</sup> 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  $\theta_i$  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  $x = 0, 1, \dots, K$ . 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:

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

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

<sup>9</sup>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.

<sup>10</sup>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.

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  $\lambda \in \mathbb{R}^v$ .<sup>11</sup>

Once  $F$  is assumed to belong to some parametric family, we can develop a Maximum Likelihood estimator for  $F$ . When cheating behavior is independent over time, the number of reported wins conditional on  $\theta_i$  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  $\theta_i$  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  $Y_1, Y_2, \dots, Y_N$  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  $\lambda$  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

---

<sup>11</sup>We discuss the specific choice of parametric family in Section A.2.1.

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 career. 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  $m$ .<sup>12</sup>

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  $\lambda$ . 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  $\theta$  and a real-valued vector  $\zeta \in \mathbb{R}^q$ , so that estimation of  $m$  is simply equivalent to estimation of  $\zeta$ .<sup>13</sup>

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

Next, to derive the likelihood function, we note that conditional on the cheat rate,  $\theta_i$ , 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 :

---

<sup>12</sup>Joint knowledge of both of these of course permits one to compute the joint distribution of the variables  $\theta_i$  and  $X_i$  as well as any corresponding conditional or marginal distribution.

<sup>13</sup>When  $F$  is discrete, each elements of  $\zeta$  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  $\zeta$  will instead be the parameters of the functional form imposed on  $m$ . In both cases, estimation of  $m$  is simply equivalent to estimation of  $\zeta$ .

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

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

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

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

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

Maximization of the log likelihood function with respect to  $\lambda$  and  $\zeta$  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 family 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 individual 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 continuous. 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.<sup>14</sup> 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

---

<sup>14</sup>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).<sup>15</sup>

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

<sup>15</sup> 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 (0.058)	0.274 (0.063)	0.288 (0.037)
Mean	0.975 (0.052)	0.975 (0.038)	0.975 (0.018)
Variance	0.001 (0.018)	0.001 (0.015)	0.001 (0.006)
	Beta-mixture component II:		
Weight	0.725 (0.058)	0.672 (0.060)	0.712 (0.037)
Mean	0.214 (0.018)	0.205 (0.056)	0.205 (0.038)
Variance	0.049 (0.008)	0.052 (0.021)	0.052 (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 assess 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) ( $p < 0.01$ ).<sup>16</sup>

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

---

<sup>16</sup>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 (0.037)	0.171 (0.040)	0.166 (0.042)
	Intermediate type I:		
Share in population	0.273 (0.030)	0.272 (0.032)	0.271 (0.035)
Cheat rate	0.105 (0.020)	0.100 (0.019)	0.096 (0.018)
	Intermediate type II:		
Share in population	0.196 (0.022)	0.195 (0.023)	0.193 (0.027)
Cheat rate	0.355 (0.021)	0.343 (0.026)	0.333 (0.031)
	Intermediate type III:		
Share in population	0.074 (0.012)	0.071 (0.015)	0.058 (0.023)
Cheat rate	0.680 (0.022)	0.642 (0.046)	0.589 (0.058)
	Intermediate type IV:		
Share in population	0.015 (0.015)	0.060 (0.015)	0.032 (0.015)
Cheat rate	0.952 (0.007)	0.892 (0.053)	0.755 (0.052)
	Intermediate type V:		
Share in population	-	0.177 (0.040)	0.068 (0.027)
Cheat rate	-	0.986 (0.012)	0.922 (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 (0.015)	0.054 (0.052)	<0.001
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.<sup>17</sup>

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.

