Peer Effects on Partisanship, Registration, and Voter Turnout: Evidence from Roommate Assignments

Kenneth Coriale, Ethan Kaplan, and Rachel Nesbit *

October 31, 2023

Abstract

We use conditional random assignment of roommates at the University of Maryland to estimate the causal impact of social connections on political behavior. Using university data on roommate assignments matched to voter registration data, we compare counterfactual assignments of roommates to those that actually occurred. We develop new methods to estimate exact bounds on treatment effects to account for a partially missing independent variable. We find evidence of homophilic preference formation as roommates are more likely to both be registered with the same party and to both vote, in comparison to counterfactual roommate pairs randomly not selected. Keywords: Preferences, Voting, Registration, Peer Effects.

**** Preliminary. Not for circulation. Please do not cite. ****

*We have benefited from helpful comments from Allan Drazen and David Karol. We particularly thank Danielle Glazer, Joann Prosser, Scott Young, and the Department of Resident Life at the University of Maryland for matching voter registration data to data on roommates and for providing institutional background on the roommate matching process. This paper simply would not have been possible without their tireless efforts over a number of years. All errors and omissions are our own. Correspondence can be addressed to Coriale at `kcoriale@towson.edu`, Kaplan at `kaplan@econ.umd.edu`, or Nesbit at `rnesbit@umd.edu`.

1 Introduction

How do political preferences form? Why are some people Democrats and others Republicans? Why do some people register and turn out to vote whereas others do not? More generally, why do people adopt their political identities and their levels of political involvement? Prior literature has emphasized the roles of self-interest (Black [1948])), class interest (Marx [1926, 2005]), ideological hegemony (Gramsci [2011]), media (Herman and Chomsky [2010], DellaVigna and Kaplan [2007], Gentzkow [2006], Martin and Yurukoglu [2017], Prior [2013]), and education (Cantoni et al. [2017], Bowles and Gintis [2011], Marshall [2019]).

Another potentially important determinant of both political beliefs and political behavior is social connections. There is ample evidence that friends tend to share similar political views. For example, a 2017 Pew Research poll finds that only 14% of Republicans and 9% of Democrats say that "a lot" of their close friends are from the opposite party.[1] A 2018 PRRI poll finds that 35% of Republicans and 45% of Democrats would be unhappy if their child's spouse was of the opposite party.[2]

Recent research has documented that friendships are more politically segregated than news consumption. Gentzkow and Shapiro [2011] compute partisan media segregation and compare it to partisan segregation in friendships.[3] They find that online media segregation "is somewhat higher than that of a social network where individuals matched randomly within counties (5.9) and lower than that of a network where individuals matched randomly within ZIP codes (9.4). It is significantly lower than the segregation of actual networks formed through voluntary associations (14.5), work (16.8), neighborhoods (18.7), or family (24.3). The Internet is also far less segregated than networks of trusted friends (30.3) and political discussants (39.4)." In fact, some recent commentators have claimed that political segregation is one of the main reasons for increased partisan animosity and hatred (Bishop [2009], Mason [2018]).

In this paper, we estimate the causal effect of social connections on partisan orientation and political participation through voting. Though the correlation in partisan attitudes and voter turnout is quite strong among friends, it is not clear whether that reflects the impact of politics on friendship formation, spurious correlation due to other factors such as race, gender, and education which influence both friendship formation and political attitudes and behaviors, or the causal impact of friendship upon political attitudes and behaviors.

The objective of this paper, to isolate the causal impact of friendship on politics, is challenging

[1]https://www.pewresearch.org/politics/2017/10/05/8-partisan-animosity-personal-politics-views-of-trump/8_02/

[2]https://www.prri.org/research/american-democracy-in-crisis-the-fate-of-pluralism-in-a-divided-nation/

[3]Segregation indices measure the difference in interaction across two groups. The segregation index is minimized at zero in the absence of segregation and is maximized at one hundred in the presence of complete segregation.

for two reasons. First, friendships rarely form randomly. Social interactions, of course, often have a strong random component. However, random social interactions are rarely deep enough to lead to friendships important enough to substantively influence political opinions or behavior. Second, it is often very difficult to find data on friendship formation. Governments do collect records on voter behavior including registration, turnout, and partisan affiliation, but largely do not collect data on friendships.

We are far from the first to write on the impact of social connection on attitudes. In 1954, Allport et al., discussed how exposure to people from other social groups lowers discrimination. A number of recent papers consider the long run impact of schooling on political attitudes and political participation. Bergman [2018] uses a busing lottery to estimate the long run impact of being bused on voter registration and voter turnout among minorities in California. His estimates are both very small and statistically indistinguishable from zero. Billings et al. [2020] and Kaplan et al. [2019] estimate the impact of racial integration on political partisanship, registration, and turnout. They both find large impacts upon partisanship and no effects on registration and turnout many years later. Similarly, Calderon et al. [2023] show that the great migration of African-Americans northwards in the U.S. resulted in shifts in attitudes of their new white neighbors. Of course, racial integration is a composite effect, encompassing many possible channels of influence. In contrast, our paper isolates one particular potential channel of influence on political attitudes: peer effects.

Economists have estimated peer effects in other contexts including such diverse areas as academic performance at school (Carrell et al. [2013], Sacerdote [2001]), major choice at college (Sacerdote [2001]), productivity at work (Cornelissen et al. [2017], Mas and Moretti [2009]), and managerial practices (Shue [2013]). We add to this list the important outcomes of political orientation and political action.

Fortuitous for our research design, one of the only common situations in which the formation of deep friendships is, in part, governed by a random process is also one in which social connections are well documented. Colleges and universities often randomly assign roommates to each other conditional upon a small number of observed characteristics such as answers to lifestyle preference questionnaires and gender. Other papers have used randomness in college roommate assignment to estimate the effects of peers on a variety of outcomes. Sacerdote's groundbreaking paper in 2001 used the random component of roommate assignment at Dartmouth College to estimate the impact of friendship on grade point average, social group formation, and major choice. More related to our own paper, Boisjoly et al. [2006] use random assignment in UCLA dormitories to estimate the impact of having an African-American roommate on White preferences towards affirmative action. They control for answers to questions used to match roommates and estimate the impact of having an African-American roommates on preferences over affirmative action. They find that having an African-American roommate makes White students more supportive of affirmative action.

Strother et al. [2021] is the paper most similar to ours. They use the randomness in roommate assignment at two undisclosed universities to estimate the impact of roommates on ideology. Using survey responses, they find that roommates who differ politically in a pre-freshman baseline survey converge toward their roommate's ideology. They do this by regressing an ordinal measure of roommate 2's ideology in the follow-up survey on roommate 1's baseline ideology conditional upon controls. Relative to their paper, we offer a number of advantages. First, their paper focuses on ideology whereas ours focuses on partisanship (party affiliation) and political participation. As a consequence, their data comes from voluntary survey responses whereas we use administrative voter registration data. Second, their sample is subject to substantial, potentially selective, attrition. In contrast, we use administrative data for our analysis; as a result, do not encounter the same attrition issues. If individuals who had a better experience with their roommate were more likely to respond to the Strother et al. [2021] survey, and if those pairs were also more likely to be similar politically, then their paper may find political convergence simply because the politically divergent roommates left the sample. Because Strother et al. [2021] require both roommates to respond to the survey in order to be included in the sample, their sample from their first survey included only 52% of all rooms, and their follow-up survey contained only 63% of those original 52% who both answered the first survey. In the end the final sample comprises only 34% of the original pool of roommates. The non-response and attrition rates are sufficiently large as to potentially explain the size of the effects that they estimate even with modest amounts of sample selection. Third, our paper differs in that we construct treatment and control groups by modelling the roommate matching process and comparing actual roommate pairings to counterfactual pairings that were equally likely to have been formed but didn't for random reasons. In contrast, Strother et al. [2021] control linearly for variables used in roommate selection. Fourth, Strother et al. [2021] estimate effects during only the first academic year, while the pairings are usually still roommates, whereas we explore dynamic effects over time for a period of between one and five years.

Methodologically, we are closer to Sacerdote [2001]. However, we make a number of methodological innovations, including using a randomization inference design and deriving bounds on the effects due to partially missing data. Our basic strategy involves matching data on freshman roommate assignments from the University of Maryland at College Park with voter registration files from the state of Maryland, nearby states, and the District of Columbia. Roommates at the University of Maryland are assigned to each other based upon answers to five lifestyle survey questions, gender, and participation in a living-learning community (a residential academic community, henceforth LLP). We analyze voter turnout, voter registration, and party of registration as outcomes for pairs of individuals matched as roommates, compared to the potential roommate pairings that could have occurred based upon answers to the lifestyle questions, gender, and LLP.

We present results from three different estimation strategies. We restrict our analysis to roommate pairings who have identical answers to all lifestyle questions, are in the same LLP (or are both not in any LLP), and have the same gender. In this trimmed sample, roommate assignments are random. In our first econometric specification, we fully saturate an OLS model with lifestyle answer, LLP, year, and gender fixed effects and regress the outcome of one roommate on the outcome of the other. Since we do not observe the direction of causation within rooms, interpretation of these estimates is difficult. Instead, our main estimation strategy employs a randomization inference design.[4] Different from standard randomization inference designs, we do not have variation treatment status in our data, since treatment (roommate assignment) is at the room level and all rooms are treated. Thus, our best comparison is between the actual roommates that were paired and counterfactual roommate pairings that were equally likely to have been created but were, for random reasons, not selected. We compute the fraction of rooms that have two Democrats, two Republicans, two Independents, two registered individuals, two roommates who have both voted, two roommates who have the same party affiliation. We then compute the percentile of double positive outcomes (i.e. double Democrats or double registered) relative to the distribution of counterfactual double positive or *homophilic* outcomes[5]. By permuting roommate assignment within groups of identically matched individuals, we eliminate the role of covariates.

Since we have a large number of observations, we cannot compute the full distribution of counterfactual roommates assignments. Instead we limit ourselves to a Monte Carlo simulation of 10,000 replications of roommates pairings in order to produce a distribution of potential outcomes. Randomization inference allows us to compute p-values for the probability of observing as many or more homophilic outcomes by random chance. However, traditional randomization inference does not provide us with a notion of a treatment effect. For this, we compute two measures of a treatment effect: (1.) the actual degree of homophily (fraction of outcomes which are the same) relative to the mean of the counterfactual distribution and (2.) the actual degree of homophily relative to the median of the counterfactual distribution. In a third estimation strategy, we combine the actual data and the counterfactual data in one regression and run a fully saturated OLS model where treatment is the set

[4]We are the first that we are aware of to use a randomization inference design to estimate peer effects using roommates.

[5]We use homophily and double positive outcomes synonymously. We could measure homophily as all the instances in which both roommates, as an example, shared the same partisanship. Instead, we will measure homophily as the number of double positive outcomes. So when we look at Democrats, our measure of homophily will be the number of rooms where both roommates are Democrats independent of whether the rooms without two Democrats are all double Republican and double Independent, all mixed partisan, or somewhere in between these two extremes.

of roommate pairings that actually occurred and control is the set of pairings that randomly did not occur. This uses similar variation to our main randomization inference strategy, but puts the estimates in a more traditional OLS framework.

We encounter one additional difficulty in our estimation which is the source of the largest econometric contribution of the paper. We have all of the data from the Residential Life department used in matching roommates, except for gender. The gender variable was owned by the registrar, who would not provide us with the data. However, in 84% of cases, we can recover gender from the voter registration file that matches the individual or from the voter registration file of the individual's roommate. We then derive exact p-value bounds for our randomization inference p-value estimates by filling in the remaining missing genders. Different from papers such as Lee [2009], our missing variable is an independent variable and is discrete, meaning for each individual we only have to consider two possible values that the variable could take. Nonetheless, it is computationally infeasible to compute all the counterfactual assignments to gender which could have occurred. We thus derive the formula for the probability of a certain number of double outcome rooms, given a fraction of roommates with a given outcome, and we use that to compute the distribution of the number of counterfactual rooms with a double positive outcome. We then use this distribution to compute the treatment effect minimizing and treatment effect maximizing assignment of gender for each room with unknown gender[6]. As a result of this, we can compute a sharp lower-bound effectand a sharp upper-bound effect.

Our main findings are that there are between 4% and 15% higher pairings of Republicans compared to the median counterfactual, 5-13% higher pairings of Democrats, 2-10% higher pairings of Independents, a 3-7% higher rate of both roommates voting, a -0.8-3.8% higher rate of both roommates being registered and a 6-11% probability of both roommates being from the same party. We thus provide causal evidence that friends influence each others' partisanship and political participation to a significant degree.

In Section 2, we discuss roommate assignment at the University of Maryland; in Section 3, we discuss the data that we use for our analysis; in Sections 4 and 5, we discuss our methods for estimating treatment effects and present our initial empirical results; in Sections 6 and 7, we discuss our method for bounding the treatment effects based on the partially missing data and present the corresponding empirical results; in Section 8, we conclude.

2 Institutional Background

Our paper uses data from freshmen roommate assignment at the University of Maryland at College Park from the years 2011-2015. In this section, we discuss the University of Maryland, its on-campus

[6] Where the definition of treatment effect is the actual outcome relative to the mean counterfactual outcome

housing, and the assignment process of roommates. The assignment process is important as our research design is based on the fact that aside from a small number of demographic variables and answers to lifestyle questions, assignment to roommates is mostly random.