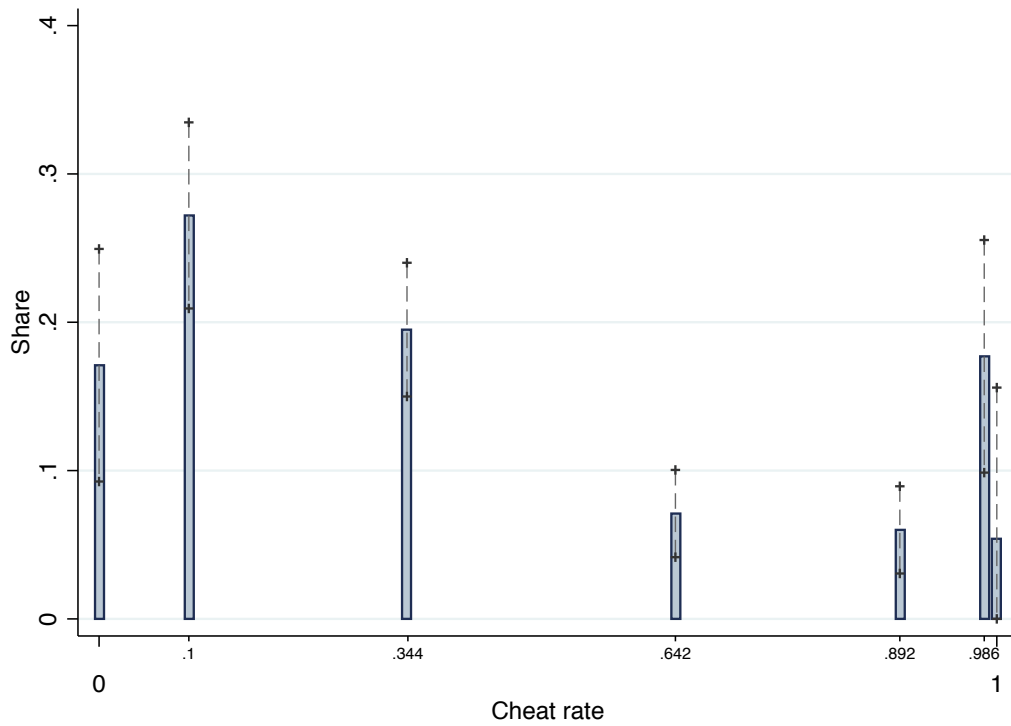
If we focus instead on the share of people who are practically completely honest or dishonest, however, we note that the discrete models tend to imply a higher fraction of

---

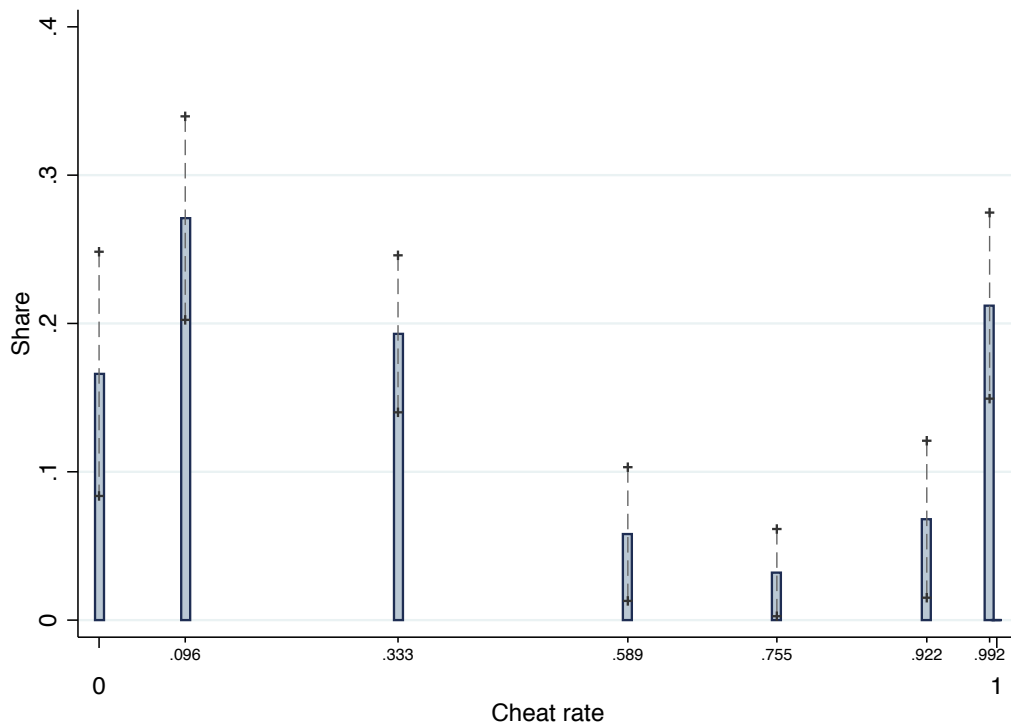
<sup>17</sup>The fact that the precise distribution of individuals across the most dishonest types is sensitive to the choice parametric model but that the share of individuals cheating more than 98 percent of the time is not, illustrates the limits to what we can identify in our data. When respondents 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)