The University of Maryland at College Park is a large, selective public research university.[7] Undergraduate enrollment over this five year period averaged 26,900 students per year. The university size of the Freshman class was 3,991, and most students live on campus during their first year. 92% of freshmen students live in dormitories, another 5% live in suites, and 3% live in apartments. We focus on those living in dorms as freshmen living in apartments are likely to have been selected by their roommates. Out of dormitories, singles do not allow for the estimation of roommate effects and thus are excluded. In triples and quadruples, conceptualizing roommate effects is substantially more complex, so we also exclude those. Our estimation sample uses only students living in doubles, which makes up 68% of freshmen. Appendix Figure AF1 shows the size distribution of living situations for dormitories, apartments, and suites.

Forty-five percent of freshmen are in a Living-Learning Program (LLP). An LLP is a section of the dormitories, ranging from a hallway to an entire building, in which students belonging to the same academic program or interested in the same academic course of study live together. The number of LLPs increases over the time period of our analysis, starting with 23 in 2011 and ending with 29 in 2015. The median LLP has around 70 freshmen members, but LLPs range in size from 2 to 526 freshmen in any given year.

A total of 2,111 out of our sample of 21,396 Freshmen are in a program called Freshmen Connection, which delays formal admissions until the Spring semester and requires students to take a fixed set of courses for their first semester. Some members of Freshmen Connection live on campus in the Fall before the Spring of their formal admittance, and those students are often matched with each other because they receive low priority in the housing process. We are able to identify these individuals in the data and treat them as members of another LLP for the purposes of defining the level of randomization.

During this period, students enrolling in the university received a housing questionnaire which asked two questions about smoking status and one each about neatness, study habits, and sleep schedule. After students returned the questionnaire, the university assigned them to rooms based on eight factors: (1.) membership in an LLP, (2.) participation in Freshmen Connection, (3.) gender, and (4.)-(8.) the five questionnaire answers. Students were assigned first to double rooms with one roommate. Once double rooms were filled, remaining students (i.e. the students who were the last to return their housing questionnaire) were assigned to triple and quadruple rooms.[8] Rooms are already pre-designated as

[7]The university's average acceptance rate over the 2011-2015 time period of our data was 46%.

[8]Everyone not in Freshmen Connection who applies before the housing lottery application deadline is treated symmet-

triples and quadruples before the housing assignment begins, which means that the entire triple or quadruple room is filled with students who applied for on-campus housing late.

Since our research design relies upon conditional random assignment, we briefly discuss other possible non-random elements of roommate assignment. A small number of roommates come from requests by people who knew each other before coming to the University of Maryland. A very small second group of freshmen are reassigned by request before the end of the semester (the time at which the data we observe on freshmen pairings is recorded). Roommate separations in the first semester are very rare; the Office of Residential Life rarely approves of within-semester transfers. However, when this does happen, the roommates usually end up living alone for the remainder of the semester and thus will be dropped from our sample. Freshmen are supposed to apply for on-campus housing in May of the year in which they are admitted for entrance in the fall and roommate announcements are made a few weeks prior to the beginning of the semester. There is little time to change roommates before the semester begins and mid-semester changes are exceedingly rare.

In our main specifications we restrict our sample of roommate pairings to those with identical data (demographics, lifestyle questions, LLP, and Freshmen Connection participation). It is very rare that roommate requests or roommate reassignments would have identical answers to the questionnaires, choices of LLPs, and Freshmen Connection participation. We compute, in Table 3, the probability of being exactly matched under different assumptions. Conditioning upon year of entry and gender, the probability of being exactly matched by random chance is 3.7%. In addition, those who did not make the official application deadline are unlikely to be in doubles and extremely unlikely to be in perfectly matched doubles. Thus, in combination, these three sources of non-random selection should reflect a trivial fraction of our sample.

3 Data

Our study necessitates matching confidential data on roommate assignment to voter registration data. In this section, we describe each of the two data sources as well as the matching process between the two. We also discuss how we narrow the sample after matching to the sample that we select for our main estimation.

3.1 Voter Registration Data

We collected data on voter registration and election participation from five states (New York, New Jersey, Pennsylvania, Maryland, and Delaware) as well as the District of Columbia. Together, these states represented the states of application for 90% of the students that attended the university in Fall

rically. Historically, there has been enough space to accommodate all who made the application deadline in doubles.

2020.[9][10] Post-graduation surveys also show these to be the main locations to which students move after graduating from the University of Maryland. One additional state, Virginia, declined to provide data to researchers.[11]

For each state, we collected data on voter registration status, gender, date of registration, party of registration, and voter participation history. From this information, we separate political parties into Democrat, Republican, and Independent. We characterize unaligned voters as Independents.[12] The voter participation data is a set of binary indicators for participation in federal presidential and federal midterm election (both primary and general). In presidential election years, we collect participation for both the earlier presidential primary and the later state primary (for cases in which such primaries are held separately) in addition to participation in the general election. The data cover all such elections from 2008 to 2015. From this data, we construct a binary variable which takes on a value of one if the individual has ever voted in a federal general election and zero otherwise.[13]

3.2 Residential Life Data

The Residential Life data contain all the information used to match roommates except for gender, which is owned by the Office of the Registrar. The data include five lifestyle questions: (1.) whether the individual wakes up early, (2.) whether the individual likes maintaining a neat room, (3.) whether the individual smokes, (4.) whether the individual is tolerant of others smoking, and (5.) whether the individual studies in the room. The data include three additional variables: the year in which the individual was a freshman, whether the individual is in Freshmen Connection, and any Living Learning Program (LLP) participation.[14]

As mentioned, the Residential Life Data does not include gender. In 86% of cases, we are able to recover gender from a match of the individual or their roommate to voter registration data which always contains gender for the six states in our sample. However, in 14% of cases, we do not know gender. This presents a non-trivial data challenge for the empirical analysis. To address the partially missing data we developed novel econometric methods which will be discussed at great length in the

[9] Note that we are using and will continue to use the word state loosely to include Washington, D.C. which is not technically a state.

[10] https://www.irpa.umd.edu/CampusCounts/Enrollments/map_us.pdf

[11] Virginia only provides data to political campaigns and interest groups

[12] There are also a small number of voters registered for third parties such as the Green party or the Libertarian party. There are a sufficiently small number of third party voters that we do not have the precision to estimate the impact of roommate pairings on third party affiliation. We code these voters as "Other" (i.e. not Democrats, Republicans or Independents). Thus our partisan categories are mutually exclusive but not exhaustive.

[13] Note that we do not see when people re-register in another state. Thus, it is possible that our measure of voter turnout is a lower bound of true voter turnout even for those who are registered, to the extent that they then move to another state.

[14] 50 pairs of individuals are in more than one LLP, 49 of these are in different LLPs; we drop all 50 of them from the sample.

methodology section.

3.3 Matching

According to FERPA (the Family Educational Rights and Privacy Act), we are not allowed to see the identified data from Residential Life. Instead, we provided the voter registration data to the university officials, who carried out the matching according to an agreed-upon procedure. They then anonymized the data and provided the resulting deidentified dataset to the researchers. The matching procedure used an exact match on birth year and last name followed by a fuzzy match using Soundex on first name. In the event of multiple matches, the record matching the zip code of initial application was preferred if available. Overall, there were only three multiple matches in which the party information differed. We drop these three individuals.

Since we do not know how matching is carried out between observed roommate pairs who have differences in lifestyle questions, LLP, Freshmen Connection participation, or gender, we restrict our sample to an exactly-matched sample of individuals who are identical in all the data we observe. We make the assumption that all individuals for whom we do not observe gender are matched on gender. Mismatch when we observe gender happens less than 3% of the time. Because the group missing gender makes up 14% of the overall sample, we believe that at most about 0.4% of individuals (9 rooms) in the missing-gender sample would have been mismatched on gender.

4 Methods

We aim to compare outcomes for actual roommate pairs relative to those for roommate pairs that were equally likely to have formed but did not due to random chance. We use two methods for estimation, OLS regression and randomization inference. Whereas OLS is more standard in the economics literature, randomization inference is better suited to our needs for two reasons. First: OLS in this setting uses variation across rooms, whereas our notion of causation derives from the comparison of outcomes in the roommate pairs that were selected relative to the ones that were not selected due to random chance. That comparison is exactly what randomization inference allows us to carry out. The second benefit of randomization inference is that we can empirically estimate p-values without having to invoke the central limit theorem and make asymptotic approximations. In addition to these two estimation methods, we present a novel hybrid method in which we estimate treatment effects using OLS but on the pooled actual and counterfactual data derived from the randomization inference. To our knowledge we are the first paper to combine these two methods.

4.1 Fixed Effects

We begin with OLS estimation, for which we follow the original paper in the economics literature that estimates peer effects using random assignment of roommates (Sacerdote [2001]). In particular we collapse the data to the room level and regress the outcome of one individual on the outcome of the other. However, this is complicated by the lack of a natural ordering of roommates. As a result, estimates may depend upon the way in which the roommates are ordered in the data set. In our main results we randomly categorize one roommate within the pair as the dependent variable roommate and the other roommate as the independent variable roommate. Of course, roommate order is then a source of randomness in our estimates. To take account of this randomness, we bootstrap our standard errors with 10,000 replications for this estimation strategy, simultaneously selecting a bootstrapped sample and a roommate ordering. The standard errors thus simultaneously account for randomness coming from both sampling and from roommate order.

In addition, a simple regression of one outcome on the other without covariates ignores the way in which roommates are randomly assigned conditional on their individual characteristics. The actual data are sorted on covariates over which people may be politically more similar. For example, women tend to register more for the Democratic party whereas men tend to register more for the Republican party. Thus roommates who are sorted on gender, due to the conditional assignment within groups defined in part by gender, will have correlated partisan registration even if there is no social influence on politics. For this reason, in our OLS estimation we additionally control for the covariates used to sort roommates.

We incorporate controls in two ways. We first control for all individual covariates simultaneously. The covariates include fixed effects for answers to each of the lifestyle questions, a fixed effect for participation in Freshmen Connection, a separate fixed effect for each of the living-learning programs (LLPs), and a year of entry fixed effect. In most, though not all, specifications we also control for a gender fixed effect. Second, we fully saturate the model by including a fixed effect for every set of covariate outcomes where we observe at least two rooms full of identical individuals.[15] We call the set of individuals with a certain list of identical covariates a group. In almost all of our regressions, we restrict to the fully-matched sample in which roommates do not differ in covariates used for selection and thus are in the same group. In this case, each room will belong to one and only one group. However, in the few specifications where we expand our sample to include imperfectly matched roommates, rooms will sometimes belong to more than one group and thus be subject to more than one group fixed effect.

[15]If there is only one room of identical individuals, the observation will be effectively removed by the fixed effect and thus will not impact the estimate. As a result, we drop singleton groups.

We interpret our estimates as the conditional correlation in political beliefs of roommates, controlling for determinants of roommate assignment. We could interpret them as the causal effect of the roommate *assignment*, but we do not interpret it as the causal effect of one roommate's views or actions on the other's since we cannot disentangle whether the first roommate influenced the second roommate, whether the second roommate influenced the first roommate, or whether they influenced each other. Our estimation strategy is given by Equation 1:

$$O_{i,g,c} = \alpha + \beta RoomOut_{i,g,c} + \gamma_{g,c} + \epsilon_{i,g,c} \tag{1}$$

where $O_{i,g,c}$ is the outcome variable for the second roommate, i indexes individuals, g is the roommate assignment group and c is the cohort. The variable $RoomOut_{i,g,c}$ is the outcome of the first roommate. This could be either a partisanship indicator or a measure of voter turnout. Finally, $\gamma_{g,c}$ is a roommate group-by-cohort fixed effect.

We present results without fixed effects, with fixed effects, and with fully saturated fixed effects. We also show the model estimated both on the full sample as well as a restricted sample of roommate pairs who are perfectly matched on roommate selection questions. Since gender is missing in 14% of observations, we (1.) show estimates where we condition upon gender and thus drop those rooms where gender is unknown and (2.) show estimates where we do not condition on gender at all. Later in the paper, we develop a method to assign gender to those missing cases in a way that either maximizes or minimizes the size of the treatment effect. We also present OLS results using these two samples (the sample with gender imputed such as to minimize the effect size, and the sample with gender imputed such as to maximize the effect size), which provides lower and upper bounds on our treatment effects.

4.2 Randomization Inference

Our second main approach uses randomization inference. Here we reassign roommate pairings within groups as defined previously. Our goal is to compute what outcomes would have looked like if a different roommate pairing had occurred. If roommates display a systematically greater degree of political similarity in the actual pairings we observe than in the counterfactual pairings that did not occur, then we can reject the hypothesis that there is no social influence on politics among roommates. In most randomization inference designs, some units are treated and others are not and the design reassigns treatment across units. However, in our context all units (rooms) are treated. We instead use randomization inference to create those counterfactual pairings and compute the degree of agreement (homophily) among all rooms that could have been paired in alternative draws of the roommate randomization process.

Our outcome measure is the fraction of time individuals have the same outcome of interest: they are both Democrats, they are both Republicans, they have the same party affiliation, they are both Independents, they both voted, or they both registered to vote. We compute this in our sample of actual roommate pairings. Then we permute roommate pairings within groups (where group is defined by individuals who are freshmen in the same year and match on all characteristics used for assignment of roommates) to form a counterfactual data set. In that data set, we compute the fraction of roommates in which we see agreement for each outcome measure. We then repeat the permutation 10,000 times, each time computing the fraction of roommates in which we see agreemnt on those outcome measures. Finally, we compare the actual to counterfactual outcomes by computing the fraction of time the counterfactual roommate pairings yield a higher degree of agreement than in the actual observed roommate pairing. In this sense, we compute a one-sided p-value. Comparing that p-value to standard levels of statistical significance tells us how confident we can be that the agreement among roommates on political outcomes is not just due to random chance.