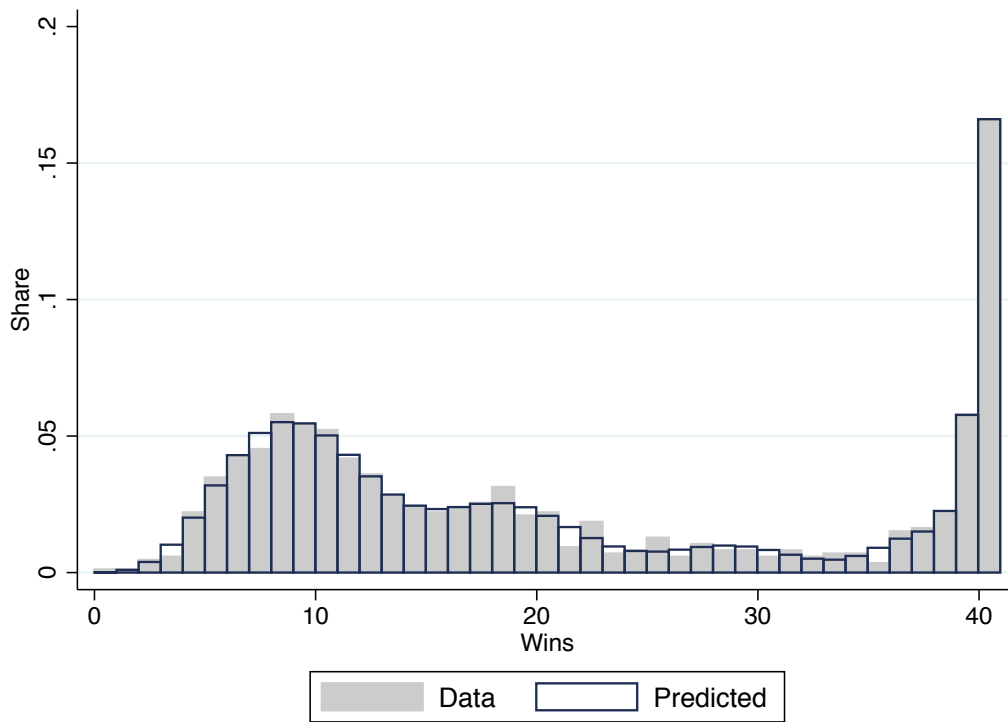
Panel B: Distribution of dishonesty, 8 discrete types, Model (3)




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 discrete 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.<sup>18</sup> 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.<sup>19</sup> 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.

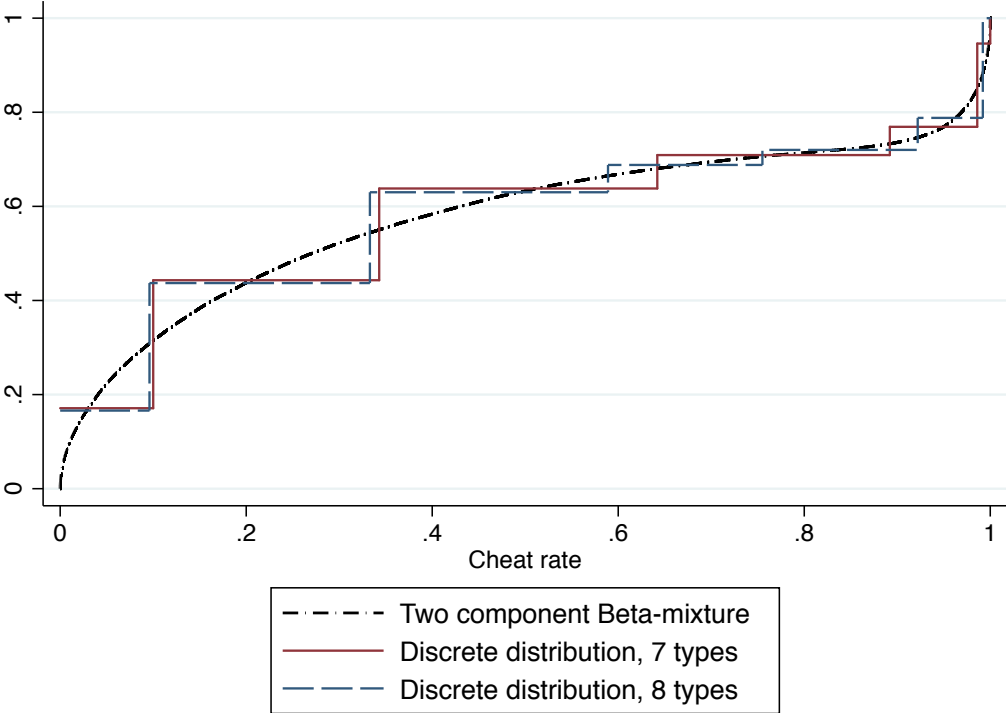
For the continuous distribution we use a two component Beta-mixture. In this case, however, we need to impose a functional form on the probability of preferring public service as function of the cheat rate. To allow for a flexible relationship between job

---

<sup>18</sup>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.

<sup>19</sup>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.

preferences 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.<sup>20</sup>

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.

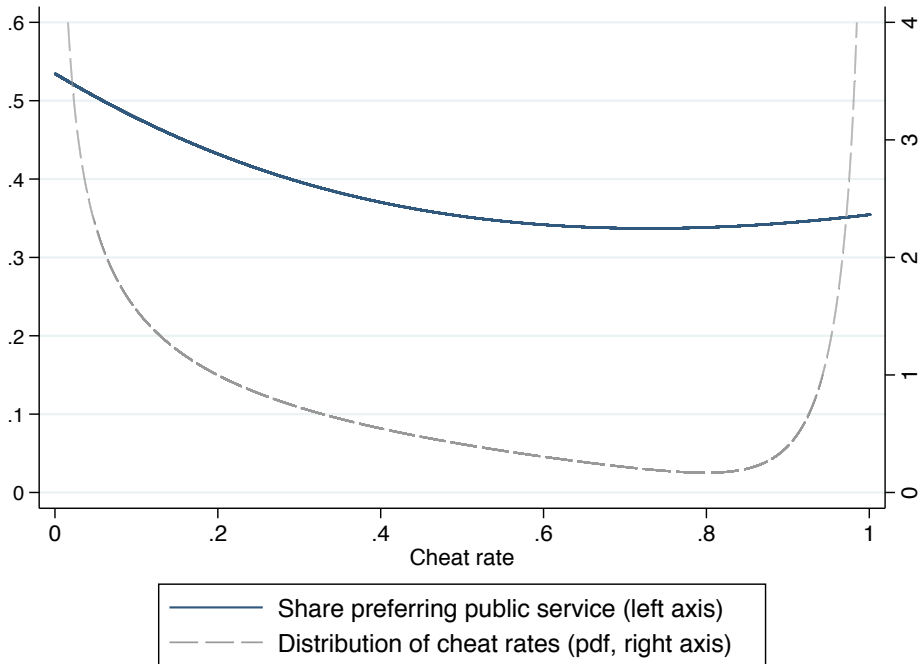
Overall, we see that the systematic self-selection of honest individuals into public service, is driven by particularly strong preferences for public service among the most honest individuals, while the preferences for public service jobs differ less across the moderately to very dishonest.