We face one insurmountable computational problem: the number of counterfactual assignments is much larger than is computationally feasible. To give a sense of the magnitude: our largest group has 180 individuals, corresponding to 90 rooms. Even with only one group of just 100 people (50 rooms), there are $\frac{100!}{50!2^{50}}$ possible ways to pair the roommates, which is a number with 141 digits.[16] Thus, we use Monte Carlo methods and limit ourselves to 10,000 randomly selected alternative partitions of the data.

4.3 Randomization OLS

We now discuss our third estimation method. Here we combine randomization inference with OLS and call the approach "randomization" OLS. As we noted above, our identification really comes from the actual assignment of roommates relative to those roommates assignments that could have happened but were not realized. Randomization inference is useful for helping us think about design-based identification. In particular, as we noted, all of our roommate pairings are treated and the control roommate pairings are ones that we don't directly observe but that we can construct under a null hypothesis that treatment does not influence outcomes. We thus take all the counterfactual roommate pairings from the randomization inference and pool them with the actual roommate pairing data. We define all the counterfactual allocations as the control observations and all the actual allocations as the

[16] If there are 100 individuals, we can think of ordering them, and then partitioning each successive group of two to form a room. In that case, there are 100! orderings. However, only the assignment of individuals to rooms matter. Any ordering which differs only by the ordering of the rooms should be considered as the same assignment. Thus, for any ordering, there are 50! redundant orderings which have the same assignment to roommates but differ in the orderings of the rooms. In addition, any reordering of roommates within rooms should also be redundant and there are $2^5 0$ of those reorderings.

treated observations. When we do this, we multiply our sample size by a factor of 10,001. We then run our OLS-type regressions but where we define the treatment variable as whether the room is a control or treatment observation. The outcome variable of interest is whether both roommates in the room have the relevant political outcome. We estimate Equation 2:

$$O_{g,c} = \alpha + \beta Actual_{g,c} + \gamma_{g,c} + \epsilon_{g,c} \tag{2}$$

where $O_{g,c}$ takes on a value of 1 if the political variable (such as a dummy variable for Democrat) is equal to 1 for both individuals and 0 otherwise and $Actual_{g,c}$ is the variable identifying whether the room is a treatment or control (counterfactual) room. The interpretation of the β coefficient is the average excess fraction of times where both individuals have a positive political outcome (i.e. both voted or both are Democrat or both are Republican) in the actual relative to in the mean counterfactual. Again, this represents how much more agreement we see in our actual data than we would expect to see by random chance.

5 Results Without Formal Bounds

We now present our first results. We begin by discussing in detail how we select our estimation sample. Then we turn to our OLS analysis in which we consider specifications conditioning on gender (and thus excluding rooms without gender from the sample) as well as not conditioning on gender (and thus keeping rooms which do not have gender in the sample).

5.1 Sample Restrictions

We start with a discussion of the sample selection process. Our initial sample includes 21,396 Freshmen who lived in University of Maryland dorms between 2011 and 2015. Our final sample will be restricted to groups of individuals living in doubles in dorms who are identical according to the information seen by the Office of Residential Life when they carried out roommate assignments. Within each of these groups, we argue that assignment was, with minor exception, random.

From our initial sample of freshmen living on campus, We first consolidate multiple matches by dropping individuals who have multiple voter registration observations matched to the same student ID. These cases occur when multiple voter registration records have the same birthday, an exact match to the last name, and a Soundex match to the first name. Some of those reflect individuals who registered in their home state and then re-registered in Maryland. For most of these multiple matches, we merge them into one individual's record. [17] In the case where two voting records are matched to

[17] We could, in principle, have preferred individuals with an exact first name match when we had multiple Soundex matches but only one exact name match; however, we opted for a conservative approach at the expense of power.

the same individual but the records differ on political party, we drop the student from the sample. If they differ in voting history but not party, we assume it is the same person. The consolidations and eliminations of multiple matches decrease our sample by 1.8%, almost exclusively due to consolidation. We additionally drop the 0.01% of observations who did not answer all the lifestyle questions. Though the student records we received from the Office of Residential Life are exclusively of freshmen, 6.40% of the sample had a roommate who was not a freshman. We drop those individuals from our sample as they reflect either non-random selection, imperfect matching, or both. We drop another 2.2% who were assigned to live in apartments rather than dorms because we expect more self-selection into apartments.

Roommates are assigned to LLPs based upon academic performance, interests, and admission to certain limited-enrollment majors such as engineering and computer science. In 18.20% of cases, individuals are paired with someone admitted to a different LLP, and we drop all those cases. We also drop the 4.72% of observations for which one roommate participated in Freshmen Connection and the other did not. We then restrict our sample to individuals who were assigned to a double, which reduces our sample by an additional 28.01%. Finally, we exclude 2.91% of the sample where the roommates differed on gender followed by an additional 2.9% of the sample with additional data quality issues.[18] Overall, these sample restrictions leave us with 47.55% of our original sample: 10,174 students in 5,087 rooms. We use this sample, which we call the “full” sample, in some analysis. However, our preferred analysis sample is further restricted to individuals who are identical to their roommate in terms of their responses to lifestyle questions, LLP membership, and Freshman Connection participation. This restriction to what we call the “exact match” or “fully matched” sample represents an additional 59.21% decrease in the sample size, leaving us with 4,150 students in 2,075 rooms. If we additionally restrict our sample to those with known gender we lose another 14.3% leaving us with 3,558 individuals. Table 1 below shows how we go from the full sample to our estimation sample.

Sample means for the fully matched sample are shown in Table 2 below. In the Appendix, we also show summary statistics for the full sample (Appendix Table AT2). While a substantial majority of college attendees nationally are women, the University of Maryland has an unusually high fraction male. Out of the fully matched sample with known gender, 56.4% are male and 43.7% female. In the fully matched sample not restricting on gender, 29.1% (31.3% of men, 37.5% of women) are registered Democrat.[19] Registered Republicans represent less than half the sample relative to registered Democrats, at 12.8% (16.3% men, 13% women) of the overall sample. A substantially higher fraction of

[18] This final group of drop observations include those where at least one roommate's listed matriculation term as the previous fall, the previous summer, or the following spring; where the listed matriculation term was blank; or where the individual was listed as being a first-time student living in on-campus housing twice in the data.

[19] The fraction of the overall sample that is registered Democrats is lower than the minimum of the fraction of men and the fraction of women registered Democrat because those without a gender in our sample are not registered to vote and none of them therefore are registered Democrat.

Table 1: Sample Restrictions

	Observations	% Dropped	% Remaining of Original
Start	21396	.	.
Drop duplicate students	21011	1.80	98.20
Drop if missing survey answers	21009	0.01	98.19
Drop if non-freshman roommate	19665	6.40	91.91
Drop if apartment	19233	2.20	89.89
Drop if in different LLPs	15733	18.20	73.53
Drop if different Fresh Connect	14991	4.72	70.06
Drop non-doubles	10792	28.01	50.44
Drop if opposite-gender roommate	10478	2.91	48.97
Drop if flag on data	10174	2.90	47.55
Drop if not exact match	10174	0.00	47.55

the sample is registered Independent, at 23.9% (28.7% men, 26.7% women). Two thirds of the sample is registered to vote, though only 27.1% had voted by 2015 when we collected our data.[20] This is because the one presidential election that individuals starting college between 2011 and 2015 had the chance to vote in by the time of our data collection was the 2012 election and most of our sample was too young to vote at that time. In terms of the lifestyle questions, there isn't as much variation. The vast majority of students are non-smokers, with 98% of the sample object to smoking and less than 0.01% claim that they smoke. 90% study in the room, 82.8% say that they are neat and 76.3% say that they get up early. Only 8% are in freshman connection. The largest components of the variation are gender and membership in an LLP. 56.8% of the sample belong to an LLP and these individuals are further split among 29 total LLPs.[21]

Most of the variables used to carry out roommate matching are not particularly correlated with each other. We show these correlations in Appendix Table AT3. The main exceptions are (1.) being a smoker is negatively correlated with objecting to smoking; this correlation isn't perfect because some non-smokers do tolerate smoking, (2.) being in Freshmen Connection is negatively correlated with being in an LLP, since the former is negatively and the latter positively correlated with prior academic performance, (3.) objecting to smoking is positively correlated with both getting up early and being neat, and (4.) getting up early is positively correlated with being neat.

[20] We compare these numbers to numbers from the 2016 Cooperative Election Survey (CES). CES does not ask if someone is in college. Instead, we use all individuals between 18 and 23 who report having some college as our population. The survey reports that 59% of college attendees are female. This is close to the national average at the time: 58%. The fraction Democrat are slightly higher in the CES than in our sample: 43.7% for women and 34.9% for men. The fraction Republican are also slightly higher: 20.1% for men and 17.4% for women. The fraction of Independents is lower at the University of Maryland. We do note that we are comparing the party ID variable in the CES to actual voter registration in our data. In terms of registration rates, we find 81% for women and 80% for men in the CES. This CES registration rate is higher than in our sample but is self reported and was collected after the beginning of a presidential election year.

[21] There are only 27 LLPs in the exactly matched sample. This is because two small LLPs have no exactly matched roommate pairs.

Table 2: Sample Descriptive Statistics

	Total	Male	Female	Gender Unknown
Democrat	0.291	0.313	0.375	.
	(0.454)	(0.464)	(0.484)	(.)
Republican	0.128	0.163	0.130	.
	(0.334)	(0.370)	(0.337)	(.)
Independent	0.239	0.287	0.267	.
	(0.426)	(0.453)	(0.443)	(.)
Voted	0.271	0.301	0.336	.
	(0.445)	(0.459)	(0.473)	(.)
Registered	0.664	0.770	0.781	.
	(0.472)	(0.421)	(0.414)	(.)
Smoker	0.001	0.001	0.001	0.000
	(0.031)	(0.032)	(0.036)	(0.000)
Objects to Smoking	0.979	0.973	0.986	0.980
	(0.144)	(0.162)	(0.118)	(0.141)
Neat	0.828	0.799	0.850	0.868
	(0.377)	(0.401)	(0.357)	(0.339)
Wakes up Early	0.763	0.769	0.743	0.794
	(0.425)	(0.422)	(0.437)	(0.405)
Studies in the Room	0.900	0.907	0.903	0.868
	(0.300)	(0.290)	(0.296)	(0.339)
Freshman Connection	0.081	0.080	0.057	0.149
	(0.273)	(0.271)	(0.231)	(0.356)
In an LLP	0.568	0.580	0.627	0.372
	(0.495)	(0.494)	(0.484)	(0.484)
N	4150	2008	1550	592

In Appendix Table AT4, we show correlations between political variables and the variables used to assign roommates. For the purposes of this table, we assign a zero rather than a missing to all political variables for those who do not appear in the voter registration data since, being unregistered, they also have not voted and are not registered with a party.[22]. The correlations of political outcomes with lifestyle questions, LLP participation, and gender are primarily small, ranging from about -0.09 to 0.15. Students who are early risers are slightly less likely to be Democrats, while students who prefer a clean dorm are slightly less likely to be registered and are less likely to vote. Being in an LLP is positively correlated with both being registered to vote and with turning out to vote, while being in Freshmen Connection is negatively correlated with both. This implies that academic performance is positively associated with both registration and turnout. Furthermore, LLP membership is positively correlated with being a Democrat and Freshmen Connection participation is negatively correlated with being a Republican. Women for whom gender is known are more likely to be Democrats and less likely to be Republicans, as is the case nationally. This table also presents correlations of political outcomes with upper and lower bound gender variables. Those rows refer to the samples we will construct in which we fill in missing gender based on the coefficient-minimizing and coefficient-maximizing assignment of gender. We will return to those correlations after we discuss in more detail how we assign gender in Section 6.

5.2 Roommate-Level Regression Analysis

The correlations discussed above are at the individual level. However, though treatment (having a roommate with a certain political outcome) occurs at the individual-level, random-assignment is at the room level. The individual-level correlations presented in the prior section are difficult to interpret as casual effects on three grounds:

1. **Full Saturation:** Assignment is not guaranteed to be random except within fully-matched groups

2. **Manski Reflection Problem:** The direction of causality across individuals within the room is not observed; worse yet, influence likely runs in both directions within a roommate pair.

3. **Counterfactual Analysis:** At the room level, at which random assignment occurs, there is no variation in treatment across the data as all rooms are treated. The 'controls' are room assignments that did not occur and are therefore unobserved.

[22] There is a possibility that they are registered but in a state not in our sample. However, the vast majority of US citizens at the University of Maryland come from states for which we have voter registration data.

In Appendix Table AT5, we present our first regression results. In the first two columns, we estimate on the fully matched sample without missing gender. This sample has at least one individual in each room who is registered to vote. In the latter two columns, we show results where we include those without known gender in the sample. In the second and column, we add fixed effects for all grouping variables involved in roommate assignment. In the fourth column, we do the same except we exclude gender from the set of grouping variables. There are 1.779 fully matched rooms with known gender and 2,075 fully matched rooms if we drop gender as a grouping variable. This means that these samples are based upon 3,558 and 4,150 roommates respectively. The sample with 2,075 individuals will be our main sample for most of the paper. We find that coefficients are uniformly much more positive when we include rooms with unknown gender than when we drop such rooms. The estimates for registration are particularly extreme. For example, in the specification without control variables, estimates increase from -0.304 to $+0.137$ from adding rooms without known gender and both estimates are significant at well below a 1% level of confidence. Though adding in controls does change some estimates, the changes are small in comparison with the changes due to adding rooms without known gender.