---

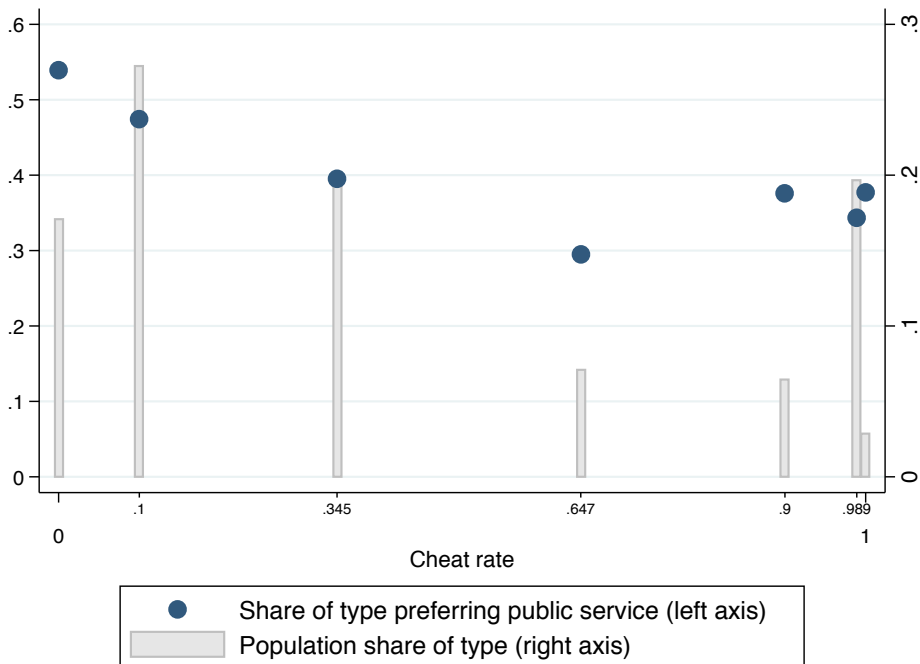
<sup>20</sup>Since we are estimating the job preferences and the distribution of cheat rates jointly, the estimated distributions will not necessarily be the same as the ones in Sections A.2.2 and A.2.3. As the Figures show, however, we see only very slight differences in the distributions.

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 preference 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 A.4 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.<sup>21</sup> 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.

---

<sup>21</sup>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.<sup>22</sup> 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

---

<sup>22</sup>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 (0.051)	0.014 (0.035)	0.050 (0.052)	0.032 (0.033)
Dictator donation, standardized	-0.192** (0.054)	-0.269** (0.032)	-0.189** (0.052)	-0.271** (0.033)
Male	0.125* (0.058)	0.076* (0.034)	0.170 (0.111)	0.055 (0.046)
Locus of control, standardized	0.015 (0.036)	-0.045 (0.035)	0.005 (0.036)	-0.058 (0.034)
N	614	862	614	862
R <sup>2</sup>			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.<sup>23</sup>

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.<sup>24</sup> 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

---

<sup>23</sup>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.