In Appendix Table AT6, we fully saturate our models on all grouping variables used for roommate assignment. We see this as an improvement in identification over the prior estimation strategies used in papers which utilize conditional random assignment of roommates. Again, in the first two columns, we estimate on the fully matched sample without missing gender and in the last two columns, we estimate on the fully matched sample including rooms with unknown gender. In the second and fourth columns, we estimate effects by group and then aggregate using sample frequency weights. We do this because, as shown in (Solon et al. [2015]), in the presence of both heterogeneous effects by groups and heterogeneity in the variance of the main independent variable, OLS will not correctly produce the average treatment effect. We again find that dropping rooms without known gender leads to substantially more negative effects. Full saturation does alter some of the coefficients though not in a systematic manner. Weighting also changes individual coefficients and some times to a significant degree. However, weighted estimates are not systematically more negative or more positive than unweighted estimates. We will discuss the reasons for the stark differences between estimates dropping rooms with missing gender and estimates including rooms with missing gender below.

To address two of the threeconcerns mentioned above, we now turn to our randomization inference results. These estimates are at the room level, based upon reassignment of roommates within fully-matched groups, and involve creating counterfactual control pairings to compare to the treated factual roommate pairings. The results are presented in Appendix Figures AF3 and AF4. Appendix Figure AF3 includes gender as a variable used in the formation of the groups and thus excludes the rooms with no gender. This is comparable to the "conditioning on gender" columns in the OLS regression tables. Appendix Figure AF4 does not include gender as a variable used in the formation of groups and is comparable to the "not conditioning on gender" columns. The figures show each of the six outcome variables, including same party defined by roommates being both registered Democrats or both registered Republicans. Each of the six figures presents a non-parametric distribution of how many times both roommates share the outcome (both are Democrats, both are Republicans, both are Independents, both are from the same party, both are registered, and both voted). The mean counterfactual is shown in the purple dashed line and the actual is shown in a red dashed line. The figures also show the p-value for the actual (i.e. the fraction of counterfactual assignments that produce an outcome greater than the actual) and the median treatment effect, which we define as the difference between the actual fraction of agreement and the median counterfactual fraction of agreement.

Similar to the OLS results, in the figures based on the sample conditioning on gender, the actual fraction of agreement for a given outcome is very close to the median counterfactual - with the exception of registration, in which the actual fraction of agreement is a large negative outlier. On the other hand, in the figures which do not condition on gender, the actual fraction of agreement for both roommates is uniformly above the meancounterfactual fraction. In fact, in four of the six outcomes, the actual shares are within the top 1%. In other words, the one-sided p-values are below a 0.01 significance level.

What is the intuition behind these results? It is essentially the same as the intuition for the OLS results. When we condition upon gender in order to define the groups within which randomization occurs, we exclude all those missing gender from the counterfactual pairings. And the students missing gender are specifically those who have a zero for all political outcomes since they are not registered. This exclusion falsely increases the fraction of two positive outcomes in the counterfactual distribution and thus lowers the relative position of the actual in the distribution of the counterfactuals. On the other hand, when we don't condition upon gender when defining the groups within which randomization occurs, then many of our counterfactual room assignments are pairing individuals with different genders which, as we have shown in Table AT4, have different associations with political outcomes (for example, female students are more likely to be Democrats). This falsely decreases the fraction of two positive outcomes in the counterfactual distribution, making it seem like the amount of agreement we see in the actual sample is very unlikely to happen by random chance. As a result, when we condition on gender the estimated mean treatment effects are too low, and when we do not condition on gender, the estimated median treatment effects are too high. The impact of missing gender is particularly stark when we look at registration as an outcome because being registered is the most correlated with missing gender.[23] Appendix Figure AF5 highlights those dramatic differences in the estimates for registration to show how important including versus excluding gender is for our analysis.

In summary, the randomization inference results are broadly consistent with the OLS results. We now address the main challenge with our estimation thus far: how to incorporate the potential bias introduced by the missing gender variable. We do this by constructing strict bounds on our estimates as is popular in the partial identification literature (Manski [1990, 2003], Tamer [2010]). We describe our approach in the next section.

6 Methods II: Bounding the Missing Gender Variable's Impact

The discussion in the previous section reveals the consequences of two simple ways to approach the partially missing gender variable. We can either drop all the gender variable from the analysis or restrict our analysis to observations for whom gender is known, but each of those methods brings a bias to the estimates because missing gender is mechanically related to the political outcomes due to the data construction. Dropping observations without gender raises the fraction of double positive outcome rooms in the counterfactuals, since all the rooms being dropped are not double positive rooms. This results in a negative bias. On the other hand, not including the gender variable when defining the randomization groups artificially decreases the fraction of double positive outcome rooms in the counterfactuals by allowing matches across gender to form. Cross-gender matches have lower homophily, which results in a positive bias.

In this portion of the paper, we develop a novel methodology to provide sharp bounds for computing coefficients in a randomization inference design given a missing binary independent variable.

6.1 Univariate Distribution of Roommate Double Positive Outcomes

With K individuals missing gender, there are 2^K different ways to fill in that missing value. One approach to addressing the missing gender would be to run regressions or use randomization inference to compute p-values for all 2^K potential assignments of missing gender and then present both the highest and lowest estimates from among that distribution. This approach would yield sharp bounds on our treatment effects. However, the solution has an exponential asymptotic and thus quickly becomes computationally infeasible. Instead, our method to compute bounds relies upon the fact that we only need to run regressions or randomization inference among the samples that will yield the maximum and minimum treatment effects.

[23] In fact, there would be a one to one relationship between registration and missing gender except that we use the roommate's gender to impute missing gender in cases where one roommate is registered (and therefore has gender data) and the other does not). Therefore, in our analysis sample some people with a gender are not registered, which breaks the perfect correlation between the two variables.

We begin by introducing some terminology. Remember that by a *group*, we mean a set of individuals with identical characteristics from the perspective of the housing authority. This means a set of students with the same answers to the lifestyle questions, in the same LLP or all not in any LLP, with the same Freshmen Connection status, the same gender, and who enter the dorms as a freshman in the same academic year. As noted, we see all these characteristics that the housing authority uses for assignment except, in less than 15% of cases, we do not have information on gender. Rooms in which students are missing gender are mechanically always going to be double negative outcome rooms since we obtain gender from the voter registration files. Thus missing-gender rooms contain students who both are not registered, both have not voted, and both have not registered with any party.

We define a *metagroup* as the set of individuals who share identical responses to the lifestyle questions, identical LLPs, and identical Freshmen Connection status, but may differ on gender. Then by definition each metagroup contains two true groups, one group in which the roommates are male and one group in which the roommates are female.[24] Our objective is to assign a gender to the set of "ungendered" individuals within a given metagroup so that we can place them in one of the two groups within the metagroup. Assigning these ungendered individuals to a given group will not change the actual number of double positive outcomes for that group because all ungendered individuals mechanically do not have a positive outcome for any of our measures. It will, however, reduce the number of counterfactual double positive outcomes within a group, by injecting potential roommate matches in which one of the roommates does not have a positive outcome. That would lower the probability of the counterfactual having more double-positive outcomes than the actual, which would increase the expected treatment effect. The question then becomes in which of the two groups within each metagroup would assigning ungendered students reduce the expected counterfactual matches more (less, for the lower bound). The answer will depend upon the number of male and female individuals in the metagroup (N_{gm}, N_{gf}) and the number of males and females with a positive outcome (k_{gm}, k_{gf}) in each metagroup.

Our objective is to maximize (or minimize) the expected number of counterfactual double positive outcomes by assigning individuals to one of the two groups in each metagroup. Note that, the expectation operator is linear and the number of counterfactual double positive outcomes in one metagroup does not depend on information from other metagroups. As a result, our allocation will be independent across metagroups. This allows us to allocate gender metagroup by metagroup without needing to consider the ordering of the metagroups. Thus if there are u_g individuals to be allocated genders in metagroup g, computationally, we can consider the impact of allocating u_{gf} to the female group

[24] Note that one of the groups within each metagroup could in theory be empty if there are no male or female students who fit the category defined by the lifestyle answers, LLP, Freshman Connection, and year of matriculation.

within metagroup g and $u_g - u_{gf}$ to the male group, where u_{gf} ranges from 0 to u_g. For each potential allocation u_{gf}, we compute the expected number of double-positive-outcome rooms in metagroup g. As shown in Equation 3, this is equivalent to the sum of the expected number of double-positive-outcome rooms in the male group and in the female group. We denote the number of double-positive-outcome rooms as X. This produces a distribution of expected counterfactual double-positive-outcome rooms for each potential allocation u_{gf} of ungendered students.[25] We then compare the expected number of counterfactual double-positive-outcome rooms in order to choose the u_{gf} allocation that either maximizes or minimizes that expectation.

$$max_{u_{gf}} \quad X_{gf}\mathbb{P}\left[X_{gf}|k_{gf}, u_{gf}, N_{gf}\right] + X_{gm}\mathbb{P}\left[X_{gm}|k_{gm}, u_g - u_{gf}, N_{gm}\right] \tag{3}$$

In order to compute the expectation given by Equation 3, we need a formula for the probability function. Dropping the metagroup (g) and group (f or m) subscripts for simplicity, we define N as the number of individuals in a particular group and thus $\frac{N}{2}$ as the number of rooms, k as the number of individuals with a positive outcome (which we will interchangeably refer to as a "success"), and x as the number of double-positive-outcome rooms. In order to describe the derivation of the probability formula, consider different ways of putting individuals in an order, where rooms are given by adjacent individuals so that individuals 1 and 2 belong to the first room, 3 and 4 belong to the second room, etc. The set of all orderings of N people is $N!$. Then we can compute the number of ways to arrange k individuals with positive outcomes into those rooms, given by Equation 4.

$$\underbrace{A_{kN} = \frac{k!(N-k)!}{N!}}_{\text{Ways to arrange k individuals among N}} \tag{4}$$

Now we need to place the double-success rooms within the ordering, as given by Equation 5. Note that any rearrangement of the x double success rooms among the $\frac{N}{2}$ total rooms will add to the number of potential orderings but leave the number of double-success rooms intact.

$$\underbrace{A_{Nx} = \frac{x!(\frac{N}{2}-x)!}{\frac{N}{2}!}}_{\text{Ways to arrange x double success rooms among N/2 total rooms}} \tag{5}$$

After placing the double success rooms within the ordering of all rooms, we need to consider the placement of the rooms which have only one individual with a success. Again, note that any

[25]Note that u_{gf}, the number of ungendered students assigned to the female group within metagroup g, is a sufficient statistic to characterize the allocation because all ungendered students not allocated to the female group will be allocated to the male group.

rearrangement of single-success rooms among the remaining rooms will add to the number of potential orderings but leave the number of double-success rooms unchanged. Since there are $\frac{N}{2} - x$ rooms left after accounting for the double-success rooms and there are $k - 2x$ single-success rooms, the number of permutations of single-success rooms which leave the total number of double-success rooms unaltered is given by Equation 6:

$$\underbrace{A_{Nxk} = \frac{(\frac{N}{2} - k + x)!(k - 2x)!}{(\frac{N}{2} - x)!}}_{\text{Ways to arrange N/2-x rooms among k-2x individuals}} \tag{6}$$

Finally, we need to address the ordering of roommates within a mixed-success room (i.e. in which a pairing of roommate 1 with roommate 2 is a different ordering than a pairing of roommate 2 with roommate 1). Changing the ordering of the roommates does not alter the number of double success rooms, but does add to the potential permutations. This matters when roommate 1 and roommate 2 are mixed, that is, one has a positive outcome and the other does not. Since there are $k - 2x$ success individuals paired with a non-success roommates, there are 2^{k-2x} rearrangements of those roommate pairs which do not alter the number of double success rooms. Equation 7 presents that number of ways to arrange roommate order among mixed roommates:

$$\underbrace{A_{kx} = 2^{k-2x}}_{\text{Ways to arrange roommate orders among mixed roommates}} \tag{7}$$

We now combine those four components to construct the probability mass function for the number of double success rooms x, given the number of individual successes k and the number of individuals N:

$$\mathbb{P}\left[X = x | k, N\right] = \frac{A_{kN} * A_{kx} * A_{Nx}}{A_{kNx}} = \frac{\left(\frac{N}{2}\right)! k! \left(N - k\right)! 2^{k-2x}}{N! x! \left(\frac{N}{2} - k + x\right)! \left(k - 2x\right)!} \tag{8}$$

6.2 Bivariate Distribution of Roommate Double Positive Outcomes

The previous section described the derivation of the probability of a certain number of double-positive outcomes for use in forming the expected degree of agreement among counterfactual pairings. That formula fits the case where a double-positive success can only happen one way. For example, individuals are either registered as Democrat or not, and a double-positive outcome occurs when both roommates are registered Democrats. However, when our outcome of interest is whether roommates are more generally registered with the same party, there is more than one way for a double-positive success to occur. Specifically, roommates can either both be registered Democrats or both be registered Republicans in order for the same party outcome to be a double-success. In this case, we will need to characterize individuals by three outcomes: Democrat, Republican, or other and where we have two types of double success outcomes: double Democrat outcomes and double Republican outcomes. In this section we generalize our previous univariate distribution to a bivariate distribution.