<sup>24</sup>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 A.5 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 is 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 A.5. 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 administration ranked $\leq 2$	0.479** (0.0714)					-0.101 (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** (0.0463)	0.785** (0.0668)	0.144 (0.211)	0.377** (0.0474)	0.105* (0.0525)	0.489** (0.0433)	0.482** (0.0456)
<i>N</i>	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 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 administration ranked $\leq 2$	-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)
<i>N</i>	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 administration 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*** (0.018)	0.327*** (0.026)	0.756*** (0.066)	0.421*** (0.016)	0.453*** (0.025)
<i>N</i>	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 administration 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*** (0.019)	0.345*** (0.027)	0.753*** (0.070)	0.441*** (0.017)	0.479*** (0.027)
<i>N</i>	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 administration 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*** (0.019)	0.353*** (0.027)	0.837*** (0.069)	0.462*** (0.017)	0.505*** (0.027)
<i>N</i>	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 administration 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*** (0.019)	0.355*** (0.027)	0.825*** (0.070)	0.463*** (0.017)	0.487*** (0.027)
<i>N</i>	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 $\leq 2$	-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** (0.018)	0.332** (0.026)	0.762** (0.067)	0.435** (0.016)	0.470** (0.026)
<i>N</i>	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 administration ranked $\leq 2$	-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** (0.017)	0.249** (0.023)	0.537** (0.065)	0.322** (0.015)	0.351** (0.023)
<i>N</i>	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 (0.020)						0.028 (0.020)
Picks risky lottery		0.053 (0.041)					0.025 (0.042)
Job security rank $\leq 2$			0.007 (0.063)				0.009 (0.062)
Donation				-0.011** (0.003)			-0.011** (0.003)
Wage rank $\leq 2$					0.034 (0.045)		-0.006 (0.045)
Male						0.155** (0.040)	0.143** (0.042)
Constant	0.319** (0.020)	0.291** (0.029)	0.317** (0.022)	0.396** (0.030)	0.308** (0.024)	0.235** (0.029)	0.305** (0.043)
<i>N</i>	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 (0.014)						0.019 (0.013)
Picks risky lottery		0.030 (0.027)					0.030 (0.028)
Job security ranked $\leq 2$			0.004 (0.040)				0.001 (0.039)
Donation				-0.015*** (0.002)			-0.014*** (0.002)
Wage ranked $\leq 2$					0.069** (0.029)		0.037 (0.029)
Male						0.057** (0.027)	0.035 (0.028)
Constant	0.401*** (0.014)	0.385*** (0.019)	0.399*** (0.015)	0.500*** (0.019)	0.380*** (0.016)	0.369*** (0.019)	0.454*** (0.028)
<i>N</i>	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 (0.015)						0.010 (0.014)
Picks risky lottery		0.036 (0.028)					0.036 (0.029)
Job security ranked $\leq 2$			-0.023 (0.043)				-0.027 (0.041)
Donation				-0.017*** (0.002)			-0.016*** (0.002)
Wage ranked $\leq 2$					0.081*** (0.030)		0.045 (0.030)
Male						0.060** (0.028)	0.032 (0.029)
Constant	0.421*** (0.014)	0.402*** (0.019)	0.423*** (0.015)	0.533*** (0.020)	0.397*** (0.017)	0.388*** (0.020)	0.487*** (0.029)
<i>N</i>	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 (0.015)						0.017 (0.015)
Picks risky lottery		0.038 (0.029)					0.040 (0.029)
Job security ranked $\leq 2$			0.020 (0.042)				0.017 (0.040)
Donation				-0.017*** (0.002)			-0.017*** (0.002)
Wage ranked $\leq 2$					0.084*** (0.031)		0.048 (0.031)
Male						0.065** (0.029)	0.036 (0.029)
Constant	0.435*** (0.014)	0.415*** (0.020)	0.432*** (0.015)	0.552*** (0.020)	0.410*** (0.017)	0.400*** (0.020)	0.494*** (0.030)
<i>N</i>	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 (0.016)						0.009 (0.015)
Picks risky lottery		0.038 (0.029)					0.039 (0.030)
Job security ranked $\leq 2$			0.005 (0.044)				0.003 (0.042)
Donation				-0.017*** (0.002)			-0.016*** (0.002)
Wage ranked $\leq 2$					0.097*** (0.032)		0.063** (0.032)
Male						0.061** (0.029)	0.031 (0.030)
Constant	0.433*** (0.015)	0.414*** (0.020)	0.432*** (0.016)	0.547*** (0.021)	0.405*** (0.017)	0.400*** (0.021)	0.490*** (0.030)
<i>N</i>	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 (0.014)						0.016 (0.014)
Picks risky lottery		0.024 (0.027)					0.028 (0.028)
Job security rank $\leq 2$			0.016 (0.040)				0.013 (0.038)
Donation				-0.016** (0.002)			-0.016** (0.002)
Wage rank $\leq 2$					0.089** (0.030)		0.057 (0.029)
Male						0.045 (0.027)	0.020 (0.028)
Constant	0.409** (0.014)	0.397** (0.019)	0.407** (0.015)	0.517** (0.019)	0.383** (0.016)	0.385** (0.019)	0.472** (0.029)
<i>N</i>	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 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 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 (0.013)						0.006 (0.012)
Picks risky lottery		-0.005 (0.025)					0.004 (0.024)
Job security rank $\leq 2$			0.030 (0.036)				0.030 (0.035)
Donation				-0.014** (0.002)			-0.013** (0.002)
Wage rank $\leq 2$					0.100** (0.027)		0.076** (0.027)
Male						0.013 (0.025)	-0.001 (0.024)
Constant	0.307** (0.012)	0.309** (0.017)	0.303** (0.013)	0.406** (0.019)	0.278** (0.014)	0.300** (0.017)	0.375** (0.026)
<i>N</i>	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 2-5 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 job measures the coefficient drops by between 32 and 36 percent when the controls are added. For the last 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)
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 (0.100)
Donation	-0.015** (0.002)	-0.014** (0.002)	-0.015** (0.002)	-0.015** (0.002)
Wage ranked $\leq 2$	0.043 (0.029)	0.038 (0.029)	0.048 (0.029)	0.049 (0.029)
Constant	0.464** (0.031)	0.758** (0.068)	0.529** (0.024)	0.545** (0.030)
<i>N</i>	862	860	862	858

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: Political 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 administration ranked $\leq 2$	-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 (0.024)						0.041 (0.024)
Picks risky lottery		0.045 (0.035)					0.044 (0.030)
Job security ranked $\leq 2$			0.007 (0.044)				0.013 (0.041)
Donation				-0.014** (0.003)			-0.014** (0.002)
Wage ranked $\leq 2$					0.082* (0.035)		0.046 (0.032)
Male						0.077* (0.034)	0.046 (0.031)
Constant	0.390** (0.015)	0.363** (0.030)	0.384** (0.019)	0.474** (0.032)	0.361** (0.023)	0.344** (0.027)	0.423** (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 (0.025)						0.029 (0.028)
Picks risky lottery		-0.043 (0.046)					-0.036 (0.042)
Job security ranked $\leq 2$			-0.007 (0.060)				-0.032 (0.058)
Donation				0.011** (0.003)			0.009** (0.003)
Wage ranked $\leq 2$					-0.178** (0.046)		-0.163** (0.042)
Male						-0.117* (0.048)	-0.085 (0.046)
Constant	0.402** (0.021)	0.419** (0.037)	0.399** (0.024)	0.330** (0.032)	0.449** (0.029)	0.459** (0.036)	0.458** (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 button.
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 button 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 terningsspillet 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*




Terningslaget blev en 2'er.

Hvilket tal tænkte du på? Indtast i feltet:

Rapportér dit svar her

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

*Screenshot*



Terningslaget blev en 1'er.

Hvilket tal tænkte du på? Indtast i feltet:

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

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 forinden for at slå med den digitale terning.

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.



Terningenslaget blev en 6'er.

Hvilket tal gættede du på? Indtast i feltet:

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 terningenslaget. Du vinder 1 kr.

Din samlede gevinst for undersøgelsen hidtil er nu 16 kr.

### 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
<b>Mulighed A</b>	0 DKK	0 DKK	<b>0 DKK</b>
<b>Mulighed B</b>	5 DKK	3 DKK	<b>8 DKK</b>
<b>Mulighed C</b>	10 DKK	4 DKK	<b>14 DKK</b>
<b>Mulighed D</b>	15 DKK	4 DKK	<b>19 DKK</b>

Hvilken donationsmulighed vælger du?

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

### 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"
<b>Mulighed A</b>	200 DKK	0 DKK
<b>Mulighed B</b>	160 DKK	30 DKK
<b>Mulighed C</b>	140 DKK	40 DKK
<b>Mulighed D</b>	120 DKK	50 DKK
<b>Mulighed E</b>	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.

## 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 Statistic for the Number of Components in a Normal Mixture.” *Applied Statistics* 36 (3): 318.