We now differentiate between the number of positive outcomes separately for Republicans and Democrats. Let k_D be the number of individuals with the Democrat outcome equal to one and let k_R be the number of individuals with the Republican outcome equal to one. We also differentiate the number of double-success rooms. Let X_D be the number of double Democrat successes and let X_R be the number of double Republican successes.

6.2.1 Conditional Probabilities

First, note that we want to compute the expected number of double-successes summed across Democrat and Republican successes:

$$\sum_{X_D} \sum_{X_R} (X_D + X_R) \mathbb{P}(X_D, X_R) \tag{9}$$

which we can write as:

$$\sum_{X_D} \sum_{X_R} (X_D + X_R) \mathbb{P}(X_D) \mathbb{P}(X_R | X_D) \tag{10}$$

We break that third term down further by conditioning on the number of mixed Democrat-Republican rooms X_{DR}.

$$\sum_{X_D} \sum_{X_R} (X_D + X_R) \mathbb{P}(X_D) \sum_{X_{DR}} \mathbb{P}(X_R | X_D, X_{DR}) \mathbb{P}(X_{DR} | X_D) \tag{11}$$

Note that we would not need to do this except that there are a third group of individuals who are neither Democrat nor Republican. If there were not a third group of individuals, conditional upon the number of double Democrat and double Republican rooms, the number of mixed Democrat and Republican rooms would be mechanically determined. However, in our case there are both mixed Democrat and Republican rooms as well as neither Democrat nor Republican rooms left over after conditioning on the number of double successes, and we need to consider the ways both of those types of rooms can be permuted.

We know how to compute $\mathbb{P}(X_D)$ from the univariate distribution derived above. Here we turn to deriving $\mathbb{P}(X_{DR} | X_D)$ and $\mathbb{P}(X_R | X_D, X_{DR})$.

6.2.2 Computing $\mathbb{P}(X_{DR}|X_D)$

We begin with $\frac{N}{2}$ rooms and X_D double Democrat rooms, so we must have $\frac{N}{2} - X_D$ rooms left to fill. This corresponds to $N - 2X_D$ individuals. In addition, there is a constraint that the $k_D - 2X_D$ Democrats cannot be in a room together, otherwise there would be more than X_D double Democrat rooms. Thus the $\frac{k_D}{2} - 2X_D$ rooms are already allocated. As a result, the number of permutations becomes the number of ways to order the X_{DR} Democrats matched to Republicans multiplied by the number of ways to order the $k_D - 2X_D - X_{DR}$ Democrats who get matched to non-Republicans multiplied by the allocation of the other $N - 2X_D - 2X_{DR}$ individuals who are not allocated to double-Democrat or mixed-Democrat-Republican rooms. Note that X_{DR} must satisfy the feasibility constraint $X_{DR} \leq \min(k_D - 2X_D, X_R)$ - that is, there cannot be more mixed-Democrat-Republican rooms than there are Republicans or leftover Democrats.

The total number of possible orderings is the number of non-Democrats, $(N - k_D)!$, since we are conditioning on the allocation of Democrats. Since there are X_{DR} Republicans allocated to mixed rooms, the number of ways to allocate the Republicans who get matched to Democrats is given by $\frac{(k_R)!}{(k_R - X_{DR})!}$. The number of ways to allocate the others to Democrats given that there are $N - k_D - k_R$ others is: $\frac{(N-k_D-k_R)!}{(N-k_D-k_R-(k_D-2X_D-X_{DR}))!}$. Finally, there are $(N - 2X_D - 2(k_D - 2X_D))!$ remaining Republicans and others to be paired together.

Putting together these four components (allocations of Republicans to mixed-Democrat rooms, allocations of others to mixed-Democrat rooms, allocations of the remaining, and the total number of orderings), we get:

$$\mathbb{P}(X_{DR}|X_D) = \frac{\frac{(k_R)!}{(k_R-X_{DR})!}\frac{(N-k_D-k_R)!}{(N-k_D-k_R-(k_D-2X_D-X_{DR}))!}(N - 2X_D - 2(k_D - 2X_D))!}{(N - k_D)!} \tag{12}$$

6.2.3 Computing $\mathbb{P}(X_R|X_D, X_{DR})$

Finally, we compute the probability of having X_R double-Republican rooms given that we have X_D double-Democrat rooms and X_{DR} mixed rooms. Since there are now only Republicans and others who have not been allocated and there are no additional restrictions on room allocations, this conditional probability is the same as our univariate distribution, where the number of rooms is $N - 2X_D - 2X_{DR}$ and the number of remaining Republicans is $k_R - X_{DR}$. Thus, the probability of getting X_R double-Republican rooms conditional upon the allocation of double-Democrat and mixed rooms is given by:

$$\mathbb{P}_{univariate}(N - 2X_D - 2X_{DR}, k_R - X_{DR}, X_R) \tag{13}$$

6.2.4 Final Equation

Putting these two components together, we arrive at our final equation for $\mathbb{P}(X_D = x_D, X_R = x_R|k_D, k_R, N)$. Recall that $\mathbb{P}(X_D = x_D, X_R = x_R|k_D, k_R, N) = \mathbb{P}(X_D)\mathbb{P}(X_R|X_D)$. We have shown that the first term $\mathbb{P}(X_D) = \mathbb{P}_{univariate}(N, k_D, X_D)$, and the second term $\mathbb{P}(X_R|X_D)$ is given by:

$$\sum_{X_{DR}} \underbrace{\frac{\frac{(k_R)!}{(k_R - X_{DR})!}\frac{(N-k_D-k_R)!}{(N-k_D-k_R-(k_D-2X_D-X_{DR})!}(N - 2X_D - 2(k_D - 2X_D))!}{(N - k_D)!}}_{\mathbb{P}(X_{DR}|X_D)} \underbrace{\mathbb{P}_{univariate}(N - 2X_D - 2X_{DR}, k_R - X_{DR}, X_R)}_{\mathbb{P}(X_R|X_D, X_{DR})} \tag{14}$$

7 Results with Formal Bounds

We now present our results using the samples in which gender has been assigned in order to achieve the lower and upper bounds of the treatment effects. For each outcome we create a separate sample allocating gender for all individuals missing gender information. We allocate in a way that achieves the lower bound on the expected treatment effect and in a way that achieves the upper bound on the expected treatment effect, as detailed in Section 6. This involves creating twelve different data sets: a lower bound and upper bound version for each of our 6 outcomes. We first walk through a simple computational method for approximating our treatment effects. We then turn to our randomization inference results, and finally discuss the randomization OLS results which are based on the counterfactuals created via randomization inference.

7.1 Back-of-the-Envelope Estimates

In Table 3, we carry out a simple back of the envelope calculation of our likely treatment effects. The rows list our outcomes. The first column shows the fraction of individuals with a positive outcome. The second column shows that fraction squared. If we were randomly pulling one room out of an infinite population, the first column gives the probability of an individual in that room having a positive outcome and the second column would give the probability of a double positive outcome in the room. It would also give the mathematical expectation of the fraction of randomly selecting a room with a double positive outcome from our sample. The third column uses our actual derived probability distribution to compute the expected fraction of rooms with a double positive outcome. In the final column, we show the actual fraction of double positive outcome rooms in our sample. Comparing the fraction of double positive outcome rooms to the expected number should give us a back-of-the-envelope calculation for a treatment effect. However, there is assortative matching in roommate allocation due to the random assignment of roommates within groups, which likely accounts for some of the gap between the expected and the actual fraction of double positive outcome rooms. In order to account for this, we use our derived probability distribution from Section 6 to compute the expected fraction of double-success rooms, as in column 3, except we do that computation within groups defined by the level at which roommates were randomly assigned. We then average the predicted fraction of double-success rooms, weighting by group size. Because the groups are in part defined by gender, we do these calculations in our lower bound (Column 4) and upper bound (Column 5) gender assigned samples. Thus we can interpret the gap between the unconditional expected and the lower bound expected (Column 3 - Column 4) to capture the number of extra double positive outcomes due to assortative allocation and the gap between the actual and the lower bound expected (Column 6 - Column 4) to capture the treatment effect. The same can be said for the upper bound, replacing Column 4 with Column 5.

Table 3: Approximating Treatment Effects

	Probability(%)
Drop unknown genders	0.379
Unconditional	2.249
Conditional on year	1.875
Conditional on gender	0.748
Conditional on gender and year	3.686

We highlight several takeaways from Table 3. First, Column 2 is identical to Column 3. This is because in the case where there is only one group (i.e. the entire sample), our formula for the expected number of double-success outcomes should reduce to the Bernoulli distribution with Bernoulli probability p^2. Second, the lower bound expected fraction of double-positive-outcome rooms (Column 4) is uniformly higher than the unconditional expected number of double-positive-outcome rooms (Column 3). This means there is positive selection (assortative matching of roommates) for all the outcomes we consider. Third, the upper bound expected fraction of double-positive-outcome rooms (Column 5) is always larger than the lower bound (Column 4), as expected since the upper bound was chosen to minimize the treatment effect. Fourth, the actual fraction of double-positive rooms (Column 6) is always higher than both the lower and upper bound expected fractions (Columns 4 and 5) with the exception of the upper bound for registration, suggesting that our treatment effects are almost uniformly positive. Finally, echoing our earlier OLS results, the largest treatment effects calculated by comparing the last two columns are for being registered Democrat and for voting.

7.2 Randomization Inference at the Lower and Upper Bounds

Next, we formalize the estimation of the treatment effects using randomization inference, allowing us also to perform statistical inference. We describe the details of randomization inference in Section 4, but now we run the procedure using our lower and upper bound samples for each outcome. In Figure

1 we present histograms of counterfactual fractions of double-positive-outcome rooms where the lower bound counterfactuals are shown in blue and the upper bound counterfactuals are shown in green. The median of each counterfactual distribution is shown by a vertical dashed line in the relevant color. The vertical dashed red line shows the actual fraction of double-positive-outcome rooms in each sample. For each outcome we report the mean of the treatment effect at the lower bound and at the upper bound. These are given the difference between the actual and the mean of the lower bound (or upper bound) counterfactual distribution. We also report the p-value associated with the lower bound and the upper bound respectively.

Figure 1: Randomization Inference: Gender Bounds

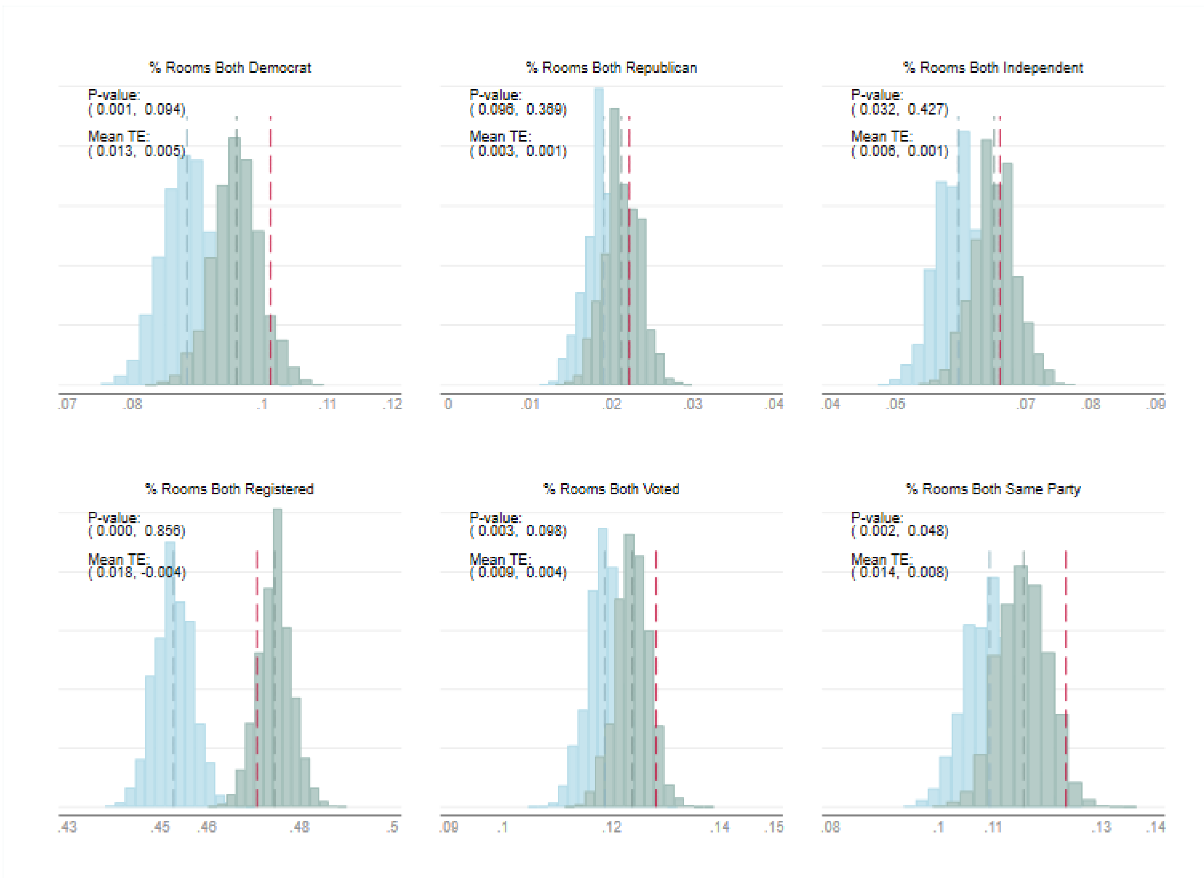

Note: This figure shows the randomization inference results for 10,000 iterations.

Figure 1 shows that the bounding exercise is important. Though the outcomes differ in terms of how far apart their lower bound and upper bound counterfactual distributions are, in all cases the distributions are clearly far from identical. The difference is particularly striking for registration, in which the two distributions only overlap in the tails. The counterfactual lower bound distributions are always first order stochastically dominated by their upper bound counterparts, and the mean treatment effects for the lower bound distributions are always higher than for the upper bound. We expect this since the lower bound minimizes the expected counterfactual, and a smaller counterfactual fraction of double-positive-outcome rooms corresponds to a larger treatment effect.

All of the p-values at the lower bound are below 0.1 and in most cases below 0.05. The largest p-value, for the Republican treatment effect, is 0.096; the next largest is that for Independents at 0.032. However, most p-values are above 0.05 at the upper bound and half are above 0.1. The only outcomes with p-values below a 10% significance level at the upper bound are for both roommates voting (0.098), for both roommates registering as Democrats (0.09) and for both roommates registering as the same party (0.048).

We find that the upper bound treatment effect for Democrats is 1.3; this means that at most there are 1.3 percentage point rooms with two Democrats in them relative to the mean of the counterfactual room assignments. At eh lower bound, that number falls to 0.5 percentage points and the p-value increases from 0.001 to 0.094. For Republicans, effects are 0.003 and 0.001 at the upper and lower bounds respectively. The lower bound p-value is 0.369. The smaller effects for Republicans, however, are largely due to the fact that there are far fewer Republicans on campus. As we shall show later in the paper, these small percentage point effects are similar in magnitude when evaluated in percent rather than percentage point terms. Overall, three variables are significant at below the 10% level even at the lower bound: Democrat, voted and same party. There is a lower bound of 0.4 percentage points greater rooms with two individuals who have voted relative to the mean counterfactual room. Our most significant results are on belonging to the same party. Even at the lower bound, roommates are 0.8 percentage points more likely to be of the same party. These results are significant at slightly below a 5% level of confidence. Interestingly, despite the fact that there are almost as many independents as there are Democrats, estimates for Independents are more similar to those for Republicans than those for Democrats. This potentially suggests that social influence is greater for membership in a political party than for non-membership in any political party.

7.3 OLS Regressions at the Lower Bound and the Upper Bound

We now turn to our OLS regressions based on Equation 1. We consider two samples for each outcome, one sample with missing gender assigned to minimize the treatment effect (Table 5) and one with missing gender assigned to maximize the treatment effect (6). one based on lower bound gender assignment (odd columns) and the other based on upper bound gender assignment (even columns). We show estimates without control variables in the first two columns and estimates with fully saturated controls in Columns 3 and 4. The estimates are presented in Tables 4.

These specifications, while properly correcting for unknown gender, still rely on correlations between roommates rather than differences between roommates pairs that were actually allocated to live together

Table 4: OLS Regressions at the Lower and Upper Bound

	Effect upper bound		Effect lower bound	
	Unweighted	Freqeuncy weighted	Unweighted	Freqeuncy weighted
Democrat	0.0859	0.0529	0.0501	0.0435
	(0.0293)	(0.0186)	(0.0311)	(0.0202)
Republican	0.0380	0.0190	0.0163	0.00315
	(0.0290)	(0.0168)	(0.0305)	(0.0137)
Independent	0.0435	0.0268	0.00406	0.00609
	(0.0303)	(0.0189)	(0.0301)	(0.0176)
Voted	0.0741	0.0297	0.0392	0.0201
	(0.0335)	(0.0160)	(0.0325)	(0.0166)
Registered	0.108	0.0584	-0.0148	-0.0209
	(0.0280)	(0.0173)	(0.0276)	(0.0156)
Same party	0.171	0.106	0.173	0.102
	(0.0385)	(0.0253)	(0.0389)	(0.0242)

Standard errors in parentheses

in contrast to pairs that were not chosen.

We instead turn to a novel method combining the estimation technique of OLS with the research design of randomization inference. In particular, we pool the actual data with the 10,000 counterfactual data sets created via randomization inference into one regression, where the outcome is an indicator for both individuals having a positive outcome and the treatment variable is an indicator for the data coming from the actual roommate pairings. This new data set dramatically increases the sample size to 8.3 million. It is analogous to what we do in our randomization inference specifications in which we compare the mean counterfactual outcomes with the actual outcomes. However, we implement the comparison between actual and mean counterfactual using OLS. As a result, the standard errors are also analytical errors based upon asymptotic approximation. We present the results for the lower bound in Table 5 and for the upper bound in Table 6.

Table 5: Randomization OLS including Counterfactuals - Lower Bound

	Both Democrats	Both Republicans	Both Independent	Both Voted	Both Registered	Same Party
Treatment Effect	0.013^{***}	0.003	0.006^{*}	0.010^{*}	0.018^{**}	0.014^{***}
	(0.004)	(0.002)	(0.004)	(0.005)	(0.008)	(0.005)
Counterfactual	0.088^{***}	0.019^{***}	0.060^{***}	0.119^{***}	0.453^{***}	0.109^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	8304150	8304150	8304150	8304150	8304150	8304150

The treatment effects are identical to the randomization inference results (up to rounding error) except that the standard errors are OLS standard errors. The standard errors incorporate two differences. First, the randomization inference standard errors are exact and do not rely upon asymptotic approximation based upon the Central Limit Theorem, and second, with randomization inference we

Table 6: Randomization OLS including Counterfactuals - Upper Bound

	Both Democrats	Both Republicans	Both Independent	Both Voted	Both Registered	Same Party
Treatment Effect	0.005	0.001	0.001	0.004	-0.004	0.008
	(0.005)	(0.002)	(0.004)	(0.005)	(0.008)	(0.005)
Counterfactual	0.096***	0.021***	0.065***	0.124***	0.475***	0.116***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	8304150	8304150	8304150	8304150	8304150	8304150

can use the non-parametric distribution of the errors in lieu of standardizing them.

In Table 7, we express the treatment effects as a fraction of the control mean. For example, there are substantially fewer Republicans than Democrats at the University of Maryland at College Park. As it turns out, while the effect size at the upper bound is approximately 4 times larger for Democrats than for Republicans, the effects in percent terms are substantially closer in size. In fact, the Republican percent effect is slightly higher in percent terms.

Table 7: OLS Treatment Effects Using Counterfactuals as Controls: Conversion to Percent Effects

	Upper Bound		Lower Bound	
	Percentage Point	Percent	Percentage Point	Percent
Democrat	0.013	0.129	0.005	0.052
Republican	0.003	0.152	0.001	0.043
Independent	0.006	0.095	0.001	0.015
Registered	0.018	0.038	-0.004	-0.008
Voted	0.010	0.075	0.004	0.034
Same Party (R, D)	0.014	0.113	0.008	0.063

7.4 Effects Over Time

Finally, we examine how the treatment effects change over time. Our sample is comprised of five cohorts of freshmen, the first entering college in 2011 and the last in 2015, the same year in which we collected the data. This means that individuals in our first cohort were treated (by being roommates) five years prior to our collection of the voter registration data, whereas for individuals in the most recent cohort were still living with each other and had only been roommates for one semester at the time of data collection. Exploring the effect of treatment for these different cohorts may give us some information about how the treatment effect evolves over time, though we cannot distinguish between cohort-specific effects and dynamic time effects. In Figure 2 we present the randomization inference treatment effects estimated separately by the amount of time since the individuals were roommates in their freshman year. The solid lines show the median treatment effect (the difference between the actual and the median counterfactual fraction of double-outcome rooms) which is analogous to the main treatment effect presented in our other randomization inference figures. In addition, we plot with dotted lines the treatment effects associated with the 2.5^{th} and 97.5^{th} percentiles of the counterfactual distribution. We include estimated treatment effects from both the lower bound and upper bound samples.

We take our results as suggestive in part due to the inability to distinguish between cohort and dynamic effects and in part due to the wide confidence bands resulting from small sample sizes. We see somewhat higher estimated treatment effects among the oldest cohorts for both roommates being registered as Independent and both roommates voting. However, particularly for voting, the likelihood that these estimates reflect a cohort rather than dynamic effect is high. The older cohorts are the only ones who had the opportunity to vote in a presidential election and thus it is possible that those cohorts were differentially impacted by their roommate's political activity. It is also possible that marginal registrants for a presidential election come from individuals with weak attachment to parties. Thus, cohort effects related to the timing of the presidential election could explain the limited heterogeneity we *do* see across cohorts. Interestingly, we do not see any evidence of a decline in effects as more time elapses since the individuals were treated, which provides suggestive evidence that social effects are long-lasting.

Since the annual estimates are statistically weak due to sample size, we additionally estimate a linear time trend. We then perform 10,000 iterations of randomization inference where, for each sample, we estimate a linear time trend. Appendix Figure AF7 presents the randomization inference results for the linear time trend estimates. We denote the time trends based upon actual roommate assignments by a red dashed line. For all six outcomes, these estimates are negative. This is unsurprising as cohorts who started in later years are younger and thus less likely to be registered to vote, less likely to have voted, and less likely to be affiliated with any particular party (or to be unaffiliated yet registered). Nonetheless, we see that for all outcomes except for registration, p-values are well above 0.5 for both the upper and lower bound estimates. Even for registration, the lower bound p-value is close to 0.2 and the upper bound close to 0.95. In sum, we see little statistical evidence for the time trend in homophily being different for the treated relative to the counterfactuals who were not treated. As a result, our results are consistent with a persistent effect over the span of five years.

8 Conclusion

In this paper, we have presented causal evidence that social connections influence both political partisanship and political participation. We do this by introducing a randomization inference approach to an empirical design that leverages conditionally random roommate assignment. We also introduce new methods for estimating sharp lower and upper bounds on treatment effects in the presence of a discrete missing variable. In addition, we introduce another methodological innovation by implementing

Figure 2: Roommate Effect By Years Since Enrollment

Note: This figure shows the randomization inference results for 10,000 iterations.

a randomization inference type design using OLS.

Substantively, we find evidence that people allocated to the same room are more likely to both be registered as Democrats and more likely to both have voted. Even the lower bound of these effects, incorporating uncertainty in gender in 14% of observations, are statistically significant at below a 10% level. We find highly an increase in probability of being registered for the same party of between 6.3% and 11.3%. These results are statistically significant at below a 5% level even at the lower bound. We see no evidence that effects fade over time though our results over time are statistically weak. Future research should confirm the importance of social connections for a broader swathe of the population than those in college. In addition, the results of this paper bring up questions about whether social effects play a role in ideology formation in addition to partisanship, and how important this shaping of political preferences by peers is for affective polarization. As the United States becomes more polarized along partisan lines and as concern over the determinants of polarization increases (Levy [2021]), the questions raised by this paper can only become more important.

References

Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. *The nature of prejudice.* Addison-wesley Reading, MA, 1954.

Peter Bergman. The risks and benefits of school integration for participating students: Evidence from a randomized desegregation program. *IZA Discussion Paper*, 2018.

Stephen B Billings, Eric Chyn, and Kareem Haggag. The long-run effects of school racial diversity on political identity. Technical report, National Bureau of Economic Research, 2020.

Bill Bishop. *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart*. Houghton Mifflin Harcourt, 2009.

Duncan Black. On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34, 1948.

Johanne Boisjoly, Greg J Duncan, Michael Kremer, Dan M Levy, and Jacque Eccles. Empathy or antipathy? the impact of diversity. *American Economic Review*, 96(5):1890–1905, 2006.

Samuel Bowles and Herbert Gintis. *Schooling in capitalist America: Educational reform and the contradictions of economic life*. Haymarket Books, 2011.

Alvaro Calderon, Vasiliki Fouka, and Marco Tabellini. Racial diversity and racial policy preferences: the great migration and civil rights. *The Review of Economic Studies*, 90(1):165–200, 2023.

Davide Cantoni, Yuyu Chen, David Y Yang, Noam Yuchtman, and Y Jane Zhang. Curriculum and ideology. *Journal of political economy*, 125(2):338–392, 2017.

Scott E Carrell, Bruce I Sacerdote, and James E West. From natural variation to optimal policy? the importance of endogenous peer group formation. *Econometrica*, 81(3):855–882, 2013.

Thomas Cornelissen, Christian Dustmann, and Uta Schönberg. Peer effects in the workplace. *American Economic Review*, 107(2):425–456, 2017.

Stefano DellaVigna and Ethan Kaplan. The fox news effect: Media bias and voting. *The Quarterly Journal of Economics*, 122(3):1187–1234, 2007.

Matthew Gentzkow. Television and voter turnout. *The Quarterly Journal of Economics*, 121(3):931–972, 2006.

Matthew Gentzkow and Jesse M Shapiro. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839, 2011.

Antonio Gramsci. *Prison Notebooks Volume 2*, volume 2. Columbia University Press, 2011.

Edward S Herman and Noam Chomsky. *Manufacturing consent: The political economy of the mass media*. Random House, 2010.

Ethan Kaplan, Jörg L Spenkuch, and Cody Tuttle. School desegregation and political preferences: Long-run evidence from kentucky. *Working Paper*, 2019.

David S Lee. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies*, 76(3):1071–1102, 2009.

Ro'ee Levy. Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3):831–70, 2021.

Charles F Manski. Nonparametric bounds on treatment effects. *The American Economic Review*, 80 (2):319–323, 1990.

Charles F Manski. *Partial identification of probability distributions*, volume 5. Springer, 2003.

John Marshall. The anti-democrat diploma: How high school education decreases support for the democratic party. *American Journal of Political Science*, 63(1):67–83, 2019.

Gregory J Martin and Ali Yurukoglu. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–99, 2017.

Karl Marx. *The eighteenth brumaire of Louis Bonaparte*. International Publishers, 1926.

Karl Marx. *Grundrisse: Foundations of the critique of political economy*. Penguin UK, 2005.

Alexandre Mas and Enrico Moretti. Peers at work. *American Economic Review*, 99(1):112–145, 2009.

Lilliana Mason. *Uncivil Agreement: How Politics Became Our Identity*. University of Chicago Press, 2018.

Markus Prior. Media and political polarization. *Annual Review of Political Science*, 16:101–127, 2013.

Bruce Sacerdote. Peer effects with random assignment: Results for dartmouth roommates. *The Quarterly Journal of Economics*, 116(2):681–704, 2001.

Kelly Shue. Executive networks and firm policies: Evidence from the random assignment of mba peers. *The Review of Financial Studies*, 26(6):1401–1442, 2013.

Gary Solon, Steven J Haider, and Jeffrey M Wooldridge. What are we weighting for? *Journal of Human Resources*, 50(2):301–316, 2015.

Logan Strother, Spencer Piston, Ezra Golberstein, Sarah E Gollust, and Daniel Eisenberg. College roommates have a modest but significant influence on each other's political ideology. *Proceedings of the National Academy of Sciences*, 118(2), 2021.

Elie Tamer. Partial identification in econometrics. *Annu. Rev. Econ.*, 2(1):167–195, 2010.

9 Appendix

Table AT1: Matching Versions

	Probability(%)
Drop unknown genders	0.379
Unconditional	2.249
Conditional on year	1.875
Conditional on gender	0.748
Conditional on gender and year	3.686

Table AT2: Summary Statistics in Full Sample

	Total	Male	Female	Gender Unknown
Democrat	0.286	0.291	0.383	.
	(0.452)	(0.454)	(0.486)	(.)
Republican	0.120	0.156	0.121	.
	(0.325)	(0.363)	(0.326)	(.)
Independent	0.237	0.291	0.247	.
	(0.425)	(0.454)	(0.431)	(.)
Voted	0.253	0.265	0.329	.
	(0.435)	(0.441)	(0.470)	(.)
Registered	0.650	0.747	0.757	.
	(0.477)	(0.435)	(0.429)	(.)
Smoker	0.007	0.008	0.004	0.009
	(0.081)	(0.087)	(0.064)	(0.097)
Objects to Smoking	0.909	0.895	0.932	0.899
	(0.288)	(0.307)	(0.252)	(0.302)
Neat	0.770	0.737	0.800	0.795
	(0.421)	(0.440)	(0.400)	(0.403)
Wakes up Early	0.652	0.646	0.651	0.666
	(0.476)	(0.478)	(0.477)	(0.472)
Studies in the Room	0.753	0.755	0.760	0.737
	(0.431)	(0.430)	(0.427)	(0.441)
Freshman Connection	0.100	0.084	0.099	0.148
	(0.301)	(0.277)	(0.299)	(0.355)
In an LLP	0.496	0.514	0.533	0.374
	(0.500)	(0.500)	(0.499)	(0.484)
N	21011	9689	7602	3720

Table AT3: Correlations Between Grouping Variables

	Smoker	Objects to Smoking	Neat	Wakes up Early	Studies in the Room	Freshman Connection	In an LLP	Female	Gender Known
Smoker	1.000	-0.227	0.016	-0.059	0.011	-0.009	-0.041	0.004	0.013
Objects to Smoking	-0.227	1.000	0.063	0.207	0.125	0.040	0.118	0.044	-0.003
Neat	0.016	0.063	1.000	0.212	0.090	0.082	-0.045	0.067	-0.044
Wakes up Early	-0.059	0.207	0.212	1.000	0.176	0.083	-0.025	-0.030	-0.030
Studies in the Room	0.011	0.125	0.090	0.176	1.000	0.066	-0.016	-0.007	0.044
Freshman Connection	-0.009	0.040	0.082	0.083	0.066	1.000	-0.335	-0.045	-0.101
In an LLP	-0.041	0.118	-0.045	-0.025	-0.016	-0.335	1.000	0.048	0.161
Female	0.004	0.044	0.067	-0.030	-0.007	-0.045	0.048	1.000	.
Gender Known	0.013	-0.003	-0.044	-0.030	0.044	-0.101	0.161	.	1.000

Sample is exactly matched on all grouping variables except for gender. N = 4150 for all variables except Female, for which N = 3558.

Table AT4: Correlations of Outcomes with Grouping Variables

	Democrat	Republican	Independent	Registered	Voted
Smoker	0.014	-0.012	-0.017	-0.011	-0.001
Object to smoking	0.002	0.001	0.000	0.005	-0.016
Neat	-0.031	0.023	-0.046	-0.055	-0.101
Wakes up early	-0.065	0.061	0.004	-0.017	-0.047
Studies in the room	-0.025	0.029	0.017	0.007	-0.012
Freshman connection	-0.060	0.003	-0.038	-0.088	-0.143
In an LLP	0.080	-0.010	0.068	0.131	0.083
Female	0.065	-0.046	-0.022	0.013	0.037
Female(LB)	0.028	-0.062	-0.043	-0.026	-0.002
Female(UB)	0.003	-0.091	-0.094	-0.152	-0.035
Gender known	0.262	0.156	0.228	0.574	0.249

Sample is exactly matched on all grouping variables except for gender. The definition of Female (LB) and Female (UB) varies by column. For each outcome, gender is imputed such that the RI p-value is minimized (LB) or maximized (UB). N = 4150 for all variables except Female, for which N = 3558.

Table AT5: OLS Estimates of Roommate Correlations with Linear Fixed Effects

	Conditioning on gender		Not conditioning on gender	
	No controls	With controls	No controls	With controls
Democrat	0.012	-0.004	0.078***	0.055**
	(0.024)	(0.025)	(0.024)	(0.024)
Republican	0.029	0.006	0.052**	0.032
	(0.024)	(0.024)	(0.022)	(0.020)
Independent	-0.003	-0.006	0.049*	0.041
	(0.027)	(0.028)	(0.025)	(0.025)
Voted	0.227***	-0.003	0.275***	0.057***
	(0.025)	(0.025)	(0.025)	(0.022)
Registered	-0.304***	-0.335***	0.137***	0.078***
	(0.014)	(0.016)	(0.023)	(0.024)
Observations	1779	1779	2075	2075

Regressions are of the outcome variable for one roommate on the outcome variable of the other roommate. Sample is exactly matched on all grouping variables in Columns 1 and 2; those without known gender are dropped from estimation. Sample is exactly matched on all grouping variables excluding gender in Columns 3 and 4; those without known gender are including in estimation. The first and third columns do not contain any controls. The second column controls linearly for all grouping variables including gender. The fourth column control linearly for all grouping variables except for gender. Standard errors are bootstrapped at the room level. Errors incorporate both randomness from sample composition and from roommate order within pair.

Table AT6: OLS Estimates of Roommate Correlations with Fully Saturated Fixed Effects

	Conditioning on gender		Not conditioning on gender	
	Matching	Matching, weighted	Matching	Matching, weighted
Democrat	0.004	0.004	0.058*	0.088**
	(0.039)	(0.043)	(0.033)	(0.042)
Republican	0.001	-0.002	0.031	0.029
	(0.037)	(0.040)	(0.027)	(0.030)
Independent	-0.020	-0.057	0.034	0.006
	(0.042)	(0.058)	(0.031)	(0.036)
Voted	0.001	-0.006	0.082***	0.069**
	(0.038)	(0.048)	(0.028)	(0.031)
Registered	-0.293***	-0.304***	0.093***	0.102**
	(0.025)	(0.032)	(0.029)	(0.048)
Observations	1779	1779	2075	2075

Regressions are of the outcome variable for one roommate on the outcome variable of the other roommate. Sample is exactly matched on all grouping variables in Columns 1 and 2; those without known gender are dropped from estimation. Sample is exactly matched on all grouping variables excluding gender in Columns 3 and 4; those without known gender are including in estimation. In the second and fourth columns, estimates are computed at the group level and then aggregated according to sample frequency weights to account for aggregation bias in the presence of heterogeneity in treatment variance across groups.Standard errors are bootstrapped at the room level. Errors incorporate both randomness from sample composition and from roommate order within pair.

Figure AF1: Distribution of Room Sizes and Types

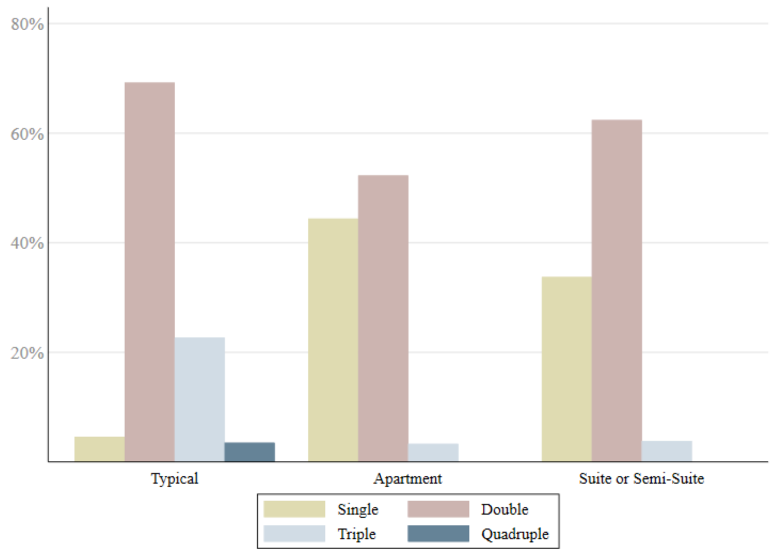

Figure AF2: Number of Votes Vary by Election

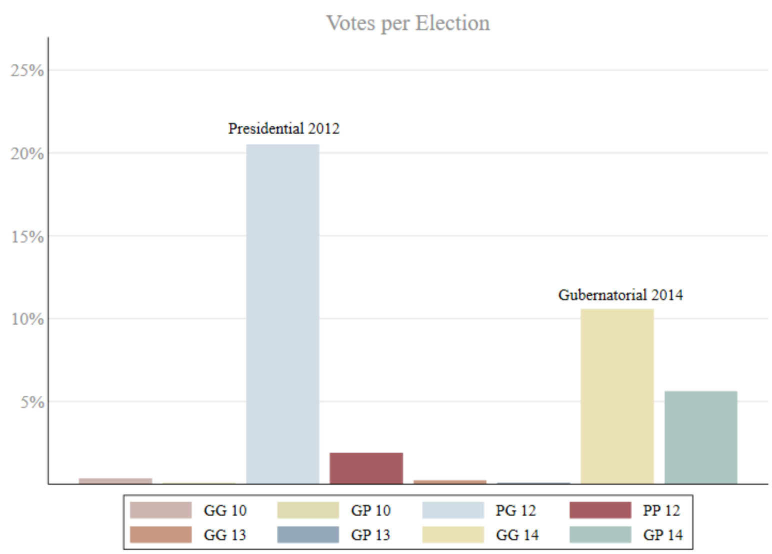

Figure AF3: Randomization Inference: Gender Groups

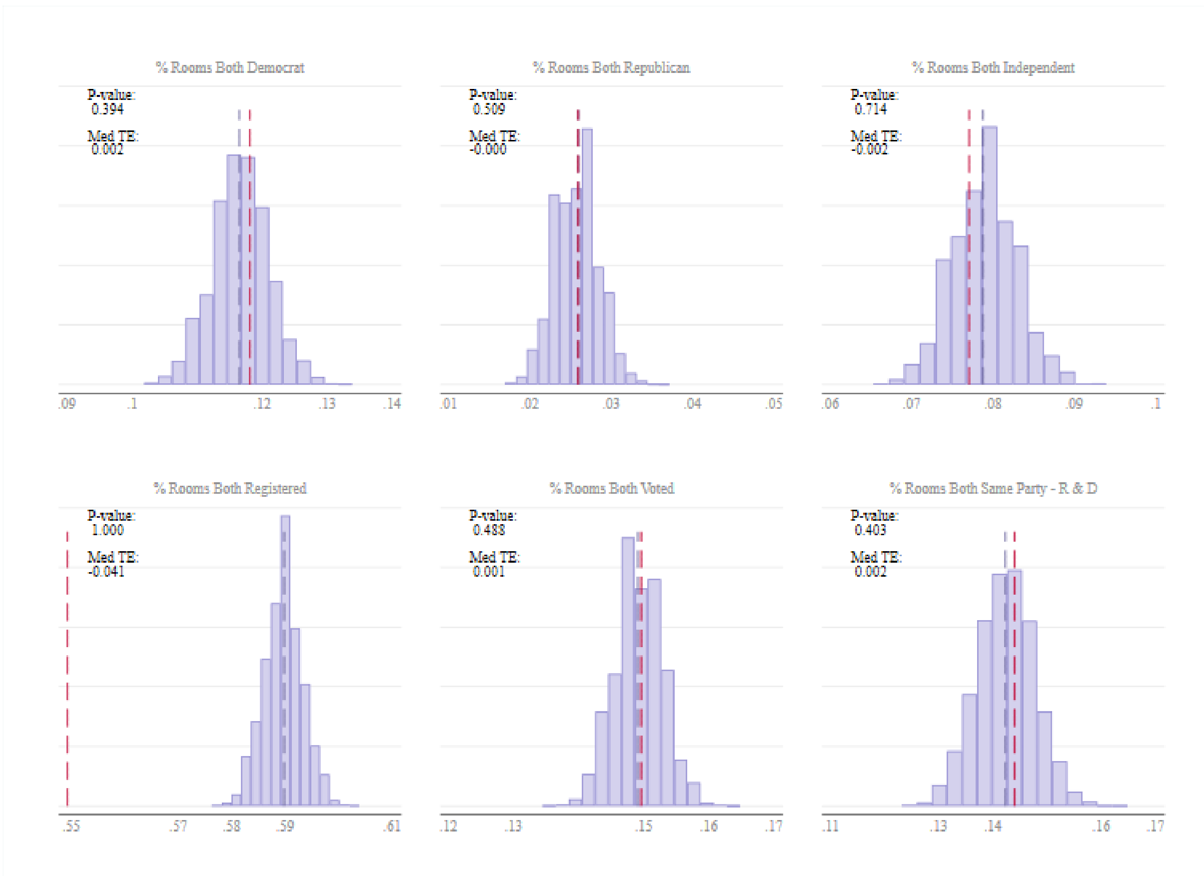

Note: This figure shows the randomization inference results for 10,000 iterations. The sample includes grouping by all LLPs and gender, in addition to year and answers to the lifestyle questions. There are 3688 individuals in this sample.

Figure AF4: Randomization Inference: No Gender Groups

Note: This figure shows the randomization inference results for 10,000 iterations. The sample includes grouping by all LLPs in addition to year and answers to the lifestyle questions. This sample does not group by gender. There are 4286 individuals in this sample.

Figure AF5: Registration Status with and without Grouping on Gender

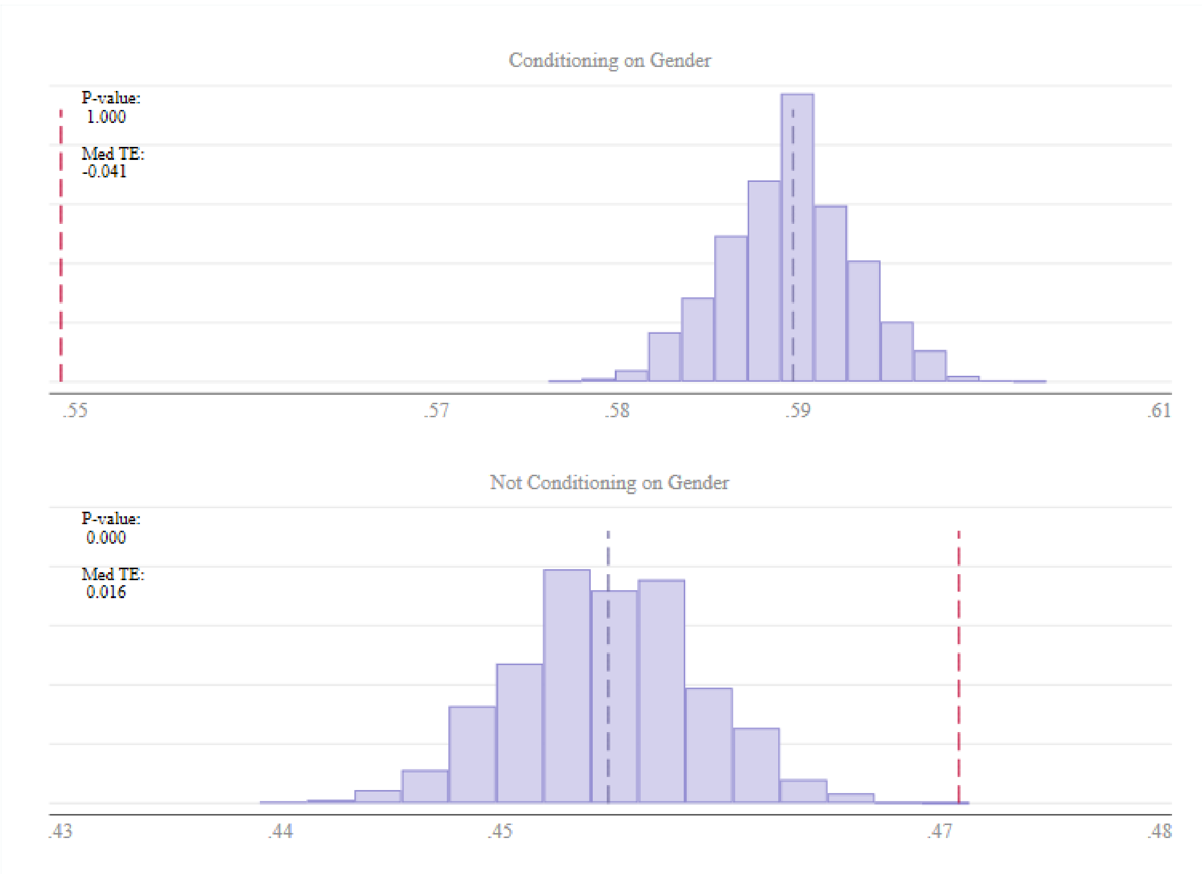

Note: This figure shows the randomization inference results for 10,000 iterations.

Figure AF6: Bootstrapped OLS Estimates

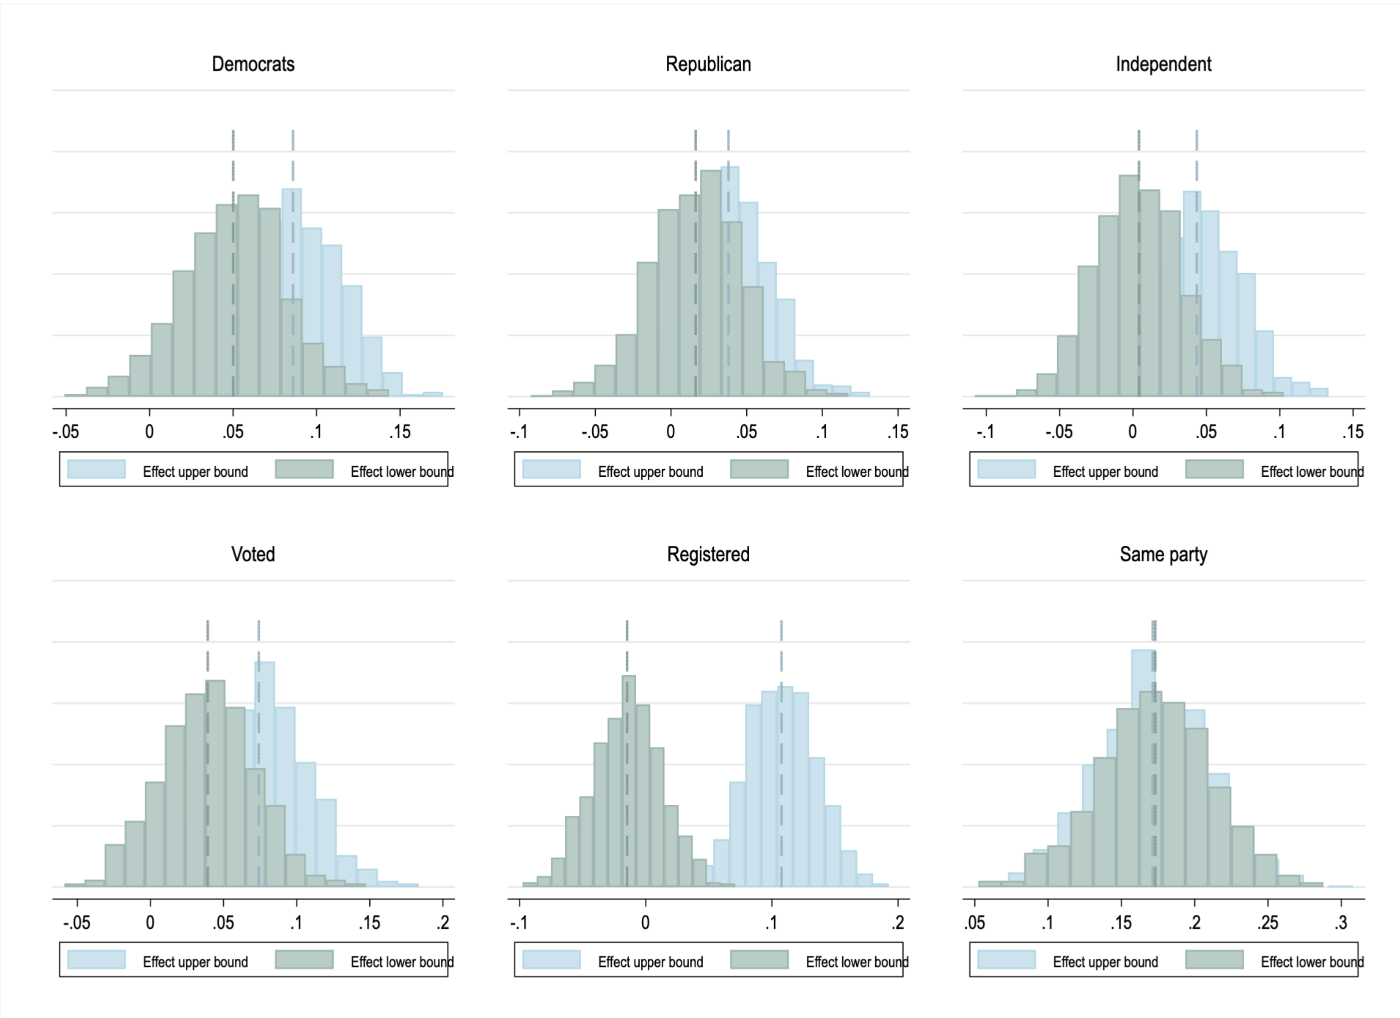

Note: This figure shows the randomization inference results for 10,000 iterations. The bootstrap both randomly draws a sample with replacement and randomly assigns a roommate order within the room.

Figure AF7: Randomization Inference: Linear Time Trend

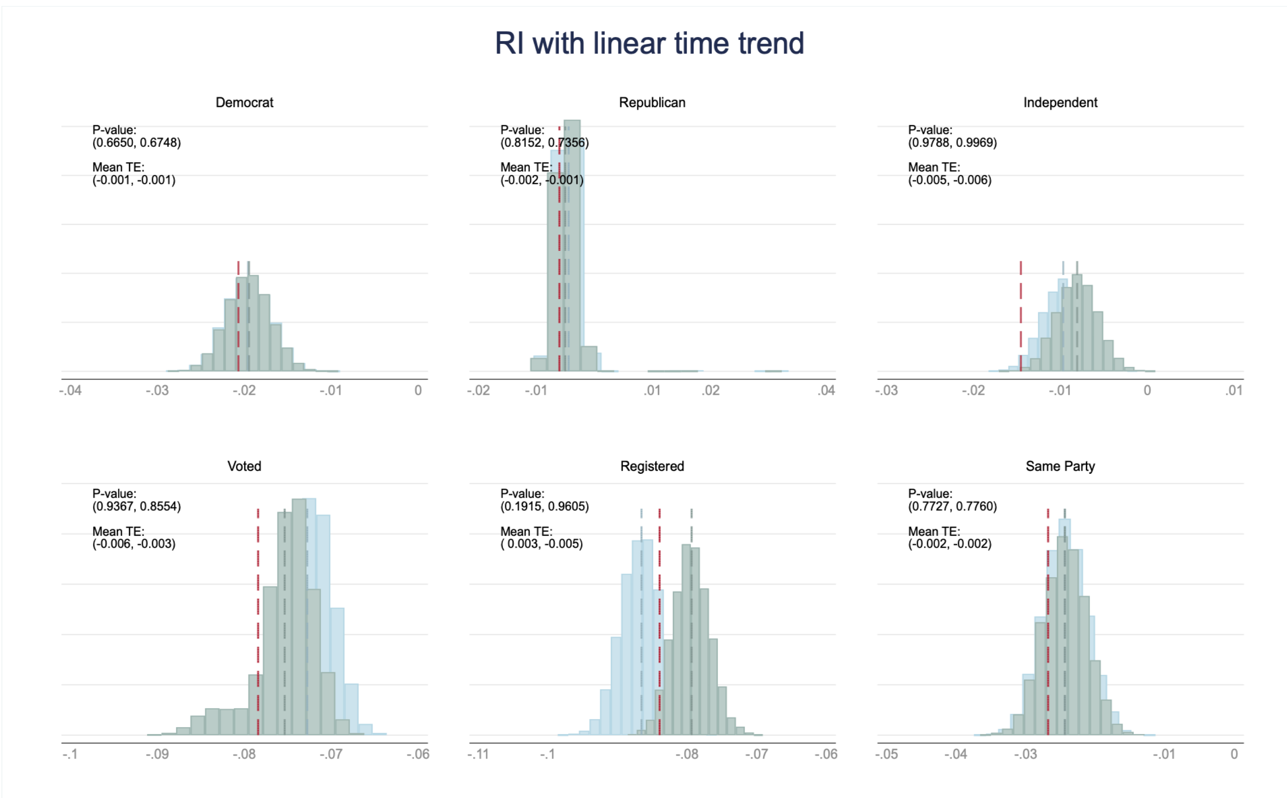

Note: This figure shows the randomization inference results for 10,000 iterations.